From: D. J. Bernstein <djb@cr.yp.to>
Sent: Saturday, April 21, 2018 6:16 PM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: OFFICIAL COMMENT: Frodo
Attachments: signature.asc

Summary: The "FrodoKEM" submission claims that various theorems support the security of the submission. This claim is incorrect for at least two of the stated theorems: these two theorems do not, in fact, support the security of the submission.

The exact quote at issue is in Section 5.1:

> 5 Justification of security strength
>
> The security of FrodoKEM is supported both by security reductions and by analysis of the best known cryptanalytic attacks.
>
> 5.1 Security reductions
>
> A summary of the reductions supporting the security of FrodoKEM is as follows: ... Theorem 5.2 gives a non-tight, classical reduction against quantum adversaries in the quantum random oracle model. ... Theorem 5.8 gives a non-tight classical reduction against classical or quantum adversaries (in the standard model).

Qualitatively, the message of the two theorems is that the following two attack avenues are ruled out against FrodoKEM:

(1) generic-hash quantum chosen-ciphertext attacks that don't break IND-CPA security of the underlying PKE;

(2) attacks against the relevant LWE parameters that don't break various worst-case lattice problems.

In fact, the theorems do _not_ rule out these attack avenues. A user of FrodoKEM is not guaranteed security against attacks of these two types.
(Please note that I'm _not_ saying that attacks have been shown to exist along these two lines.) For #1 this is because the stated theorem is too loose to say anything regarding the KEMs that were actually submitted.
For #2 the theorem statement lacks the quantification necessary to even see how loose the proof is; I'm skeptical of the idea that a tight enough theorem is possible, and in any case the claim of applicability of the stated theorem is incorrect.

Terminology: It's important to be aware of a conflict between two different definitions of "KEM". For concrete literature such as

http://shoup.net/iso/std6.pdf

a KEM has a concrete size, a concrete security level, etc. This is the definition that matters for standards, implementations, users, and this message. Most encryption submissions specify more than one KEM, but, ultimately, whichever particular KEM the user ends up using is the KEM whose security and performance matter for that user.

Old-fashioned asymptotic definitions instead use the word "KEM" to refer to a particular type of infinite _family_ of KEMs. Specifically, key generation in a "KEM" takes a nonnegative-integer security parameter as input---so one can't even talk about, e.g., the performance of commonly used KEMs built from Curve25519, or the security of ntruhrss701. This definition is also not compatible with statements such as

FrodoKEM’s communication size is sufficiently small that it is still
compatible with many existing deployments

in the FrodoKEM submission: the supporting details provided for this statement are specific to the KEMs actually contained in the submission.

Details regarding the FrodoKEM security claims: The submission contains two KEMs, FrodoKEM-640 and FrodoKEM-976 (plus a further division between AES and SHAKE, not relevant here). Table 1 lists the following failure probabilities for these two KEMs:

* 1/2^148.8 for FrodoKEM-640;
* 1/2^199.6 for FrodoKEM-976.

"Theorem 5.2 (IND-CPA PKE => IND-CCA KEM in quantum ROM)" refers to the failure probability as "delta" and states that various types of attacks succeed with probability at most

9*q_RO*sqrt(q_RO^2*delta+q_RO*sqrt(Adv_PKE^{ind-cpa}(B)+1/|M|)).

The "q_RO" term is the number of hash queries carried out inside the attack. Let's focus specifically on an attack carrying out 2^50 hash queries. What the theorem states is that such an attack succeeds with probability at most

9*2^50*sqrt(2^100*delta+2^50*sqrt(Adv_PKE^{ind-cpa}(B)+1/|M|)).

In particular, for the delta = 1/2^148.8 claimed for FrodoKEM-640, the theorem states that the attack succeeds with probability at most

457731076.16... + More

where "More" is a nonnegative term reflecting the addition of Adv_PKE^{ind-cpa}(B)+1/|M| inside the square root. Similarly, for the delta = 1/2^199.6 claimed for FrodoKEM-976, the theorem states that the attack succeeds with probability at most

10.338... + More.

Both of these statements are superseded by the content-free observation that an attack succeeds with probability <=1. (Even stronger, <=1/2 for security definitions that define success as |Pr-1/2| without doubling.
Both 1/2 and 1 are adequate for my main point in this message.)

As noted above, this theorem is cited in a list of "reductions supporting the security of FrodoKEM". In fact, the theorem does not support the security of FrodoKEM. For all submitted FrodoKEM parameters, the theorem becomes content-free before 2^50 hash queries.

The submission notes that the theorem is "non-tight" (and that this was ignored in parameter selection), but this does not correct the error in the "supporting the security of FrodoKEM" claim. There's a big difference between

* quantitatively weak but still saying something about security and
* quantitatively so weak as to say _nothing_ about security.

One could change parameters to build KEMs for which the theorem _would_ say something, depending on Adv_PKE^{ind-cpa}(B) and |M|. However, those KEMs would be slower and larger than what the submission proposes. Other statements in the submission such as

> FrodoKEM’s communication size is sufficiently small that it is still
> compatible with many existing deployments

would have to be modified accordingly. A hypothetical submission S that glues together

* S-640: small enough but theorem does not apply
* S-976: small enough but theorem does not apply
* S-MuchBigger: theorem applies but not small enough

cannot make the exaggerated claim that S is small enough and backed by a theorem; the submission has to be clear about

* the inapplicability of the theorem to S-640 and S-976, and
* the inapplicability of the smallness claims to S-MuchBigger.

More to the point, if S-MuchBigger is omitted and the submission has

* S-640: small enough but theorem does not apply
* S-976: small enough but theorem does not apply

then the not-small-enough disclaimer regarding S-MuchBigger becomes unnecessary but the theorem becomes irrelevant to what was actually submitted.

The FrodoKEM submission also claims that the tightness gap "seems to be an artifact of the proof technique" since there is "no known attack that takes advantage of the tightness gap". There is no citation to previous work giving examples of the nightmare scenario that tightness gaps are exploitable, such as the following:

> https://www.youtube.com/watch?v=l56ORg5xXkk

There is also a puzzling inconsistency between how FrodoKEM treats two specific gaps between known attacks and theorems:

* For the _tightness_ gap discussed above, FrodoKEM chooses parameters where the theorems are inapplicable, and points to the lack of known attacks as justification.

* For the gap in _error width_, FrodoKEM insists on choosing parameters that are claimed to make some theorems applicable, and seems to portray these large parameters as an advantage, even when this is not required for protection against known attacks.

In any case, whether or not attacks are known, the simple fact is that the stated theorem is content-free for all submitted parameters and therefore does not support the security of the submission.

Another theorem, Theorem 5.8, is also claimed to support the security of FrodoKEM. However, the theorem uses words such as "negligible" and "poly" without stating quantitative details, and therefore says nothing about FrodoKEM-640, FrodoKEM-976, or any other specific KEM.

Previous work https://eprint.iacr.org/2016/360.pdf by Chatterjee, Koblitz, Menezes, and Sarkar filled in many quantitative details of an older theorem (also cited in the FrodoKEM submission) of a similar type, relating an LWE-type problem to a lattice-type problem. The conclusion was that, for typical system parameters, the unstated monomial factor was 2^504, presumably making the theorem useless.

(The reason I'm saying "presumably" is that the unstated _constant_ factor wasn't analyzed, and could conceivably be far enough below 1 to outweigh the monomial factor. However, experience suggests that the constant factor is also above 1.)

Of course it's possible that a new theorem is quantitatively better, but this needs to be stated and proven. The statement presented as "Theorem 5.8" is missing the quantification required for applicability to any particular KEM, so it does not support the security of the FrodoKEM submission.

Elsewhere in the FrodoKEM submission there is an admission that the "known worst-case reduction does not yield any meaningful 'end-to-end'
security guarantee for our concrete parameters based on the conjecture [sic] hardness of a worst-case problem". However, the "Justification of security strength" section incorrectly includes Theorem 5.8 in its list of "reductions supporting the security of FrodoKEM".

Formally, similar problems also exist in at least

   * Theorem 5.1 ("about"),
   * Theorem 5.2 ("about"),
   * Theorem 5.3 ("approximately"), and
   * Theorem 5.4 ("about").

One might guess that the quoted words are alluding specifically to polynomial slowdowns, and that these slowdowns aren't big problems, but a security audit requires these polynomials to be quantified.

---Dan

From: Mike Hamburg <mike@shiftleft.org>
Sent: Sunday, April 22, 2018 3:14 PM
To: D. J. Bernstein
Cc: pqc-comments; pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: Frodo

Hi Dan and other PQCers,

I don't agree that Theorem 5.2 is content-free, or that it doesn't qualitatively support the security of Frodo.

Evidence for security comes on a spectrum from tight proofs of practical security on the one side, down to semi-formal arguments that you didn't screw up on the other side. Theorem 5.2 is a semi-formal argument that the authors didn't screw up. Anything in the ROM or QROM must be taken with a grain of salt anyway, because computers don't come with magic black boxes instantiating a uniformly random function. Also we don't know whether SHAKE would preserve security against a quantum adversary even if Keccak-p were a uniformly random function.

You pointed out that theorem 5.2 is vacuous at 2^50 queries. But it's not vacuous at 2^20 or 2^30 queries, which would likely be enough for a Simon- or Shor-style attack, or if the CCA security mode just didn't work at all. What's more, as of mid last year, nobody knew how to prove anything tighter in the QROM, and theorem 5.1 establishes practical security of the CCA mode in the classical ROM.

Of course, it could be broken if the nightmare scenario obtains and there really is a 6th-degree loss in tightness from FO against a quantum adversary. But nobody expected this to happen, and as of just before submission time, we had the tools available to show that at least in some cases it doesn't happen. The techniques in [SXY'17], [JZCWM'17] should work for Frodo, and if they work they would narrow the gap to something like

q*(sqrt(Adv-IND-KPA) + sqrt(delta) + sqrt(1 / message space) + sqrt(1 / keyspace))

times a small constant, between maybe 2 and 12. I don't recall any slowdown in that method beyond a constant overhead per query, at least for single-target attacks. With the semiclassical O2H techniques I sketched last year (also in the ThreeBears writeup), the q sqrt(Adv-IND-KPA) term should improve to something like sqrt((q+1) Adv-IND-KPA). Again the slowdown should be constant per query for single-target attacks, but not for multiple-target attacks (it's something like an extra comparison per RO query per challenge ciphertext). Even without this improvement, the results would not be vacuous for 2^64 queries.

We don't have a proof yet that SXY or JZCWM applies to Frodo (but they should because it uses implicit rejection), and we don't have a proof that it is secure against multi-target attacks (but it does hash the public key as a countermeasure to such attacks). But coming back to the original chain of reasoning:

* Theorem 5.1 showed that Frodo's CCA transform is practically secure in the classical ROM.

* Theorem 5.2 showed that Frodo's CCA transform matched the best available in mid-2017, which had an enormous tightness loss and didn't prove practical security, but wasn't completely vacuous.

* The community expected, and SXY and JZCWM partially proved, that the tightness gap can typically be improved enough to prove some degree of practical security.

This is support for the security claims. So, at least for Theorem 5.2, the authors weren't bullshitting.

I am not nearly expert enough to comment on Theorem 5.8.

Regards,
— Mike

[JZCWM’17] https://eprint.iacr.org/2017/1096
[SXY’17] https://eprint.iacr.org/2017/1005.pdf

> On Apr 21, 2018, at 3:15 PM, D. J. Bernstein <djb@cr.yp.to> wrote:
>
> Summary: The "FrodoKEM" submission claims that various theorems support
> the security of the submission. This claim is incorrect for at least two
> of the stated theorems: these two theorems do not, in fact, support the
> security of the submission.
>
> The exact quote at issue is in Section 5.1:
>
>   5 Justification of security strength
>
>   The security of FrodoKEM is supported both by security reductions and
>   by analysis of the best known cryptanalytic attacks.
>
>   5.1 Security reductions
>
>   A summary of the reductions supporting the security of FrodoKEM is as
>   follows: ... Theorem 5.2 gives a non-tight, classical reduction
>   against quantum adversaries in the quantum random oracle model. ...
>   Theorem 5.8 gives a non-tight classical reduction against classical
>   or quantum adversaries (in the standard model).
>
> Qualitatively, the message of the two theorems is that the following two
> attack avenues are ruled out against FrodoKEM:
>
>   (1) generic-hash quantum chosen-ciphertext attacks that don't break
>   IND-CPA security of the underlying PKE;
>
>   (2) attacks against the relevant LWE parameters that don't break
>   various worst-case lattice problems.
>
> In fact, the theorems do _not_ rule out these attack avenues. A user of
> FrodoKEM is not guaranteed security against attacks of these two types.
> (Please note that I'm _not_ saying that attacks have been shown to exist
> along these two lines.) For #1 this is because the stated theorem is too
> loose to say anything regarding the KEMs that were actually submitted.
> For #2 the theorem statement lacks the quantification necessary to even
> see how loose the proof is; I'm skeptical of the idea that a tight
> enough theorem is possible, and in any case the claim of applicability
> of the stated theorem is incorrect.
>
> Terminology: It's important to be aware of a conflict between two
> different definitions of "KEM". For concrete literature such as
>
>   http://shoup.net/iso/std6.pdf
>
> a KEM has a concrete size, a concrete security level, etc. This is the
> definition that matters for standards, implementations, users, and this
> message. Most encryption submissions specify more than one KEM, but,

From: Jacob Alperin-Sheriff <jacobmas@gmail.com>
Sent: Sunday, April 22, 2018 4:37 PM
To: Mike Hamburg
Cc: D. J. Bernstein; pqc-comments; pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: Frodo

I haven't read the more recent paper by Peikert Regev and Stephens-Davidowitz nearly as carefully.

Also, I'm going to ignore the dispute over whether "asymptotic-only" reductions are still meaningful here.

Regarding Regev's original reduction and the Koblitz paper doing concrete calculations, there are a lot of things going on there that would be extraordinarily weird/unlikely for a "natural" algorithm breaking LWE to do in a manner that can't easily be "unwound" if you will.

Lemma 3.7 in Regev05 for instance would be utterly shocking in my view if an algorithm solving decision LWE for error rate beta could somehow not naturally differentiate between LWE with error rate beta/2 and uniform (minus the obvious way of a subroutine that solves it and then gives you a bad answer for smaller error because it hates you)

Actually there there may be a tighter reduction possible.

Anyway long story short I'd like to see an analysis of Regevs reduction that takes all of this into account.

On Apr 22, 2018 3:13 PM, "Mike Hamburg" <mike@shiftleft.org> wrote:

> Hi Dan and other PQCers,
>
> I don’t agree that Theorem 5.2 is content-free, or that it doesn’t qualitatively support the security of Frodo.
>
> Evidence for security comes on a spectrum from tight proofs of practical security on the one side, down to semi-formal arguments that you didn’t screw up on the other side. Theorem 5.2 is a semi-formal argument that the authors didn’t screw up. Anything in the ROM or QROM must be taken with a grain of salt anyway, because computers don’t come with magic black boxes instantiating a uniformly random function. Also we don’t know whether SHAKE would preserve security against a quantum adversary even if Keccak-p were a uniformly random function.
>
> You pointed out that theorem 5.2 is vacuous at 2^50 queries. But it’s not vacuous at 2^20 or 2^30 queries, which would likely be enough for a Simon- or Shor-style attack, or if the CCA security mode just didn’t work at all. What’s more, as of mid last year, nobody knew how to prove anything tighter in the QROM, and theorem 5.1 establishes practical security of the CCA mode in the classical ROM.
>
> Of course, it could be broken if the nightmare scenario obtains and there really is a 6th-degree loss in tightness from FO against a quantum adversary. But nobody expected this to happen, and as of just before submission time, we had the tools available to show that at least in some cases it doesn’t happen. The techniques in [SXY’17], [JZCWM’17] should work for Frodo, and if they work they would narrow the gap to something like
>
> q*(sqrt(Adv-IND-KPA) + sqrt(delta) + sqrt(1 / message space) + sqrt(1 / keyspace))
>
> times a small constant, between maybe 2 and 12. I don’t recall any slowdown in that method beyond a constant overhead per query, at least for single-target attacks. With the semiclassical O2H techniques I sketched last year (also in the ThreeBears writeup), the q sqrt(Adv-IND-KPA) term should improve to something like sqrt((q+1) Adv-IND-KPA). Again the slowdown should be constant per query for single-target attacks, but not for multiple-target attacks (it’s something like an extra comparison per RO query per challenge ciphertext). Even without this improvement, the results would not be vacuous for 2^64 queries.

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Sunday, April 22, 2018 5:05 PM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: Frodo
Attachments: signature.asc

If X is doable, then a theorem saying "An attack of this type must do X"
is content-free. It is not a theorem supporting security.

In the case of Theorem 5.2 from the Frodo submission, the doable thing is (under) 2^50 hash queries inside a quantum computation---far below NIST's minimum allowed security level. There's nothing in the theorem statement prohibiting 2^50 _parallel_ queries, so the latency limits permitted by NIST aren't relevant.

The followup comment dated 22 Apr 2018 12:13:51 -0700 seems to be arguing that a theorem supporting _security against cheap attacks_ should be of interest since some attacks in the literature are cheap.
Here's why I disagree: adopting this approach would leave people free to propose submissions

* that turn out to flunk NIST's security requirements, and
* that, if deployed, will presumably be broken by attackers, but
* that are in the meantime advertised as being "supported" by a proof of security, where
* the attack turns out to exploit exactly the fact that the proof is limited to _cheap_ attacks.

This bait-and-switch treatment of security levels is exactly what has led to, e.g., widespread deployment of the "provably secure" protocols broken in https://sweet32.info. The problem there wasn't a break in the proof hypotheses; the problem is that the security conclusion was only against cheap attacks.

I suppose the commenter would speculate at this point that the people advertising those proofs "weren't bullshitting".
Perhaps. What matters is that the resulting security levels were quantitatively unacceptable.
The theorems needed larger system parameters, and it was wrong to advertise them as supporting smaller system parameters. It was also wrong to wait for attacks to be demonstrated.

Aside from the core question of what it means for a theorem to support security, the followup comment makes several points that, as far as I can tell, aren't relevant to my comment:

* It "should" be possible to prove better theorems.

  Hopefully, yes, but so what? The Frodo submission states theorems and makes incorrect claims of applicability of those theorems. This is not the same as expressing hope of better theorems. The errors need to be corrected.

* Nobody knew how to do better in mid-2017.

  This seems accurate regarding the general shape of Theorem 5.2, but it's not accurate regarding the parameter selection---Frodo could have selected parameters big enough for the theorem to be meaningful.

More to the point, given that the parameters _aren't_ large enough, there's a big difference between saying "we don't know how to prove a theorem supporting security" and falsely saying "this theorem supports security".

- * There are other attack avenues ruled out by other theorems: e.g., Theorem 5.1 rules out generic-hash _non-quantum_ attacks that don't break IND-CPA security of the underlying PKE.

  It seems so, yes, but so what? The submission claims applicability of various theorems, and I'm pointing out errors in two specific claims. The accuracy of other claims isn't relevant.

- * There are other attack avenues that aren't ruled out by _any_ of the theorems, such as specific-hash attacks.

  Obviously, yes, but so what? My objection is to two incorrect claims that certain attack avenues _are_ ruled out.

---Dan

From: Mike Hamburg <mike@shiftleft.org>
Sent: Sunday, April 22, 2018 7:53 PM
To: D. J. Bernstein
Cc: pqc-comments; pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: Frodo

Hi Dan,

I think I don’t properly understand what you are arguing. Here is what I understand you to be saying; please correct me if I’m wrong about this.

* Proofs of security-against-cheap-attacks are meaningless if they don’t reach NIST’s required parameters. It doesn’t matter if such proofs are expected to be loose, or if they are the best available.

* Asymptotic proofs are meaningless in a concrete context. (Or “content-free”, if that means something different.)

* A design element motivated by asymptotic analysis, or motivated by a security proof that only rules out cheap attacks with the proposed parameters, is worthless.

* To knowingly claim otherwise is academic fraud, even if you are honest about what your proof actually says.

Is this correct? Because if so, I think we simply disagree at a philosophical level.

— Mike

> On Apr 22, 2018, at 2:04 PM, D. J. Bernstein <djb@cr.yp.to> wrote:
>
> If X is doable, then a theorem saying "An attack of this type must do X"
> is content-free. It is not a theorem supporting security.
>
> In the case of Theorem 5.2 from the Frodo submission, the doable thing
> is (under) 2^50 hash queries inside a quantum computation---far below
> NIST's minimum allowed security level. There's nothing in the theorem
> statement prohibiting 2^50 _parallel_ queries, so the latency limits
> permitted by NIST aren't relevant.
>
> The followup comment dated 22 Apr 2018 12:13:51 -0700 seems to be
> arguing that a theorem supporting _security against cheap attacks_
> should be of interest since some attacks in the literature are cheap.
> Here's why I disagree: adopting this approach would leave people free to
> propose submissions
>
>   * that turn out to flunk NIST's security requirements, and
>   * that, if deployed, will presumably be broken by attackers, but
>   * that are in the meantime advertised as being "supported" by a proof
>   of security, where
>   * the attack turns out to exploit exactly the fact that the proof is
>   limited to _cheap_ attacks.
>
> This bait-and-switch treatment of security levels is exactly what has
> led to, e.g., widespread deployment of the "provably secure" protocols

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Monday, April 23, 2018 11:42 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: Frodo
Attachments: signature.asc

The systems broken in https://sweet32.info are, obviously, not secure.
However, they have proofs of security. The hypotheses of the theorems (e.g., that the ciphers are PRP-secure) are not challenged by the breaks; the breaks are within the classes of attacks considered in the theorems; and no errors have been discovered in the proofs.

How is this possible? The answer is purely quantitative, a switch in security levels. The theorems are talking about 2^32 security, and what the users actually need is far beyond that.

Standardization organizations should never have allowed theorems talking about 2^32 security to be portrayed as saying something about real security. The quantitative security levels are important.

In the NIST PQC process, submissions are required to specify concrete parameters, and to define their security targets for those parameters.
The rules do _not_ allow security levels below the (conjectured) post-quantum difficulty of finding an AES-128 key.

Submissions can---after analysis---claim security against known attacks.
Submissions can also---after analysis---claim that theorems guarantee security against _all_ attacks of certain types.
However,

 * the word "security" is defined to be >=AES-128 security, so
 * if the analyses are actually substituting a substandard notion of
 "security" then that's an error that needs to be fixed.

The notion that (e.g.) 2^32 security is "support" for real security is dangerous for cryptographic users---again, see https://sweet32.info.
Similarly, Theorem 5.2 in the Frodo submission is below the minimum security level allowed in this process.

 [attempted restatement:]
> * Proofs of security-against-cheap-attacks are meaningless if they
> don’t reach NIST’s required parameters.

Meaningless for this NIST process, yes. Maybe such proofs would be meaningful for a lightweight pre-quantum competition, but that's not relevant here.

I'll refrain from commenting on issues of "design" and "fraud". The issue at hand is errors in security claims; one should be able to point out a dangerous conflation of two different notions of security without getting bogged down in questions of whether this was intentional.

---Dan

From: cpeikert@gmail.com on behalf of Christopher J Peikert <cpeikert@alum.mit.edu>
Sent: Wednesday, April 25, 2018 2:48 PM
To: pqc-forum; pqc-comments
Subject: Re: [pqc-forum] OFFICIAL COMMENT: Frodo

The FrodoKEM submission distinguishes between:

1. the freely parametrizable FrodoPKE/KEM constructions (Section 2.2), whose asymptotic security is indeed supported by a collection of tight and non-tight reductions (Section 5.1), and

2. the concrete instantiations FrodoKEM-640 and -976 (Section 2.4), whose concrete security is estimated by cryptanalysis, e.g., using the "core-SVP" methodology (Section 5.2).

(Please note that in Section 6.2 ("Compatibility with existing deployments"), the references to FrodoKEM can only make sense as referring to the concrete instantiations, but this should have been completely explicit to avoid any possibility of confusion. We believe there aren't any such ambiguities in the rest of the submission.)

Our approach, of starting from a parametrizable construction with asymptotic security supported by (possibly loose) reductions and then instantiating its parameters via cryptanalysis, is motivated and explained in detail in Section 1.2.2. Moreover, we explicitly disclaim any use of loose reductions as supporting the concrete security of our instantiations. For example, Section 1.2.2 says,

"We stress that we use the worst-case reduction only for guidance in
choosing a narrow enough error distribution for practice that still
has some theoretical support, and not for any concrete security claim.
... Instead, as stated in the above quote from [85], we choose
concrete parameters using a conservative analysis of the best known
cryptanalytic attacks, as described next."

Therefore, we believe there should not be any confusion about what the submission does and does not claim (and even disclaims) as justification for the concrete security of the FrodoKEM instantiations.

Sincerely, the FrodoKEM team

On Sat, Apr 21, 2018 at 6:15 PM, D. J. Bernstein <djb@cr.yp.to> wrote:
> Summary: The "FrodoKEM" submission claims that various theorems
> support the security of the submission. This claim is incorrect for at
> least two of the stated theorems: these two theorems do not, in fact,
> support the security of the submission.
>
> The exact quote at issue is in Section 5.1:
>
>   5 Justification of security strength
>
>   The security of FrodoKEM is supported both by security reductions and
>   by analysis of the best known cryptanalytic attacks.
>
>   5.1 Security reductions
>
>   A summary of the reductions supporting the security of FrodoKEM is as
>   follows: ... Theorem 5.2 gives a non-tight, classical reduction

From: Le Trieu Phong <letrieu.letrieuphong@gmail.com>
Sent: Wednesday, April 25, 2018 10:42 PM
To: Christopher J Peikert
Cc: pqc-forum; pqc-comments
Subject: Re: [pqc-forum] OFFICIAL COMMENT: Frodo

Dear FrodoKEM's submitters,

On page 31 of the specification (Theorem 5.6 on Renyi divergence and the paragraph below it), it is written that "During a single run of FrodoKEM the parties sample from the distribution (8+8)x640+64 = 10,304 times."

I cannot get the number 10,304 above correctly with my computations. Details are below:

- If one-way attack game is considered, then I think FrodoKEM.KeyGen should be run. Looking into FrodoKEM.KeyGen, I see two runs of Frodo.SampleMatrix of sizes n x \bar{n} and n x \bar{n}. Therefore I think the total samples should be n x \bar{n} + n x \bar{n} = 2n x \bar{n} = 2x640x8 = 10240, which does not match yours.
- If IND-CPA attack game is considered, then the samples in FrodoKem.Encaps should be additionally included for the challenge ciphertext. In FrodoKem.Encaps, there are 3 runs of Frodo.SampleMatrix of sizes \bar{m}xn + \bar{m}xn + \bar{m}x\bar{n} = \bar{m} x (2n + \bar{n}) = 8 x (2x640 + 8) = 10304. Therefore, the total samples should be 10240 (KeyGen) + 10304 (one encryption) = 20544.
- If IND-CCA attack game is considered, then the number of decryption queries should be additionally counted, as each decryption query will lead to Gaussian sampling as in Frodo.Decaps. Concretely, each decryption query will generate \bar{m} x (2n + \bar{n}) = 10304 samples so I think the total number of Gaussian samples should be:
 10240 (for KeyGen) + 20544 (one encryption for the challenge ciphertext) + 10304*q_D (decryption queries)
where q_D is the number of decryption queries.

Thank you in advance for your explanations and best regards,
Phong

On Thu, Apr 26, 2018 at 3:47 AM, Christopher J Peikert <cpeikert@alum.mit.edu> wrote:

> The FrodoKEM submission distinguishes between:
>
> 1. the freely parametrizable FrodoPKE/KEM constructions (Section 2.2), whose asymptotic security is indeed supported by a collection of tight and non-tight reductions (Section 5.1), and
>
> 2. the concrete instantiations FrodoKEM-640 and -976 (Section 2.4), whose concrete security is estimated by cryptanalysis, e.g., using the "core-SVP" methodology (Section 5.2).
>
> (Please note that in Section 6.2 ("Compatibility with existing deployments"), the references to FrodoKEM can only make sense as referring to the concrete instantiations, but this should have been completely explicit to avoid any possibility of confusion. We believe there aren't any such ambiguities in the rest of the submission.)
>
> Our approach, of starting from a parametrizable construction with asymptotic security supported by (possibly loose) reductions and then

From: Le Trieu Phong <letrieu.letrieuphong@gmail.com>
Sent: Thursday, April 26, 2018 12:00 AM
To: pqc-forum
Cc: pqc-comments
Subject: Re: [pqc-forum] OFFICIAL COMMENT: Frodo

I would revise the final computation as:

10240 (for KeyGen) + 10304 (one encryption for the challenge ciphertext) + 10304*q_D (decryption queries)

Thank you!

On Thu, Apr 26, 2018 at 11:41 AM, Le Trieu Phong <letrieu.letrieuphong@gmail.com> wrote:

> Dear FrodoKEM's submitters,
>
> On page 31 of the specification (Theorem 5.6 on Renyi divergence and the paragraph below it), it is written that "During a single run of FrodoKEM the parties sample from the distribution (8+8)x640+64 = 10,304 times."
>
> I cannot get the number 10,304 above correctly with my computations. Details are below:
>
> - If one-way attack game is considered, then I think FrodoKEM.KeyGen should be run. Looking into FrodoKEM.KeyGen, I see two runs of Frodo.SampleMatrix of sizes n x \bar{n} and n x \bar{n}. Therefore I think the total samples should be n x \bar{n} + n x \bar{n} = 2n x \bar{n} = 2x640x8 = 10240, which does not match yours.
> - If IND-CPA attack game is considered, then the samples in FrodoKem.Encaps should be additionally included for the challenge ciphertext. In FrodoKem.Encaps, there are 3 runs of Frodo.SampleMatrix of sizes \bar{m}xn + \bar{m}xn + \bar{m}x\bar{n} = \bar{m} x (2n + \bar{n}) = 8 x (2x640 + 8) = 10304. Therefore, the total samples should be 10240 (KeyGen) + 10304 (one encryption) = 20544.
> - If IND-CCA attack game is considered, then the number of decryption queries should be additionally counted, as each decryption query will lead to Gaussian sampling as in Frodo.Decaps. Concretely, each decryption query will generate \bar{m} x (2n + \bar{n}) = 10304 samples so I think the total number of Gaussian samples should be:
> 10240 (for KeyGen) + 20544 (one encryption for the challenge ciphertext) + 10304*q_D (decryption queries)
> where q_D is the number of decryption queries.
>
> Thank you in advance for your explanations and best regards,
> Phong
>
> On Thu, Apr 26, 2018 at 3:47 AM, Christopher J Peikert <cpeikert@alum.mit.edu> wrote:
>
>> The FrodoKEM submission distinguishes between:
>>
>> 1. the freely parametrizable FrodoPKE/KEM constructions (Section 2.2), whose asymptotic security is indeed supported by a collection of tight and non-tight reductions (Section 5.1), and
>>
>> 2. the concrete instantiations FrodoKEM-640 and -976 (Section 2.4), whose concrete security is estimated by cryptanalysis, e.g.,