From: 赵运磊 <ylzhao@fudan.edu.cn>
Sent: Monday, April 22, 2019 3:09 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: ROUND 2 OFFICIAL COMMENT: Frodo: Comparison between CCA-secure AKCN-LWE and FrodoKEM
Attachments: AKCN-LWE.pdf

Dear FrodoKEM Team:
Recently, we tested the performance of our AKCN-LWE protocol (available from https://arxiv.org/abs/1611.06150). Based on the same discrete distributions and the same CCA transformation used by FrodoKEM, the performance and comparison are briefly summarized in the following table. From the table, the generality of AKCN-LWE allows for more flexible parameter selection: smaller bandwidth (about 13% reduction for Frodo-976), and seemingly better balance between error probability and pq-security.

Best regards
Yours sincerely
Yunlei

	n	q	m	g	t	dist	ciphertext	err.	$\|K\|$	C	Q
Frodo-640	640	2^{15}	2^2	2^{15}	0	$\chi_{\text{Frodo-640}}$	9720	$2^{-148.8}$	128	144	103
AKCN-640	640	2^{15}	2^2	2^{10}	1	$\chi_{\text{Frodo-640}}$	9040	$2^{-132.7}$	128	144	103
Frodo-976	976	2^{16}	2^3	2^{16}	0	$\chi_{\text{Frodo-976}}$	15744	$2^{-199.6}$	192	209	150
AKCN-976	976	2^{16}	2^3	2^8	2	$\chi_{\text{Frodo-976}}$	13728	$2^{-164.1}$	192	209	150

Brief comparison between CCA-secure AKCN-LWE and FrodoKEM.

The ciphertext size is the total length of bytes sent by Bob. For AKCN-640, its ciphertext is 7% smaller than Frodo-640. While its error probability is larger than Frodo-640, it's still under 2^{-130} that is sufficiently smaller for 103-bit pq-security. For AKCN-976, its ciphertext is 12.8% smaller than Frodo-976, and its error probability is still under 2^{-160} that is sufficiently smaller for 150-bit pq-security.

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Friday, May 24, 2019 4:33 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: ROUND 2 OFFICIAL COMMENT: Frodo
Attachments: signature.asc

Page 30 of the Frodo submission claims that Theorem 5.1 of the submission is a tight ROM reduction showing that "FrodoKEM is an IND-CCA-secure KEM under the assumption that FrodoPKE is an OW-CPA-secure public-key encryption scheme".

The theorem statement is somewhat more general, and claims that "Adv^ind-cca" of an attack against the KEM in Definition 2.19 is at most

* 3 times "Adv^ow-cpa" of an attack against the underlying PKE, plus
* the number q of hash queries times the decryption failure rate, plus
* 3q+1 divided by the size of the message space.

There is then a claim that "The proof of Theorem 5.1 is analogous to the proofs of Theorems 3.2 and 3.4 of Hofheinz, Hövelmanns, and Kiltz (HHK) [63]". This is followed by a few plausible comments about tweaks to the hashing; these tweaks aren't relevant to what I'm going to say here.

I've looked for the claimed proof of Frodo Theorem 5.1 in some obvious spots and haven't found it. Has this theorem in fact been proven?

It's also not clear to me what level of review has taken place regarding this proof, or the other Frodo proofs. In general, does the Frodo team vouch for the correctness of the proofs used in the Frodo submission?
And, to be clear, this includes proofs of every "Theorem" claimed in the submission?

If "ow-cpa" were replaced by "ind-cpa" then it would be clear (modulo various smaller details that I haven't checked---I don't vouch for correctness here!) how Frodo Theorem 5.1 relates to HHK Theorems 3.2 and 3.4. But switching from OW-CPA to IND-CPA is a huge change:

* The HHK paper presents a way to construct ROM IND-CPA tightly from OW-CPA---but this is at the expense of much larger ciphertexts. Frodo doesn't use this.

* The notion that lattice problems have been "well studied" consists primarily of pointers to the literature for algorithms to attack various _search_ problems such as SVP. Even if one leaps all the way to assuming hardness of search-LWE for the relevant parameters, there isn't a _tight_ proof of hardness of decision-LWE.

* OW-CPA has an easy proof of robustness against moderate changes in distributions (as measured by "Renyi divergence"; for some simple examples see https://ntruprime.cr.yp.to/divergence-20180430.pdf). Frodo seems to rely critically on this in its choice of error distribution (see Section 5.1.3), but why should one believe that

an assumption different from OW-CPA has the same robustness? I haven't found any proofs in the Frodo submission on this point.

Is there a way to rescue the Frodo Theorem 5.1 claim of a proof that starts from OW-CPA? There's HHK Theorem 3.1, which starts from OW-CPA, but this theorem isn't tight. There's

https://eprint.iacr.org/2018/526

which starts from OW-CPA and _is_ tight, but this needs the underlying PKE to be deterministic. The tightness gap here is exactly in the FO derandomization step, which produces a deterministic PKE by choosing the randomness in encryption as a hash of the message. Does someone have a way to avoid this gap?

---Dan

From: daniel.apon <daniel.apon@nist.gov>
Sent: Friday, May 24, 2019 2:03 PM
To: pqc-forum
Cc: pqc-comments; djb@cr.yp.to
Subject: Re: ROUND 2 OFFICIAL COMMENT: Frodo

Hi Dan,

"I've looked for the claimed proof of Frodo Theorem 5.1 in some obvious spots and haven't found it. Has this theorem in fact been proven?"

Sure, let's verify it together.

Note that FrodoKEM's specification document's Theorem 5.1 is a generic theorem, in that doesn't refer to any particular PKE (in particular, it doesn't refer to FrodoKEM's PKE).

Theorem 5.1 in the FrodoKEM spec is obtained by beginning with an IND-CPA PKE, applying HHK's -
- https://eprint.iacr.org/2017/604.pdf -- Theorem 3.2 (HHK, page 12) with q_V set to 0 to obtain a OW-PCA PKE via Transformation T (page 10) -- not OW-PCVA, since q_V is 0; and then applying HHK's Theorem 3.4 (HHK, page 16) via transform U^\notbot (page 15; it's U^\bot with implicit rejection) to obtain an IND-CCA KEM.

The 'change' between HHK and FrodoKEM's Theorem 5.1 is that this transformation is done 'in one step,' which basically just means using a single hash function with longer output -- but I think we agree that that tweak is inconsequential to security.

If I read what you're saying correctly, then we seem to agree that if FrodoKEM's spec is changed to say "IND-CPA" instead of "OW-CPA," that it would be more clear. To me, it looks like the "summary" paragraph titled "Security reductions" in Section 5.1 of FrodoKEM's spec (including Footnote 5) is mis-explained. Similarly, it looks like Theorem 5.1 of FrodoKEM should refer to IND-CPA.

My interpretation for this is due to two points: (i) the claimed security loss of Theorem 5.1 would either need to begin at IND-CPA if it follows HHK, or it would need to prove something fresh (rather than just giving a reference, and (ii) the fact that the 'proof' of Theorem 5.1 refers you to HHK's Theorem 3.2 rather than HHK's Theorem 3.1. (HHK-3.2 starts with IND-CPA while HHK-3.1 starts with OW-CPA; both arrive at OW-PCA when q_V = 0, but HHK-3.2 is 'tight' while HHK-3.1 is not quite tight).

As a tertiary point, it looks to me like FrodoKEM's Theorem 5.1 is -- in fact -- *correct* if you begin at IND-CPA and go through HHK-3.2 and HHK-3.4. (Interpretations that lead to the claimed outcome are probably the best interpretations to use.)

My conclusion so far: Yes, it looks like the FrodoKEM document needs to be cleaned up / re-written on page 30, but it looks like the technical matter is nonetheless correct at its core.

Regarding your separate *'d (starred) points:

* #1: I think the point is that FrodoPKE is 'naturally' IND-CPA under LWE. No need to use a transform to boost some OW-

CPA PKE to an IND-CPA PKE. (One can, of course, argue separately against the IND-CPA-ness of FrodoPKE, or argue separately against the validity of LWE itself, or even argue that lattice cryptography is fundamentally poorly-studied -- if one likes.)

* #2: The assumption being made by FrodoKEM is uniform-secure DLWE (Decisional Learning With Errors), so any looseness of search-to-decision reductions for LWE don't necessarily come into play. But this isn't about the tightness of the reduction underlying Theorem 5.1, where the tightness claim originally under question here was made. (One can, of course, argue separately against the usefulness of assuming DLWE vs search-LWE as the underlying assumption, given that practical attacks generally aim to recover keys, or similar -- if one likes.)

* #3: I'm not sure I understand what your point is here; feel free to clarify if this was something critical in your mind.

--Daniel

On Friday, May 24, 2019 at 4:33:40 AM UTC-4, D. J. Bernstein wrote:

> Page 30 of the Frodo submission claims that Theorem 5.1 of the
> submission is a tight ROM reduction showing that "FrodoKEM is an
> IND-CCA-secure KEM under the assumption that FrodoPKE is an
> OW-CPA-secure public-key encryption scheme".
>
> The theorem statement is somewhat more general, and claims that
> "Adv^ind-cca" of an attack against the KEM in Definition 2.19 is at most
>
> * 3 times "Adv^ow-cpa" of an attack against the underlying PKE, plus
> * the number q of hash queries times the decryption failure rate, plus
> * 3q+1 divided by the size of the message space.
>
> There is then a claim that "The proof of Theorem 5.1 is analogous to the
> proofs of Theorems 3.2 and 3.4 of Hofheinz, Hövelmanns, and Kiltz (HHK)
> [63]". This is followed by a few plausible comments about tweaks to the
> hashing; these tweaks aren't relevant to what I'm going to say here.
>
> I've looked for the claimed proof of Frodo Theorem 5.1 in some obvious
> spots and haven't found it. Has this theorem in fact been proven?
>
> It's also not clear to me what level of review has taken place regarding
> this proof, or the other Frodo proofs. In general, does the Frodo team
> vouch for the correctness of the proofs used in the Frodo submission?
> And, to be clear, this includes proofs of every "Theorem" claimed in the
> submission?
>
> If "ow-cpa" were replaced by "ind-cpa" then it would be clear (modulo
> various smaller details that I haven't checked---I don't vouch for
> correctness here!) how Frodo Theorem 5.1 relates to HHK Theorems 3.2 and
> 3.4. But switching from OW-CPA to IND-CPA is a huge change:
>
> * The HHK paper presents a way to construct ROM IND-CPA tightly from
> OW-CPA---but this is at the expense of much larger ciphertexts.
> Frodo doesn't use this.

From: daniel.apon <daniel.apon@nist.gov>
Sent: Friday, May 24, 2019 2:10 PM
To: pqc-forum
Cc: pqc-comments; djb@cr.yp.to
Subject: Re: ROUND 2 OFFICIAL COMMENT: Frodo

Of course, if it's actually important that FrodoKEM's security proof begin at OW-CPA instead of IND-CPA (as I assumed was intended, given that IND-CPA is what is shown of FrodoPKE via appeal to **D**LWE), then I don't immediately see how to recover full tightness using HHK out of the box. Is this the case?

--Daniel

On Friday, May 24, 2019 at 2:02:51 PM UTC-4, daniel.apon wrote:

> Hi Dan,
>
> *"I've looked for the claimed proof of Frodo Theorem 5.1 in some obvious spots and haven't found it. Has this theorem in fact been proven?"*
>
> Sure, let's verify it together.
>
> Note that FrodoKEM's specification document's Theorem 5.1 is a generic theorem, in that doesn't refer to any particular PKE (in particular, it doesn't refer to FrodoKEM's PKE).
>
> Theorem 5.1 in the FrodoKEM spec is obtained by beginning with an IND-CPA PKE, applying HHK's -
> - https://eprint.iacr.org/2017/604.pdf -- Theorem 3.2 (HHK, page 12) with q_V set to 0 to obtain a OW-PCA PKE via Transformation T (page 10) -- not OW-PCVA, since q_V is 0; and then applying HHK's Theorem 3.4 (HHK, page 16) via transform U^\notbot (page 15; it's U^\bot with implicit rejection) to obtain an IND-CCA KEM.
>
> The 'change' between HHK and FrodoKEM's Theorem 5.1 is that this transformation is done 'in one step,' which basically just means using a single hash function with longer output -- but I think we agree that that tweak is inconsequential to security.
>
> If I read what you're saying correctly, then we seem to agree that if FrodoKEM's spec is changed to say "IND-CPA" instead of "OW-CPA," that it would be more clear. To me, it looks like the "summary" paragraph titled "Security reductions" in Section 5.1 of FrodoKEM's spec (including Footnote 5) is mis-explained. Similarly, it looks like Theorem 5.1 of FrodoKEM should refer to IND-CPA.
>
> My interpretation for this is due to two points: (i) the claimed security loss of Theorem 5.1 would either need to begin at IND-CPA if it follows HHK, or it would need to prove something fresh (rather than just giving a reference, and (ii) the fact that the 'proof' of Theorem 5.1 refers you to HHK's Theorem 3.2 rather than HHK's Theorem 3.1. (HHK-3.2 starts with IND-CPA while HHK-3.1 starts with OW-CPA; both arrive at OW-PCA when q_V = 0, but HHK-3.2 is 'tight' while HHK-3.1 is not quite tight).
>
> As a tertiary point, it looks to me like FrodoKEM's Theorem 5.1 is -- in fact -- *correct* if you begin at IND-CPA and go through HHK-3.2 and HHK-3.4. (Interpretations that lead to the claimed outcome are probably the best interpretations

From: Mike Hamburg <mike@shiftleft.org>
Sent: Friday, May 24, 2019 5:22 PM
To: Apon, Daniel C. (Fed)
Cc: pqc-forum; pqc-comments; djb@cr.yp.to
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: Frodo

We discussed IND vs OW some at the Oxford PQC workshop. My conclusion from that was as follows:

IND is a much stronger assumption. OW is weaker because requires the adversary to be completely right in a many-bit guess, and it gives much more information in general than a 1/2 +- epsilon distinguisher (see [1], and for quantum reductions, [2]).

OW is a weaker assumption for a randomized scheme than for a deterministic one, to the point that it can't be used without a large tightness loss. This is because the simulator is extracting potential OW answers from the adversary's queries. if the scheme is also IND-CPA secure, then the simulator can't tell which query is the actual preimage, so it pretty much unavoidably loses a factor of q in tightness. By contrast, for a deterministic scheme OW is a slightly stronger assumption, and can be used without tightness loss.

This can be avoided by using an "OW-Conf" assumption, which is OW but the adversary has an oracle (classical, semi-classical or quantum) that checks a guess for the answer. This follows tightly, even in the sense of [1], from IND-CPA (and from IND-KPA), and it can be used in most (q)ROM proofs without as much tightness loss as OW-Passive. So possibly if you want an OW-type assumption for Frodo, then OW-Conf would be the way to do it.

For the specific case of LWE, neither IND nor OW assumptions have been studied in depth as far as I know. Search-LWE has been studied, but OW-Passive or OW-Conf security of LWE encryption doesn't follow from search-LWE unless the parameters are large enough for search-to-decision reduction.

Cheers,
— Mike

[1] Daniele Micciancio and Michael Walter. On the bit security of crypto- graphic primitives. In Jesper Buus Nielsen and Vincent Rijmen, editors, EUROCRYPT 2018, Part I, volume 10820 of LNCS, pages 3–28. Springer, Heidelberg, April / May 2018. doi:10.1007/978-3-319-78381-9_1.

[2] Haodong Jiang, Zhenfeng Zhang, and Zhi Ma. On the non-tightness of measurement-based reductions for key encapsulation mechanism in the quantum random oracle model. Cryptology ePrint Archive, Report 2019/494, 2019. https://eprint.iacr.org/2019/494.

On May 24, 2019, at 11:09 AM, 'daniel.apon' via pqc-forum <pqc-forum@list.nist.gov> wrote:

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Friday, May 24, 2019 8:38 PM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: Frodo
Attachments: signature.asc

'daniel.apon' via pqc-forum writes:
> * #3: I'm not sure I understand what your point is here; feel free to
> clarify if this was something critical in your mind.

Frodo uses distributions that don't match the underlying problem. Why is this supposed to avoid a massive security loss? Section 5.1.3 claims to answer this as follows:

(1) The Renyi divergence is limited. (I haven't checked this but
let's assume that the calculations are correct.)
(2) This preserves the relevant reductions.

The problem with #2---as explained in, e.g., the introduction of

https://ntruprime.cr.yp.to/divergence-20180430.pdf

---is that standard divergence arguments apply to OW-CPA and _not_ to IND-CPA. This is why you can't simply say that replacing OW-CPA with IND-CPA in Theorem 5.1 will magically fix the proof structure. Maybe intermediate notions such as OW-PCVA handle this issue, but this needs proof, and the Frodo submission doesn't have a proof.

This issue would disappear if the starting assumption were OW-CPA rather than IND-CPA. OW-CPA is what Theorem 5.1 claims, and what the summary on page 30 claims. Footnote 5 says "OW-CPA is for example defined in [63] and is implied by IND-CPA", which makes it very difficult to believe that the authors were merely confusing OW-CPA with IND-CPA. But I don't see how to prove the claim starting from OW-CPA.

> I think the point is that FrodoPKE is 'naturally' IND-CPA under LWE

Sure, this is an example of the trivial generic split of IND-CPA attacks into key distinguishers and ciphertext distinguishers. But then where's the proof that limited divergence is adequate?

Theorem 5.1 instead takes an extra detour through OW-CPA:

IND-CPA => OW-CPA (this is what footnote 5 says)
=> OW-CPA for modified distribution (divergence argument)
=> IND-CCA2 for the KEM (this is what Theorem 5.1 says)

But then where's the proof of Theorem 5.1?

> we seem to agree that if FrodoKEM's spec is changed to say "IND-CPA"
> instead of "OW-CPA," that it would be more clear.

I don't agree with this at all. Changing "OW-CPA" to "IND-CPA" makes a different (wimpier) statement with exactly the same level of clarity.

I also don't agree that this change is an "interpretation". There's a big difference between (1) filling in missing details of an unclear statement and (2) modifying an alleged "theorem" from something clear and unproven to something different. Obviously both situations raise questions regarding the level of review, but the second situation is clearly moving the goalposts while the first situation might not be.

> One can, of course, argue separately against the usefulness of
> assuming DLWE vs search-LWE as the underlying assumption

Yes. Page 4 claims that the relevant LWE parameters are "well studied", but this claim is even less well supported for the decisional assumption than it is for the search assumption.

---Dan

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Saturday, May 25, 2019 7:48 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: Frodo
Attachments: signature.asc

Checking the round-1 Frodo submission, I see that Theorem 5.1 in the
round-1 Frodo submission assumed IND-CPA. The round-2 submission

* changed the IND-CPA assumption in Theorem 5.1 to OW-CPA,
* added a footnote saying that IND-CPA implies OW-CPA---the
  probability gap here is of course O(1/M)---and
* made a O(q/M) change to the probability gap in Theorem 5.1.

The added round-2 detour through OW-CPA makes sense to me as a way to resolve the Renyi-divergence issue (even if it sounds "unnatural" from the perspective of a trivial IND-based key-ciphertext split). But it's then critical to have a proof that starts from OW-CPA (as round-2 Frodo Theorem 5.1 claims), not just IND-CPA (the round-1 version).

Is there a proof of round-2 Frodo Theorem 5.1 somewhere?

---Dan

From: daniel.apon <daniel.apon@nist.gov>
Sent: Saturday, May 25, 2019 11:49 AM
To: pqc-forum
Cc: pqc-comments; djb@cr.yp.to
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: Frodo

Hi Dan,

Thanks for clarifying; I appreciate it.
Allow me to re-state, just to properly frame the context of my comment again: I do not yet see a tight reduction for FrodoKEM passing through OW-CPA.

In any case, there is obviously incorrect information in the Round 2 spec for FrodoKEM, as the Round 2 spec refers to OW-CPA PKE, then references the IND-CPA form of HHK -- Theorem 3.2.

One 'fix' for the incorrect information would be to do as you (independently) recommended: Stick with OW-CPA. Then, the theorem statement could be modified to not be tight.

Another possible 'fix' would be give a fresh proof (of any type) beginning at OW-CPA that is tight (as you have asked for between four and six times in this thread, depending on how you count).

One route for the above might be to explore Mike's idea of proving security via OW-Conf type notions.

Another possible 'fix' would be to explore Bai et al.'s comment in https://eprint.iacr.org/2015/483.pdf on Bai-page-21 (paper cited as [15] in FrodoKEM's Round 2 spec, on page 8) which reads as follows:

"We remark that the search-decision equivalence idea in the proof of Theorem 5.1 could be extended to show the hardness of the decision LWE problem with any noise distribution ψ, with respect to the hardness of LWE with Gaussian noise D_α if ... ψ is 'close' to D_α in the sense of RD (i.e., $R(\psi||D_\alpha)$ is 'small') ..."

However, I would point out that
(i) FrodoKEM's Round 2 spec does not take this approach (instead referring to OW-CPA in Round 2), so this would require reverting to FrodoKEM's Round 1 approach, then extending it; and
(ii) Bai et al. have only an informal comment about this direction, but not a formal proof which lays out all of the nitty-gritty details (the latter of which is very important to have written down and reviewed before it could be accepted more broadly).

In particular, it's clear from FrodoKEM's Section 5.1.3 "Approximating the error distribution" that their error distribution and the related discrete Gaussian are close via RD. What would be needed here (for this direction to be convincing) is a clean proof of IND-CPA under this alternative distribution, which has not yet appeared. But -- if so -- then the original route of HHK-3.2 then HHK-3.4 could be taken via IND-CPA, and give a tight reduction.

I suspect we will likely end up needing the FrodoKEM Team to chime in to resolve these issues further (or someone to write a nice research paper -- P.S. NIST's 2nd PQC Standardization Conference submission deadline is in a few days :-)).

--Daniel

On Saturday, May 25, 2019 at 7:48:16 AM UTC-4, D. J. Bernstein wrote:

> Checking the round-1 Frodo submission, I see that Theorem 5.1 in the
> round-1 Frodo submission assumed IND-CPA. The round-2 submission
>
> * changed the IND-CPA assumption in Theorem 5.1 to OW-CPA,
> * added a footnote saying that IND-CPA implies OW-CPA---the
> probability gap here is of course O(1/M)---and
> * made a O(q/M) change to the probability gap in Theorem 5.1.
>
> The added round-2 detour through OW-CPA makes sense to me as a way to
> resolve the Renyi-divergence issue (even if it sounds "unnatural" from
> the perspective of a trivial IND-based key-ciphertext split). But it's
> then critical to have a proof that starts from OW-CPA (as round-2 Frodo
> Theorem 5.1 claims), not just IND-CPA (the round-1 version).
>
> Is there a proof of round-2 Frodo Theorem 5.1 somewhere?
>
> ---Dan

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Monday, June 24, 2019 1:14 PM
To: pqc-comments
Cc: pqc-forum
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: Frodo
Attachments: signature.asc

Update after a month: I still don't see how to prove Frodo's claimed Theorem 5.1. However, the claimed theorem still hasn't been withdrawn.

The official "Submission Requirements and Evaluation Criteria" say that "proofs will be considered if they are available". If the Frodo team is still trying to find a proof (or, maybe more productive, a way to rescue the Renyi-divergence claims if Theorem 5.1 is replaced by something
weaker) then in the meantime surely it has to withdraw the submission's claim that a proof is available. Alternatively, if the Frodo team claims that I'm mistaken and that a proof has been available for Theorem 5.1 all along, surely this claim should also be made public.

---Dan

From: Douglas Stebila <dstebila@gmail.com>
Sent: Tuesday, June 25, 2019 8:27 AM
To: D. J. Bernstein
Cc: pqc-comments
Subject: Re: [pqc-forum] ROUND 2 OFFICIAL COMMENT: Frodo

We're finalizing our revisions and plan to have a response and revision in the next couple of days. Sorry for the delay.

Douglas

> On Jun 24, 2019, at 1:14 PM, D. J. Bernstein <djb@cr.yp.to> wrote:
>
> Update after a month: I still don't see how to prove Frodo's claimed
> Theorem 5.1. However, the claimed theorem still hasn't been withdrawn.
>
> The official "Submission Requirements and Evaluation Criteria" say
> that "proofs will be considered if they are available". If the Frodo
> team is still trying to find a proof (or, maybe more productive, a way
> to rescue the Renyi-divergence claims if Theorem 5.1 is replaced by
> something
> weaker) then in the meantime surely it has to withdraw the
> submission's claim that a proof is available. Alternatively, if the
> Frodo team claims that I'm mistaken and that a proof has been
> available for Theorem 5.1 all along, surely this claim should also be made public.
>
> ---Dan
>
> --
> 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/20190624171410.24604.qmail%40cr.yp.to.

From: daniel.apon <daniel.apon@nist.gov>
Sent: Tuesday, June 25, 2019 10:35 AM
To: pqc-forum
Cc: pqc-comments; djb@cr.yp.to
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: Frodo

"surely it has to withdraw the submission's claim that a proof is available"

The outstanding question was whether the proof could be made tight (or tighter), not whether a proof exists.

Thanks.

On Monday, June 24, 2019 at 3:15:56 PM UTC-4, D. J. Bernstein wrote:

> Update after a month: I still don't see how to prove Frodo's claimed
> Theorem 5.1. However, the claimed theorem still hasn't been withdrawn.
>
> The official "Submission Requirements and Evaluation Criteria" say that
> "proofs will be considered if they are available". If the Frodo team is
> still trying to find a proof (or, maybe more productive, a way to rescue
> the Renyi-divergence claims if Theorem 5.1 is replaced by something
> weaker) then in the meantime surely it has to withdraw the submission's
> claim that a proof is available. Alternatively, if the Frodo team claims
> that I'm mistaken and that a proof has been available for Theorem 5.1
> all along, surely this claim should also be made public.
>
> ---Dan

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Tuesday, June 25, 2019 1:18 PM
To: pqc-comments
Cc: pqc-forum
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: Frodo
Attachments: signature.asc

NIST's official evaluation criteria ask

* whether mathematical structure is "well understood",
* whether relevant research is "established",
* what the "maturity" of analysis is,

etc., making clear reference to the analysis timeline.

The timeline here appears to be that Frodo made a new provable-security claim in March 2019, round-2 Frodo Theorem 5.1, and is now obliged to withdraw the claim. This in turn could have serious consequences for round-2 Frodo's divergence claims (and for the different round-1 Frodo divergence claims, which were withdrawn for other reasons). If the divergence claims fail then I see no justification for the claim in NIST IR 8240 that Frodo's LWE secrets "are sampled from a discrete Gaussian distribution"---which in turn seems to be a prerequisite for applying certain theorems that are claimed to "support" the security of Frodo.

I'm saying "appears to be" because, after a month, the Frodo team has not yet spoken up to withdraw round-2 Frodo Theorem 5.1. Does this mean that the Frodo team thinks I've made a mistake, and thinks the theorem is fine as is? If this takes multiple rounds of discussion to resolve, maybe delaying each round by a month isn't the best idea.

The following questions are also unanswered: "It's also not clear to me what level of review has taken place regarding this proof, or the other Frodo proofs. In general, does the Frodo team vouch for the correctness of the proofs used in the Frodo submission? And, to be clear, this includes proofs of every 'Theorem' claimed in the submission?"

Daniel Apon writes:
> "surely it has to withdraw the submission's claim that a proof is available"
> The outstanding question was whether the proof could be made tight (or
> tighter), not whether a proof exists.

Now I'm really puzzled. Are you claiming that you see how to prove Theorem 5.1 as stated in the round-2 Frodo submission? How does the proof work?

The theorem has a clear probability bound that allows very little wiggle room---certainly not enough for an HHK17-style proof from OW-CPA. The theorem also says "the running time of B is about that of A", which violates the mathematical requirement for each statement in a theorem to have a clear definition---but nobody would accept interpreting this as allowing a massive slowdown, so you can't simply plug in a very slow high-probability B. (Another problem with such an A-independent proof strategy is that it would allow (FrodoKEM,FrodoPKE) to be replaced with (X,FrodoPKE) for an arbitrarily weak KEM X, so it would obviously say nothing about FrodoKEM's security. But, more to the point, this strategy doesn't prove the claimed theorem.)

The summary of the theorem earlier on page 30 of the Frodo submission also includes a tightness claim:

> FrodoKEM is an IND-CCA-secure KEM under the assumption that FrodoPKE
> is an OW-CPA-secure public-key encryption scheme, and where G2 and F
> are modeled as random oracles. Theorem 5.1 gives a tight, classical

reduction against classical adversaries in the classical random oracle model.

Are you claiming that you see a "tight, classical reduction" deducing "IND-CCA" security for this KEM from "OW-CPA" security for this PKE?

I see no reason to believe that a proof exists. I don't understand why these claims from the Frodo submission still haven't been withdrawn.

---Dan

From: Michael Naehrig <mnaehrig@microsoft.com>
Sent: Tuesday, July 2, 2019 4:36 PM
To: D. J. Bernstein; pqc-comments
Cc: pqc-forum
Subject: RE: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: Frodo

Dear all,

First off, we apologize for the long delay. We wanted to respond to all of the raised issues at once, and coordinating that response took longer than we expected.

There were two issues identified in the email thread Dan started.

The first issue is Theorem 5.1 about the tight IND-CCA security of FrodoKEM from the OW-CPA (instead of IND-CPA) security of FrodoPKE in the classical random oracle model.

Indeed, the change in hypothesis from IND-CPA to OW-CPA was a typo that was inadvertently introduced in the revisions between the round-1 and round-2 submissions. We did not intend to claim a new tight security proof from OW-CPA security; as stated in both submissions, we rely on the prior results of Hofheinz, Hövelmanns, and Kiltz. (In the round-2 submission we also use a more recent result of Jiang, Zhang, Chen, Wang, and Ma on IND-CCA security from OW-CPA security, in the
quantum random oracle model.)

The second issue is how the Rényi divergence argument applies in the security reductions.

FrodoKEM uses specific error distributions in its instantiations, but we claim security based on LWE with rounded Gaussian distributions.
We argued that this substitution was sound due to the results of Langlois, Stehlé, and Steinfeld, because the Rényi divergence of the new distribution from the original one is sufficiently small. However, as pointed out, the results of LSS hold for search problems, whereas the problems named in the round 2 submission (IND-CCA, IND-CPA, decision-LWE) are decision problems.

Nonetheless, there *is* a search problem in the HHK reductions (for the classical ROM): the sequence is IND-CPA -> OW-PCA -> IND-CCA, so we can apply the Rényi divergence argument at the OW-PCA step.

We have revised Section 5.1 of the specification to make this explicit. In particular, the new Theorem 5.1 shows the result of combining the IND-CPA -> OW-PCA -> IND-CCA reductions with the Rényi divergence argument at the OW-PCA step. Various other parts of Section 5.1-5.1.4 have been reorganized as part of this revision.

We emphasize that neither the FrodoKEM scheme nor any of its parameters have changed. In addition, this analysis shows that the IND-CCA security of FrodoKEM can be tightly based solely on the OW-PCA security of the "T-transformed" (deterministic) version of FrodoPKE.
In other words, IND-CPA security of FrodoPKE and hardness of decision-LWE are not *necessary* assumptions, but they can be used to show OW-PCA security.

Finally, we have added concrete calculations of FrodoKEM's IND-CCA bit security under the chain of classical reductions, assuming our bit-security estimates for the LWE problem, taking into account the various losses associated with the reductions (number of LWE samples, Rényi divergence, probability of decryption failure). The new Table 2 shows these estimates.

A few other insubstantial typos were fixed during our revisions, and various other arguments were made more precise. A list of changes appears at the end of the revised document.

Our revised specification document, as well as all previous versions, is available at https://gcc01.safelinks.protection.outlook.com/?url=https%3A%2F%2Fwww.frodokem.org%2F&data=02%7C01%7Csara.kerman%40nist.gov%7C279649de9ffc482c850f08d6ff2ce7c4%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C1%7C636976965661473236&sdata=4ohRN2gLfDQ6iVvylkqjO%2B8iZjSIjVZpAOhRSDc5h%2B0%3D&reserved=0.

The FrodoKEM team

-----Original Message-----
From: D. J. Bernstein <djb@cr.yp.to>
Sent: Tuesday, June 25, 2019 10:18 AM
To: pqc-comments@nist.gov
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: Frodo

NIST's official evaluation criteria ask

* whether mathematical structure is "well understood",
* whether relevant research is "established",
* what the "maturity" of analysis is,

etc., making clear reference to the analysis timeline.

The timeline here appears to be that Frodo made a new provable-security claim in March 2019, round-2 Frodo Theorem 5.1, and is now obliged to withdraw the claim. This in turn could have serious consequences for round-2 Frodo's divergence claims (and for the different round-1 Frodo divergence claims, which were withdrawn for other reasons). If the divergence claims fail then I see no justification for the claim in NIST IR 8240 that Frodo's LWE secrets "are sampled from a discrete Gaussian distribution"---which in turn seems to be a prerequisite for applying certain theorems that are claimed to "support" the security of Frodo.

I'm saying "appears to be" because, after a month, the Frodo team has not yet spoken up to withdraw round-2 Frodo Theorem 5.1. Does this mean that the Frodo team thinks I've made a mistake, and thinks the theorem is fine as is? If this takes multiple rounds of discussion to resolve, maybe delaying each round by a month isn't the best idea.

The following questions are also unanswered: "It's also not clear to me what level of review has taken place regarding this proof, or the other Frodo proofs. In general, does the Frodo team vouch for the correctness of the proofs used in the Frodo submission? And, to be clear, this includes proofs of every 'Theorem' claimed in the submission?"

Daniel Apon writes:
> "surely it has to withdraw the submission's claim that a proof is available"
> The outstanding question was whether the proof could be made tight (or
> tighter), not whether a proof exists.

Now I'm really puzzled. Are you claiming that you see how to prove Theorem 5.1 as stated in the round-2 Frodo submission? How does the proof work?

The theorem has a clear probability bound that allows very little wiggle room---certainly not enough for an HHK17-style proof from OW-CPA. The theorem also says "the running time of B is about that of A", which violates the mathematical requirement for each statement in a theorem to have a clear definition---but nobody would accept interpreting this as allowing a massive slowdown, so you can't simply plug in a very slow high-probability B. (Another problem with such an A-independent proof strategy is that it would allow (FrodoKEM,FrodoPKE) to be replaced with

From: Christopher J Peikert <cpeikert@alum.mit.edu>
Sent: Tuesday, July 30, 2019 12:34 PM
Cc: pqc-comments; pqc-forum
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: Frodo

Speaking for myself, I'd like to further highlight this:

> We have revised Section 5.1 of the specification to make this
> explicit. In particular, the new Theorem 5.1 shows the result of
> combining the IND-CPA -> OW-PCA -> IND-CCA reductions with the Rényi
> divergence argument at the OW-PCA step. Various other parts of Section
> 5.1-5.1.4 have been reorganized as part of this revision.
>
> We emphasize that neither the FrodoKEM scheme nor any of its
> parameters have changed. In addition, this analysis shows that the
> IND-CCA security of FrodoKEM can be tightly based solely on the OW-PCA
> security of the "T-transformed" (deterministic) version of FrodoPKE.
> In other words, IND-CPA security of FrodoPKE and hardness of
> decision-LWE are not *necessary* assumptions, but they can be used to
> show OW-PCA security.

I believe this point is relevant to Mike Hamburg's comments on IND versus OW assumptions, and decryption failures, here: https://groups.google.com/a/list.nist.gov/d/msg/pqc-forum/bW09IcPVG6E/vJ0N1FKdBAAJ

As stated above, in the classical ROM, we get (from HHK) a tight proof of FrodoKEM's IND-CCA security assuming OW-PCA security of the
(deterministic) "T-transformed" FrodoPKE -- call it T-FrodoPKE.

Now, OW-PCA is very similar to but not quite OW-CPA: the attacker also has a "plaintext-checking oracle" that, given a query (m,c), answers whether c decrypts to m. (This might coincide with the OW-Conf notion that Mike mentioned, but I couldn't find a formal definition.)

Because T-FrodoPKE's decryption algorithm returns m only if Enc_pk(m) = c, in any other case the oracle's output is predictably 0, and hence useless. So, we can think of the oracle as taking just m as input.

This oracle is useless (it always returns 1) unless the attacker queries an m whose T-FrodoPKE ciphertext decrypts *under FrodoPKE* to some m' != m. (Note that FrodoPKE decryption always returns a message, never \bot.) In other words, the encryption "coins" derived from m (via the RO) must induce an incorrect decryption.

The cost of finding such an m for several LWE/LWR-based schemes was studied in https://gcc01.safelinks.protection.outlook.com/?url=https%3A%2F%2Feprint.iacr.org%2F2018%2F1089&data=02%7C01%7Csara.kerman%40nist.gov%7Cf46cc917b58441dc812c08d7150bd436%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C1%7C637001012861084693&sdata=1%2FuBVRULLM9YP3C5gZTFabRloQg6QUar2eKfwvPEelw%3D&reserved=0 . For FrodoKEM-976 -- alone among the studied systems -- the authors found that the cost exceeded that of other attacks, i.e., there was no loss of security due to the potential for decryption failures. (See Table 1 just before Section 6.) It would be nice to see this analysis adapted to the other proposed FrodoKEM parameters.

In summary: in the classical ROM, the IND-CCA security of FrodoKEM can be based tightly on a one-wayness assumption which appears no stronger than OW-CPA for (at last some) concretely proposed parameters.

Sincerely yours in cryptography,
Chris

From: Mike Hamburg <mike@shiftleft.org>
Sent: Tuesday, July 30, 2019 1:56 PM
To: Christopher J Peikert
Cc: pqc-comments; pqc-forum
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: Frodo

On Jul 30, 2019, at 9:34 AM, Christopher J Peikert <cpeikert@alum.mit.edu> wrote:

Speaking for myself, I'd like to further highlight this:

> We have revised Section 5.1 of the specification to make this explicit. In particular, the new Theorem 5.1 shows the result of combining the IND-CPA -> OW-PCA -> IND-CCA reductions with the Rényi divergence argument at the OW-PCA step. Various other parts of Section 5.1-5.1.4 have been reorganized as part of this revision.
>
> We emphasize that neither the FrodoKEM scheme nor any of its parameters have changed. In addition, this analysis shows that the IND-CCA security of FrodoKEM can be tightly based solely on the OW-PCA security of the "T-transformed" (deterministic) version of FrodoPKE. In other words, IND-CPA security of FrodoPKE and hardness of decision-LWE are not *necessary* assumptions, but they can be used to show OW-PCA security.

I believe this point is relevant to Mike Hamburg's comments on IND versus OW assumptions, and decryption failures, here: https://groups.google.com/a/list.nist.gov/d/msg/pqc-forum/bW09IcPVG6E/vJ0N1FKdBAAJ

As stated above, in the classical ROM, we get (from HHK) a tight proof of FrodoKEM's IND-CCA security assuming OW-PCA security of the (deterministic) "T-transformed" FrodoPKE -- call it T-FrodoPKE.

Now, OW-PCA is very similar to but not quite OW-CPA: the attacker also has a "plaintext-checking oracle" that, given a query (m,c), answers whether c decrypts to m. (This might coincide with the OW-Conf notion that Mike mentioned, but I couldn't find a formal definition.)

Hi Chris,

Here is the formal definition of OW-Conf for an rPKE:

(pk, sk) <- keygen()
m* <- uniformly random from message space
c* <- Encrypt(pk, m*)
Oracle Conf(m): return m == m*

ret <- Adv^Conf(pk,c*)
Adv wins if ret == m*, or equivalently if Conf(ret).

It's a weaker assumption than OW-PCA, because you can only check whether the challenge ciphertext matches. It's stronger than OW-Passive: for an rPKE if you think you've solved OW-Passive, you can't check your answer, but for OW-Conf you can.

You can split this into at least three variants: Conf oracle is classical, Conf oracle is quantum, or Conf oracle is semiclassical. The one that slots in tightly (up to a constant factor) into BHHP'19 is the semiclassical one, which is why we didn't put it in that paper — we thought it was a little too unnatural.

The technique used in BHHP'19 should prove:

OW-Passive -> OW-Conf-semiclassical with O(q) tightness loss. Proof: lemma 4.
IND-CPA -> OW-Conf-semiclassical with constant tightness loss. Proof: roughly half of Theorem 1, but puncture on either m_0 or m_1 at random instead of both (to meet the definition of OW-Conf).
OW-Conf-semiclassical -> OW-Passive T(PKE,G) with O(d) tightness loss. Proof: the other half of Theorem 1.

The IND-CPA -> OW-Conf-semiclassical should work with constant tightness loss even under Micciancio-Walter's definition of IND-CPA [MW'18]. MW's definition is roughly that the adversary can abstain instead of guessing, but guessing wrong counts against the advantage. This should forbid useless distinguishers, including "Breaking AES with MD5" [BL'12]. I sketched these proofs and discussed them with the other BHHP authors, but we didn't formally write them up, so there's always a chance that there is a mistake.

And as Chris mentions, that's the same as OW-PCA T(PKE,G) up to finding failing messages in the ROM, but BHHP'19 instead uses finding failing ciphertexts. That approach is less natural but (with current proofs) tighter in the QROM.

Cheers,
— Mike (also speaking for myself, and not the other BHHP authors)

[BHHP'19] Nina Bindel and Mike Hamburg and Andreas Hülsing and Edoardo Persichetti. Tighter proofs of CCA security in the quantum random oracle model. https://eprint.iacr.org/2019/590

[MW'18] Daniele Micciancio and Michael Walter. On the Bit Security of Cryptographic Primitives. https://eprint.iacr.org/2018/077

[BL'12] Daniel J. Bernstein and Tanja Lange. Non-uniform cracks in the concrete: the power of free precomputation. https://eprint.iacr.org/2012/318

From: Christopher J Peikert <cpeikert@alum.mit.edu>
Sent: Tuesday, August 6, 2019 3:10 PM
Cc: pqc-comments; pqc-forum
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: Frodo

> In other words, the encryption "coins" derived from m (via the RO)
> must induce an incorrect decryption.
>
> The cost of finding such an m for several LWE/LWR-based schemes was
> studied in https://eprint.iacr.org/2018/1089
> For FrodoKEM-976 -- alone among the studied systems -- the authors found that the cost exceeded that of other attacks, i.e., there was no loss of security due to the potential for decryption failures. (See Table 1 just before Section 6.) It would be nice to see this analysis adapted to the other proposed FrodoKEM parameters.

To be clear, 2018/1089 concluded that there was no loss of *claimed* security for FrodoKEM-976.

From: Jan-Pieter D'Anvers <janpieter.danvers@esat.kuleuven.be>
Sent: Monday, August 12, 2019 4:11 AM
To: Christopher J Peikert
Cc: pqc-comments; pqc-forum
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: Frodo

I redid the calculations for the other FrodoKEM instances and it seems that there is no loss in quantum security for any of them under the attack of 2018/1089. Maybe noteworthy is the failure rate of 2^252.5 for FrodoKEM-1344, which is close to the 2^256 possible ciphertexts due to the FO transformation. In this case, there are on average only 2^3.5 or
11 failures for each secret.

Take into account that this attack scenario is under the assumption that an adversary can do an unlimited number of decryption queries, which is already a very optimistic scenario from an attackers point of view.

Best regards,

Jan-Pieter

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Tuesday, July 28, 2020 12:14 PM
To: pqc-comments
Cc: pqc-forum
Subject: ROUND 2 OFFICIAL COMMENT: Frodo
Attachments: signature.asc

The latest NIST report states that Frodo in TLS key exchange would cost "around 20,000 bytes" plus "2 million cycles" for the server. NIST appears to conclude from this that Frodo does not have "acceptable performance in widely used applications overall".

I'm filing this comment to request explanation of the basis for NIST's claim that 20000 bytes plus 2 million cycles would not be "acceptable performance" for post-quantum TLS key exchange.

Google said (https://www.imperialviolet.org/2018/04/11/pqconftls.html)
that the unstructured-lattice size is "probably not preferable for real-time TLS connections". This does not justify NIST's black-and-white claim that the size isn't "acceptable". The words "real-time" are also
important: there are many ways to avoid having a user wait for a key exchange. Google documented software bugs causing problems with these sizes for a particular way of integrating post-quantum crypto into TLS, but this can be worked around.

I should note that the above claim is my understanding of what NIST is saying regarding Frodo performance, but the text is somewhat ambiguous:

> The resulting potential security advantages of Frodo are paid for with far worse performance in all metrics than other lattice schemes. ... Use of FrodoKEM would have a noticeable performance impact on high traffic TLS servers, where each server does decapsulation which requires close to 2 million cycles for the best performing parameter set (FrodoKEM-640-AES) and receives a public key and a ciphertext (around 20,000 bytes in total) for every fresh key exchange.
>
> In NIST's view, FrodoKEM may be suitable for use cases where the high confidence in the security of unstructured lattice-based schemes is much more important than performance. NIST's first priority for standardization is a KEM that would have acceptable performance in widely used applications overall. As such, possible standardization for FrodoKEM can likely wait until after the third round.

This doesn't _directly_ say that TLS is the "widely used application" in which Frodo doesn't have "acceptable performance", so perhaps NIST meant something else, but then it's weird that TLS is the only example given.

---Dan

P.S. It's also worrisome to see NIST expressing "high confidence in the security of unstructured lattice-based schemes", as if the claimed asymptotic lattice security levels weren't 42% higher just 10 years ago and superexponentially higher just 20 years ago. Asking for the basis for NIST's claims regarding acceptable application performance should not be interpreted as endorsing overconfident security claims.

From: Perlner, Ray A. (Fed)
Sent: Wednesday, July 29, 2020 11:19 AM
To: D. J. Bernstein; pqc-comments
Cc: pqc-forum
Subject: RE: ROUND 2 OFFICIAL COMMENT: Frodo

Dear Dan,

While it is not possible to speak for what every user of our standards would or wouldn't find "acceptable", there is a pretty large difference between the performance of Frodo on the one hand and Kyber, NTRU, and Saber on the other hand. We are therefore more confident that Kyber, NTRU, or Saber will be considered "acceptable" for most users than that Frodo will.

-----Original Message-----
From: D. J. Bernstein <djb@cr.yp.to>
Sent: Tuesday, July 28, 2020 12:14 PM
To: pqc-comments <pqc-comments@nist.gov>
Cc: pqc-forum <pqc-forum@list.nist.gov>
Subject: ROUND 2 OFFICIAL COMMENT: Frodo

The latest NIST report states that Frodo in TLS key exchange would cost "around 20,000 bytes" plus "2 million cycles" for the server. NIST appears to conclude from this that Frodo does not have "acceptable performance in widely used applications overall".

I'm filing this comment to request explanation of the basis for NIST's claim that 20000 bytes plus 2 million cycles would not be "acceptable performance" for post-quantum TLS key exchange.

Google said (https://www.imperialviolet.org/2018/04/11/pqconftls.html)
that the unstructured-lattice size is "probably not preferable for real-time TLS connections". This does not justify NIST's black-and-white claim that the size isn't "acceptable". The words "real-time" are also
important: there are many ways to avoid having a user wait for a key exchange. Google documented software bugs causing problems with these sizes for a particular way of integrating post-quantum crypto into TLS, but this can be worked around.

I should note that the above claim is my understanding of what NIST is saying regarding Frodo performance, but the text is somewhat ambiguous:

> The resulting potential security advantages of Frodo are paid for
> with far worse performance in all metrics than other lattice schemes.
> ... Use of FrodoKEM would have a noticeable performance impact on
> high traffic TLS servers, where each server does decapsulation which
> requires close to 2 million cycles for the best performing parameter
> set (FrodoKEM-640-AES) and receives a public key and a ciphertext
> (around 20,000 bytes in total) for every fresh key exchange.
>
> In NIST's view, FrodoKEM may be suitable for use cases where the high