
From: Damien Stehlé <damien.stehle@gmail.com>
Sent: Friday, May 03, 2019 3:10 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: ROUND 2 OFFICIAL COMMENT: NTRU Prime

Hello NTRUPrime team,

It seems that the observation that Jan-Pieter D'Anvers made on Kyber during Round 1 (Jan 17, 2018) also applies to NTRU LPrime.

By rounding the second ring element in the public key, one cannot directly rely on the Ring-LWE hardness assumption to argue security of the encryption procedure.

With your terminology, this means Problem 3 (defined on p.36) does not perfectly model the problem of finding a plaintext from a public key and a ciphertext, in the NTRU LPrime instantiation of Product NTRU. The modelling would be more adequate if the ring element A_2 (in Problem 3) was restricted to be a multiple of 3.

Not sure what implications this remark has, though.

Best regards
Damien

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Friday, May 03, 2019 5:55 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] ROUND 2 OFFICIAL COMMENT: NTRU Prime
Attachments: signature.asc

Speaking for myself.

Damien Stehlé writes:

> It seems that the observation that Jan-Pieter D'Anvers made on Kyber
> during Round 1 (Jan 17, 2018) also applies to NTRU LPRime.

No. D'Anvers was pointing out an error in a claimed Module-LWE reduction for round-1 Kyber (fixed in round-2 Kyber, if I understand correctly), specifically a mistaken claim of uniformity of multipliers. There is no such mistake in NTRU LPRime (or in Streamlined NTRU Prime).

> By rounding the second ring element in the public key, one cannot
> directly rely on the Ring-LWE hardness assumption to argue security of
> the encryption procedure.

What exactly do you think you're disputing? NTRU Prime has never claimed Ring-LWE reductions.

> With your terminology, this means Problem 3 (defined on p.36) does not
> perfectly model the problem of finding a plaintext from a public key
> and a ciphertext, in the NTRU LPRime instantiation of Product NTRU.

The submission already says that Problem 3 isn't a perfect model. Here's the relevant quote, immediately after the models are introduced:

There are various reasons that these models could be underestimating or overestimating security. For example: ...

There are then five examples covering twenty lines. (There's no claim that the list is complete; on the contrary, the words "for example" are the opposite of claiming completeness.)

Simple models can be helpful as targets for cryptanalysis, but it's important (for NTRU Prime and for every other lattice submission) to keep in mind that attacks against the cryptosystems don't necessarily match attacks in the models. The first page of the introduction of the NTRU Prime submission lists "serious risks" including "algorithms to break cryptosystems without breaking these problems".

I think it's unfortunate that lattice-based crypto has built a tradition of hyping the limited things that have been proven, rather than warning users about the huge remaining attack surface.

---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.

Visit this group at <https://groups.google.com/a/list.nist.gov/group/pqc-forum/>.

From: 'daniel.apon' via pqc-forum <pqc-forum@list.nist.gov>
Sent: Friday, May 03, 2019 1:25 PM
To: pqc-forum
Cc: pqc-comments
Subject: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: NTRU Prime

Hi Damien,

Two comments:

1. Your observation (as I read it) that, *if* the goal were to reduce security of NTRU Prime to Ring-LWE, Problem 3 would need to be modified slightly (or some similar change made) is correct.

Generally speaking, the largest benefit in doing so would be to use Ring-LWE (or in the algebraically-unstructured case, LWE) problem as a stepping stone toward provably basing security of the cryptosystem on the hardness of finding approximately short vectors in (structured/unstructured) lattices.

But for every lattice-based KEM submission to NIST, the concrete parameters are chosen so that the various, theoretical reductions do not actually apply. This is true, for example, of Kyber when considering a (hypothetical, additional) reduction between the relevant forms of Module-LWR and Module-LWE and a (hypothetical) reduction between their chosen form of Module-LWE and the problem of finding approximate-shortest vectors in the associated lattices. The same would hold true for NTRU Prime, given its parameter choices, if one independently attempted a substantially similar line to generate a proof of security; the NTRU Prime Team did not attempt such a proof. (We would need to wait for an official comment from the NTRU Prime Team for any further clarification on the Team's motivations for not attempting such a proof.)

For clarity's sake -- in the case of Kyber, the cleanest statement along these lines that I'm aware of may be found in their conference proceeding <https://eprint.iacr.org/2017/634.pdf>, right column, page 6. That is, *"...yet for [parameters] used in practice (c.f. [10] and many submissions to the NIST post-quantum call), it is still assumed that the distribution of Module-LWR is pseudorandom despite the fact that the proof in [18] is no longer applicable."* Specifically, the extra assumption is (always?) made that some form of lattice reduction is still the best, 'direct' attack method (despite the lack of formal proof that this is necessarily the case), so concrete analysis is done with respect to lattice reduction algorithms anyway.

The point being: There are larger obstacles to explicitly deriving provable security (from the hardness of finding approximate shortest vectors in lattices) for the lattice-based submissions than this one issue.

As an exercise, you could try estimating the performance numbers for a Plain/Ring/Module/etc-LWE system that indeed passes all the way through Regev's reduction (or related, such as PRS17) in a theoretically-sound manner, but the resulting scheme(s) can be seen to be non-competitive in practice.

2. And now, ignoring the above (mostly tangential) digression into Ring-LWE, in order to get at what appears to be the point of your post:

"The modelling would be more adequate if the ring element A_2 (in Problem 3) was restricted to be a multiple of 3."

Yes, that's apparent. But the security implications are minimal (again, see the discussion in Kyber's conference paper).

Hope this helps,
--Daniel Apon

From: Damien Stehlé <damien.stehle@gmail.com>
Sent: Friday, May 03, 2019 2:27 PM
To: Apon, Daniel C. (Fed)
Cc: pqc-forum; pqc-comments
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: NTRU Prime

Hi Dan, hi Daniel,

> What exactly do you think you're disputing? NTRU Prime has never
> claimed Ring-LWE reductions.

What makes you think I am disputing something?
I must have been unclear in my email.

To clarify, if need be: I do not claim that you claimed a Ring-LWE reduction.
This is actually why I focused on "Problem 3".

The choice of name NTRU ****LPR****ime, though, may induce a user to think that the scheme may enjoy the same Ring-LWE security "proof" as the LPR encryption scheme (even if you do not claim it). I plead guilty for having thought this may be the case, for a short while.

I hope that my email helps clarifying, for other potentially careless readers, that NTRU LPRime does not enjoy a security proof that is analogous to that of the LPR scheme.

> The submission already says that Problem 3 isn't a perfect model.
> Here's the relevant quote, immediately after the models are introduced:
>
> There are various reasons that these models could be underestimating
> or overestimating security. For example: ...
>
> There are then five examples covering twenty lines. (There's no claim
> that the list is complete; on the contrary, the words "for example"
> are the opposite of claiming completeness.) Simple models can be
> helpful as targets for cryptanalysis, but it's important (for NTRU
> Prime and for every other lattice submission) to keep in mind that
> attacks against the cryptosystems don't necessarily match attacks in
> the models.

I like models too. Explaining the limits of models is also good.
Let's say that my contribution is yet another example of a discrepancy between the model and the cryptosystem.

> I think it's unfortunate that lattice-based crypto has built a
> tradition of hyping the limited things that have been proven, rather
> than warning users about the huge remaining attack surface.

My email goes into the direction of assessing the attack surface of a lattice-based crypto scheme. We are on the same page.

> Generally speaking, the largest benefit in doing so would be to use Ring-LWE (or in the algebraically-unstructured case, LWE) problem as a stepping stone toward provably basing security of the cryptosystem on the hardness of finding approximately short vectors in (structured/unstructured) lattices.

There seems to be some misunderstanding: I was not referring to worst-case lattice problems. Just the connection between the scheme security and (a potential variant of) Ring-LWE/Problem 3. In Ring-LWE/Ring-LWR samples (a_i, a_i*s+e_i) , the ring elements a_i are uniform. I am merely saying, in fact recalling D'Anvers' point, that a rounded a_i (e.g., obtained by using LWR in the key generation) is not uniform.

Best regards
Damien

On Fri, 3 May 2019 at 19:24, 'daniel.apon' via pqc-forum <pqc-forum@list.nist.gov> wrote:

>
> Hi Damien,
>
> Two comments:
>
> 1. Your observation (as I read it) that, if the goal were to reduce security of NTRU Prime to Ring-LWE, Problem 3 would need to be modified slightly (or some similar change made) is correct.
> Generally speaking, the largest benefit in doing so would be to use Ring-LWE (or in the algebraically-unstructured case, LWE) problem as a stepping stone toward provably basing security of the cryptosystem on the hardness of finding approximately short vectors in (structured/unstructured) lattices.
>
> But for every lattice-based KEM submission to NIST, the concrete
> parameters are chosen so that the various, theoretical reductions do
> not actually apply. This is true, for example, of Kyber when
> considering a (hypothetical, additional) reduction between the
> relevant forms of Module-LWR and Module-LWE and a (hypothetical)
> reduction between their chosen form of Module-LWE and the problem of
> finding approximate-shortest vectors in the associated lattices. The
> same would hold true for NTRU Prime, given its parameter choices, if
> one independently attempted a substantially similar line to generate a
> proof of security; the NTRU Prime Team did not attempt such a proof.
> (We would need to wait for an official comment from the NTRU Prime
> Team for any further clarification on the Team's motivations for not
> attempting such a proof.)
>
> For clarity's sake -- in the case of Kyber, the cleanest statement
> along these lines that I'm aware of may be found in their conference proceeding
<https://gcc01.safelinks.protection.outlook.com/?url=https%3A%2F%2Fprint.iacr.org%2F2017%2F634.pdf&data=02%7C01%7Csara.kerman%40nist.gov%7Cf62a7c2826164f32c86b08d6cff502a3%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C636925048549685209&reserved=0>, right column, page 6. That is, "...yet for [parameters] used in practice (c.f. [10] and many submissions to the NIST post-quantum call), it is still assumed that the distribution of Module-LWR is pseudorandom despite the fact that the proof in [18] is no longer applicable."
> Specifically, the extra assumption is (always?) made that some form of lattice reduction is still the best, 'direct' attack method (despite the lack of formal proof that this is necessarily the case), so concrete analysis is done with respect to lattice reduction algorithms anyway.
>
> The point being: There are larger obstacles to explicitly deriving provable security (from the hardness of finding approximate shortest vectors in lattices) for the lattice-based submissions than this one issue.
>
> As an exercise, you could try estimating the performance numbers for a Plain/Ring/Module/etc-LWE system that indeed passes all the way through Regev's reduction (or related, such as PRS17) in a theoretically-sound manner, but the resulting scheme(s) can be seen to be non-competitive in practice.
>

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Saturday, May 04, 2019 3:00 PM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: NTRU Prime
Attachments: signature.asc

Damien Stehlé writes:

> NTRU LPRime does not enjoy a security proof that is analogous to that
> of the LPR scheme.

I can't figure out the intended meaning of this statement. Can you please clarify?

What features would a proof need to have to qualify as "analogous"? Is "enjoy" simply a judgment call about those features, or is it adding extra constraints? Examples of more specific questions that a clarification should easily answer:

- * You began by writing "rely on the Ring-LWE hardness assumption". Are you excluding problems that eliminate random errors in favor of deterministic rounding, such as the problems inside NTRU LPRime, Round5, Saber, and Streamlined NTRU Prime?
- * Are you excluding the problems inside Quotient NTRU systems, such as the NTRU proposal and Streamlined NTRU Prime?
- * Would you allow a system that releases many overstretched Ring-LWE samples, given that the "Ring-LWE" name includes such problems?

I presume that you have in mind some sort of selection of "analogous" lattice problems, based on some argument that this selection reduces security risks; but both the selection and the argument are unclear.

The Lyubashevsky--Peikert--Regev paper

<https://eprint.iacr.org/2012/230>

was supposedly "proving" that Ring-LWE "enjoys very strong hardness guarantees", namely a theorem relating Ring-LWE to an approximate Ideal-SVP problem. However, a closer look shows that this "very strong hardness guarantee" provides little assurance for security reviewers:

- * The theorem is quantitatively very loose. Even if the approximate Ideal-SVP problem is imagined to be as hard as the best attacks known in 2012, the looseness makes the theorem inapplicable to serious proposals of parameters for lattice-based cryptography.
- * A few cryptanalysts have been attacking the approximate Ideal-SVP problem in the last few years, and have found ways to exploit the structure of the problem---especially the extra automorphisms provided by the usual fields---to do surprising levels of damage.

At some point the claim of a "very strong hardness guarantee" has to succumb to the facts: this "guarantee" starts from hypotheses that haven't been holding up well against cryptanalysis, and then draws security conclusions about irrelevant cryptosystem parameters.

Are you asking for a proof sharing some features with this "guarantee"?

If so, what are the features? Would you also allow a proof based upon Ring-LWR, since---for irrelevant cryptosystem parameters---Ring-LWR is provably as hard as Ring-LWE? How about a proof based upon the "NTRU problem", which---for irrelevant cryptosystem parameters---is proven hard in your paper <https://eprint.iacr.org/2013/004>? If not, why not?

You've recently been writing things like

... strongly suggests that approx-SVP for ideals ... may be weaker than Ring-LWE, for a vast family of number fields ...

It's hard to see how this is compatible with putting any reliance upon the LPR "guarantees". But then what makes you think that Ring-LWE is safer than other lattice problems? You also now write

I was not referring to worst-case lattice problems.

But then what sort of proof are you asking for? Surely you aren't talking about the generic observation that

- (1) one-wayness for the system's keys and ciphertexts follows from
- (2) presumed indistinguishability of the keys from random and
- (3) presumed one-wayness for ciphertexts for random keys,

which has nothing to do with the specific structure of Ring-LWE, or with the choice of distribution that parameterizes the word "random".

Maybe you're alluding to search-to-decision reductions. However, these still aren't tight enough for a careful security reviewer. Cryptanalysis papers are attacking lattice one-wayness problems much more often than distinguishing problems, so it's worrisome for the security reviewer to see a cryptosystem that needs to make indistinguishability assumptions.

---Dan (speaking for myself)

From: daniel.apon <daniel.apon@nist.gov>
Sent: Sunday, May 05, 2019 3:48 PM
To: pqc-forum
Cc: Apon, Daniel C. (Fed); pqc-comments
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: NTRU Prime

"I am merely saying, in fact recalling D'Anvers' point, that a rounded a_i (e.g., obtained by using LWR in the key generation) is not uniform.

*Best regards
Damien "*

Hi Damien,

Yes, this is true.

Focusing more on this particular issue that you've brought up:

Note that an avenue to "fix" this issue, which was explored by the Kyber team (but ignored in the end), is to re-randomize by adding in some fresh error term e' .

An analogous "fix" (whether this constitutes actually fixing a true issue or not) in the case of NTRU LRPrime would be to consider a Problem 3* which has a term (A_2+e') rather than A_2 on its own (and in tandem, modify the scheme accordingly).

Two comments are in order:

- 1) Of first importance, it's unclear at present that this "fix," in full context, actually improves security; perhaps, it even hurts security. (One could note that the Kyber team ignored this possible avenue, for the reasons their various publications mention.)
- 2) Secondly, such a "fix" would actually pose problems in the context of NTRU Prime, as (at least, applied naively) it would induce decryption failures with some small probability. (This is less a concern when a scheme moves from $\sim 2^{-140}$ to $\sim 2^{-120}$ failure rate, but more of a concern in the case of the NTRU LRPrime scheme, which does not have decryption failures as-is.)

Regarding comment (1) above: *asymptotically speaking*, it's clear that this has no security implication whatsoever. (To see this, one could consider e.g. the "coset sampleable" framework of <https://eprint.iacr.org/2015/220>.)

But in more concrete terms -- in particular as the issue arises in the context of lattice-based submissions to NIST -- the question is essentially unanswered.

Is the fact that an adversary can explicitly view the rounded public-key component as non-uniform (in either of these schemes) a security vulnerability?

On the face of things, it seems unlikely, though it's not impossible.

Best,
--Daniel

On Friday, May 3, 2019 at 2:27:28 PM UTC-4, Damien Stehlé wrote:

Hi Dan, hi Daniel,

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Sunday, May 05, 2019 7:08 PM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: NTRU Prime
Attachments: signature.asc

Daniel Apon writes:

> But in more concrete terms -- in particular as the issue arises in the
> context of lattice-based submissions to NIST -- the question is
> essentially unanswered.

To be clear, are you talking about the question of how the key distribution might affect an attack against $2n$ fully released samples, as in the LPR cryptosystem? Or are you talking about the question of how the key distribution might affect an attack against

- * n fully released samples each having ~ 10 bits of information and
- * 256 samples each having only ~ 4 bits of information,

as in the NTRU LPrime cryptosystem and (modulo numerical details) other modern variants of LPR?

Trying to use the 4-bit samples in place of full-width samples causes a severe problem for known attacks, because the error is much bigger. Note that the gap here, viewed as information per element of F_q , is also much bigger than the gap between rounded and unrounded elements.

For some systems, the "Estimate" numbers claim that attacks benefit from going beyond n samples---usually only one or two bits of security loss, but occasionally more. This claim is based on the assumption that more samples are fully released. The actual systems release fewer bits past n samples. (This is mentioned in the NTRU Prime submission, along with many other reasons that the estimates shouldn't be taken as gospel.)

For systems where we don't even know how to fully exploit n samples (such as NTRU LPrime with the proposed parameters---only about 80% of the first n samples are used), surely it's a higher priority to think about how attacks might exploit the unused full-width samples than to think about how attacks might exploit the 4-bit samples.

Of course there are reasons to think that releasing more samples is dangerous. One nightmare scenario for LPR and its derivatives (Frodo, Kyber, LAC, NewHope, NTRU LPrime, Round5, Saber; also ThreeBears if that's counted as a lattice submission) would be that going beyond n samples, even with only a few bits per sample, immediately starts damaging security. Surely the best defense against an attack of this type wouldn't be to play with the key distribution, but rather to release fewer samples, as in NTRU and Streamlined NTRU Prime.

As a side note, this isn't the only potential problem specific to the LPR-based systems. For example, there could be an attack exploiting FO derandomization (presumably combining symmetric and asymmetric attack techniques---how many people have the right experience for this?). Known derandomization proofs from one-wayness aren't tight (even if all attacks are assumed to be QROM attacks), and this could reflect a real security problem.

---Dan (speaking for myself)

From: daniel.apon <daniel.apon@nist.gov>
Sent: Sunday, May 05, 2019 10:14 PM
To: pqc-forum
Cc: pqc-comments; djb@cr.yp.to
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: NTRU Prime

Hi Dan,

"To be clear, are you talking about [X], or are you talking about [Y]?"

I was speaking about neither option explicitly -- other than what anyone might independently derive from the relation between the public parameter and the secret.

This relation is clearly not a trivial issue, but it was not the subject of my comment to Damien.

In particular, Damien's comment *"The modelling would be more adequate if the ring element A_2 (in Problem 3) was restricted to be a multiple of 3."* stands on its own.

That said, my intention in speaking further was to emphasize that the asymptotic argument of coset-sampling due to 2015/220 -- which resolves this question of multiples-of-a-constant-times-the-public-parameter -- indeed fails in the case of every lattice-based submission to NIST. (We appear to agree that any "rounding-like" lattice KEM submitted to NIST seems to share this property.) Note that the reason 2015/220 is relevant is that it is the best-case, theoretical treatment of this issue (and related issues) that is published to-date.

Is there an equivalent *proof* that is as strong or general as Boneh et al. in the case of the concrete-parameterized lattice cryptosystems submitted to NIST?

(We encourage the community to examine this issue in further detail.)

"Of course there are reasons to think that releasing more samples is dangerous."

We seem to be far into a fresh topic now.

This comment -- taken as such -- presupposes that all samples are equal (or close enough to equal).

But given that: Yes, of course.

NIST is eager to support the community in its search of answers to these issues.

(Note that a formal proof, which is concretely and explicitly relevant, is obviously the most desired outcome, as such a proof would end all debate. This, of course, may be asking too much.)

Cheers,
--Daniel Apon

On Sunday, May 5, 2019 at 7:08:25 PM UTC-4, D. J. Bernstein wrote:

Daniel Apon writes:

> But in more concrete terms -- in particular as the issue arises in the
> context of lattice-based submissions to NIST -- the question is
> essentially unanswered.

To be clear, are you talking about the question of how the key

From: Damien Stehlé <damien.stehle@gmail.com>
Sent: Monday, May 06, 2019 3:20 AM
To: Apon, Daniel C. (Fed)
Cc: pqc-forum; pqc-comments; D. J. Bernstein
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: NTRU Prime

Hi Dan, hi Daniel,

(From Dan, May 4)

>> NTRU LPRime does not enjoy a security proof that is analogous to
>> that of the LPR scheme.

>

> I can't figure out the intended meaning of this statement. Can you
> please clarify?

My two emails were clear, I believe, about what I am discussing:

the difficulty of proving the security of NTRU LPRime based under some type of Ring-LWE assumption, due to the fact that the encryption procedure uses "LWE left hand sides" that are multiples of 3 (because of the rounding in the key generation procedure).

Merely following an old thread from D'Anvers on a similar issue with Kyber. As you pointed out, the observation was stronger for Kyber, as Kyber claimed such a reduction. NTRU *LPR*ime does not. Though the name is a bit misleading in that respect, but that is only my personal opinion. Indeed, the main point of the LPR encryption scheme was to have a proof under the RingLWE hardness assumption.

> What features would [...]

I am not going to discuss about topics that were not covered by the email of mine that started this thread. If you are interested in such topics, then please start another thread.

I am just extracting this from your May 4th email, as it might be relevant to this thread:

> [...] proof based upon Ring-LWR [...]

Maybe the specific noise distribution makes a difference for the problem under scope. I don't really see how, but why not.

(from Daniel, May 5)

> An analogous "fix" (whether this constitutes actually fixing a true
> issue or not) in the case of NTRU LRPrime would be to consider a
> Problem 3* which has a term $(A_2 + e)$ rather than A_2 on its own (and
> in tandem, modify the scheme accordingly).

>

> Two comments are in order:

>

> 1) Of first importance, it's unclear at present that this "fix," in
> full context, actually improves security; perhaps, it even hurts security.
> (One could note that the Kyber team ignored this possible avenue, for
> the reasons their various publications mention.)

> 2) Secondly, such a "fix" would actually pose problems in the
> context of NTRU Prime, as (at least, applied naively) it would induce

> decryption failures with some small probability. (This is less a
> concern when a scheme moves from $\sim 2^{-140}$ to $\sim 2^{-120}$ failure rate,
> but more of a concern in the case of the NTRU LPRime scheme, which
> does not have decryption failures as-is.)

I globally agree. On 1), speaking for myself, I would say the reason we (Kyber) didn't opt for this option was its cost compared to the other option that we chose. In particular, you have to be careful about how to sample e' so that $e'+A_2$ is uniform.

On 2), I did not check the calculations, but it is plausible. Surely the NTRUPrime team could have a more educated answer on this than me.

> Regarding comment (1) above: asymptotically speaking, it's clear that
> this has no security implication whatsoever. (To see this, one could
> consider e.g. the "coset sampleable" framework of
> <https://gcc01.safelinks.protection.outlook.com/?url=https%3A%2F%2Fepri.nt.iacr.org%2F2015%2F220&data=02%7C01%7Csara.kerman%40nist.gov%7C20344888e2c3468fc9aa08d6d1f35872%7C2ab5d82fd8fa4797a93e054655c61dec%7C1%7C0%7C636927240424025771&sdata=l5vfZC%2FY0eqDBAh5MBXMeg2ILIOHD3ZlRbbtFJzKPE%3D&reserved=0>.) But in more concrete terms -- in
> particular as the issue arises in the context of lattice-based
> submissions to NIST -- the question is essentially unanswered.

Nice idea. It's not 'clear' to me that it is going to work well, though.
But it may work. I see at least two non-trivial obstacles in the present context: we have a single ring element (as opposed to a module-LWE approach), and the secret s is (very) small.

> Is the fact that an adversary can explicitly view the rounded
> public-key component as non-uniform (in either of these schemes) a
> security vulnerability?
> On the face of things, it seems unlikely, though it's not impossible.

I tend to agree. To me, it's the main conceptual discrepancy with an LPR-like encryption scheme, but it may indeed not be relevant for actual security.

(From Dan, May 6)

> To be clear, are you talking about the question of how the key
> distribution might affect an attack against $2n$ fully released
> samples, as in the LPR cryptosystem? Or are you talking about the
> question of how the key distribution might affect an attack against
>
> * n fully released samples each having ~ 10 bits of information and
> * 256 samples each having only ~ 4 bits of information,
>
> as in the NTRU LPRime cryptosystem and (modulo numerical details)
> other modern variants of LPR?
>
> Trying to use the 4-bit samples in place of full-width samples causes
> a severe problem for known attacks, because the error is much bigger.
> Note that the gap here, viewed as information per element of F_q , is
> also much bigger than the gap between rounded and unrounded elements.

If I may try to re-state, my understanding is that you are saying that this part of the ciphertext is rounded so much anyway, that this property of the second public-key component does not matter for concrete security.

I agree that it is not unlikely.

Best regards

Damien

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Monday, May 06, 2019 9:27 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: NTRU Prime
Attachments: signature.asc

Damien Stehlé writes:

> some type of Ring-LWE assumption

So, to be clear, the dividing line is whether a hardness assumption fits within the parameter space of the LPR definition of "Ring-LWE"? E.g.:

- * A proof for anything like Round5 or Saber would not qualify as "enjoy a security proof that is analogous", because the starting assumption is hardness of Ring-LWR/Module-LWR instead of Ring-LWE?
- * If a scheme releases many overstretched Ring-LWE samples, and has a proof based on the alleged hardness of this case of Ring-LWE, then this would qualify as "enjoy a security proof that is analogous"?

This sounds very strange. Cryptanalysis indicates that the second example is much more dangerous than the first.

Are you sure this is what you meant by "enjoy a security proof that is analogous to that of the LPR scheme"?

> due to the fact that the encryption procedure uses "LWE left hand
> sides" that are multiples of 3 (because of the rounding in the key
> generation procedure).

If you're really insisting on specifically Ring-LWE and not anything rounding-based, then you're excluding all of the proposed LPR variants that eliminate random errors in favor of deterministic rounding, such as Round5 and Saber. This isn't specific to NTRU LPRime, and the main technical point isn't new: the literature already makes clear that "Ring-LWE hard => Ring-LWR hard" relies on irrelevant parameter choices.

Tweaking the key distribution in these pure rounding schemes doesn't enable a proof from Ring-LWE, so I don't understand why you attribute your ad-hoc exclusion of NTRU LPRime to details of the key distribution.

If I wanted to write down a list of proof requirements that allows Round5 and Saber (and round-2 Kyber etc.) while disallowing NTRU LPRime, I think I'd be able to, but stating the list clearly would also show how unprincipled and artificial it is.

> Indeed, the main point of the LPR encryption scheme was to have a
> proof under the RingLWE hardness assumption.

What the LPR paper actually says that it is (1) "introducing an algebraic variant of LWE" and (2) "proving that it too enjoys very strong hardness guarantees. Specifically, we show that the ring-LWE distribution is pseudorandom, assuming that worst-case problems on ideal lattices are hard for polynomial-time quantum algorithms."

A closer look shows that these "guarantees" are for irrelevant cryptosystem parameters. (I think <https://eprint.iacr.org/2016/360> deserves the primary credit for this observation.) But then why should Ring-LWE be treated as better than

- * Ring-LWR, which has a proof "Ring-LWE hard => Ring LWR hard" for irrelevant parameters;
- * the "NTRU problem", which has a hardness proof for irrelevant parameters;
- * a rounded-multiplier variant of Ring-LWE, which has a hardness proof for irrelevant parameters; and
- * a rounded-multiplier variant of Ring-LWR, which has a hardness proof for irrelevant parameters?

I already asked for clarification of whether your "analogous" claim would allow the first and second assumptions. You didn't answer. How about the third and fourth assumptions?

(Disclaimer: I'm relying on claims that I've heard about these proofs. I'm not saying that I've verified the proofs and theorem statements. I try to keep my proof-verification efforts focused on proofs applicable to relevant cryptosystem parameters.)

The underlying "worst-case problems on ideal lattices" have not been holding up well to cryptanalysis. My understanding is that, for this reason, you no longer endorse relying on the LPR "guarantee". But, without this "guarantee", essentially nothing is left of LPR's cryptosystem "proof"---it's simply the generic observation that

- (1) one-wayness for the system's keys and ciphertexts follows from
- (2) presumed indistinguishability of the keys from random and
- (3) presumed one-wayness for ciphertexts for random keys,

modulo generic replacement of OW with IND. As I said before, this has nothing to do with the specific structure of Ring-LWE, or with the choice of distribution that parameterizes the word "random". Do you not agree that all submissions "enjoy" such a proof?

> I am not going to discuss about topics that were not covered by the
> email of mine that started this thread.

You filed an "OFFICIAL COMMENT" dated 3 May 2019 20:27:01 +0200 claiming, among other things, that "NTRU LPRime does not enjoy a security proof that is analogous to that of the LPR scheme".

The intended meaning of this claim is not clear. I have asked various clarification questions, and I would like to hear your answers. If these questions make you realize that the statement you made doesn't actually have a meaning that you endorse, then you can withdraw the statement.

Of course the withdrawal itself needs to be specific and clear!

---Dan (speaking for myself)

From: Damien Stehlé <damien.stehle@gmail.com>
Sent: Tuesday, May 07, 2019 2:18 AM
To: pqc-forum; pqc-comments
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: NTRU Prime

Dear Dan,

> You filed an "OFFICIAL COMMENT" dated 3 May 2019 20:27:01 +0200
> claiming, among other things, that "NTRU LPrime does not enjoy a
> security proof that is analogous to that of the LPR scheme".

> The intended meaning of this claim is not clear. I have asked various
> clarification questions, and I would like to hear your answers. If
> these questions make you realize that the statement you made doesn't
> actually have a meaning that you endorse, then you can withdraw the statement.
> Of course the withdrawal itself needs to be specific and clear!

I've already clarified the meaning of the claim. Again, from context, it is clear that it is related to the coefficients of A2 being multiples of 3 (aka rounded multipliers).

I have the impression that you are asking me to give a value judgement to the sentence. Or to state a value judgement that I might have. I am not going to, beyond my original value judgement:

> Not sure what implications this remark has, though.

I will write more at a later stage if I have more to say on the topic, at the moment, it is not the case.

I am not going to continue going in circles like this. You might consider replying to the technical observation at hand (if you have more to say about it, that is). Or start another thread with the numerous questions that you seem to have. I may or may not give my viewpoint in this other discussion, if I feel interested and if I have time.

Best regards
Damien

Damien

On Mon, 6 May 2019 at 15:27, D. J. Bernstein <djb@cr.yp.to> wrote:

>
> Damien Stehlé writes:
>> some type of Ring-LWE assumption
>
> So, to be clear, the dividing line is whether a hardness assumption
> fits within the parameter space of the LPR definition of "Ring-LWE"? E.g.:
>
> * A proof for anything like Round5 or Saber would not qualify as
> "enjoy a security proof that is analogous", because the starting
> assumption is hardness of Ring-LWR/Module-LWR instead of Ring-LWE?
>
> * If a scheme releases many overstretched Ring-LWE samples, and has a
> proof based on the alleged hardness of this case of Ring-LWE, then

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Tuesday, May 07, 2019 5:10 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: NTRU Prime
Attachments: signature.asc

Damien Stehlé writes:

> I've already clarified the meaning of the claim.

Do Round5 and Saber "enjoy a security proof that is analogous to that of the LPR scheme"? If you were being clear in your claim that NTRU LPRime doesn't, then surely this question would be easy to answer too.

Given your nearby quotes "rely on the Ring-LWE hardness assumption" and "the same Ring-LWE security 'proof'" and "some type of Ring-LWE assumption", I'm guessing that you're requiring the hardness assumption in the proof to fit within the parameter space of the LPR definition of "Ring-LWE". But then Round5 and Saber also don't have such a proof, and you're wrong in attributing this to details of the key distribution.

> I have the impression that you are asking me to give a value judgement

No. You made an unclear claim regarding NTRU LPRime, and I'm asking you what the intended meaning of the claim is.

You say that the keys are multiples of 3 and thus not uniform. This is clear, correct, and not new.

You then jump from this to a vague claim that "NTRU LPRime does not enjoy a security proof that is analogous to that of the LPR scheme".

This is where I'm asking for clarification.

With my top guess for what this claim is supposed to mean (see above regarding Ring-LWE), the claim applies to all of the pure rounding schemes, including Round5 and Saber. It is not specific to NTRU LPRime. This contradicts your subsidiary claim that this is "due to" the details of the key distribution in NTRU LPRime.

The lack of clarity has made this unnecessarily difficult to resolve.

The response has to be phrased in an unnecessarily complicated way (if you meant this then ...). You've been ignoring the response---making readers think that actually you meant something else. Well, okay, then what does your claim mean?

> Again, from context, it is clear that it is related to the
> coefficients of A2 being multiples of 3 (aka rounded multipliers).

Sure. I don't think I'm having any trouble understanding the meaning of your subsidiary claim. Your subsidiary claim is that your main claim--- about the alleged difficulty of coming up with some unspecified type of proof---is "due to" the fact that NTRU LPRime keys are multiples of 3.

Structurally, this is disclosing a fragment of the logic that you have in mind for a particular argument for the main claim. This is very far from clarifying what the main claim means.

---Dan (speaking for myself)

--

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

From: Mike Hamburg <mike@shiftleft.org>
Sent: Tuesday, May 07, 2019 12:36 PM
To: D. J. Bernstein
Cc: pqc-comments; pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: NTRU Prime

This is getting tiresome. Let me take a try at it.

Damien appears to mean: under what simple, concise and preferably common assumption are NTRU LPRime's public keys indistinguishable from random and its ciphertexts one-way? He doesn't like the answer "they're one-way if Problem 3 is hard" because Problem 3 is vague. In particular, he appears to object to the statement that "one can see these problems, for uniform random inputs, as models of attacks on ... NTRU" because the public key doesn't look like a uniformly random ring element: it's a multiple of 3. This may be a sore point because Kyber changed its spec in the second round, removing a rounding step so that the public key would be indistinguishable from random in the ring.

Dan's counterpoint — that all the entrants are secure if their public keys look random and their ciphertexts are one-way with a random public key — is obviously true. And it may be that for NTRU LPRime there's no simpler assumption. But most of the entrants have been designed with an eye to making those problems simple and similar to past proposals when possible, though of course that just means that they are each working with a particular variant of [RM]LW[ER] (with what ring and what noise distribution?). Still, making the assumption simpler and similar to past ones slightly reduces the risk that a new family of attacks will apply only to your variant.

The same criticism might not apply to Round5 and Saber because they do RLWR mod 2, which (if I'm thinking correctly here with my phone on the train) allows one to lift the rounded public key to a hopefully-indistinguishable-from-uniformly-random element of the original ring. Therefore the same RLWR variant with uniformly random ring elements can serve for both the public key and ciphertext.

Regards,
— Mike Hamburg

> On May 7, 2019, at 2:10 AM, D. J. Bernstein <djb@cr.yp.to> wrote:

>

> Damien Stehlé writes:

>> I've already clarified the meaning of the claim.

>

> Do Round5 and Saber "enjoy a security proof that is analogous to that

> of the LPR scheme"? If you were being clear in your claim that NTRU

> LPRime doesn't, then surely this question would be easy to answer too.

>

> Given your nearby quotes "rely on the Ring-LWE hardness assumption"

> and "the same Ring-LWE security 'proof'" and "some type of Ring-LWE

> assumption", I'm guessing that you're requiring the hardness

> assumption in the proof to fit within the parameter space of the LPR

> definition of "Ring-LWE". But then Round5 and Saber also don't have

> such a proof, and you're wrong in attributing this to details of the key distribution.

>

>> I have the impression that you are asking me to give a value

>> judgement

>

> No. You made an unclear claim regarding NTRU LPRime, and I'm asking

> you what the intended meaning of the claim is.

From: Christopher J Peikert <cpeikert@alum.mit.edu>
Sent: Tuesday, May 7, 2019 6:38 PM
To: pqc-forum; pqc-comments
Subject: Re: [pqc-forum] Re: ROUND 2 OFFICIAL COMMENT: NTRU Prime

I hesitate to lengthen a thread whose straightforward point was clear (to me at least) from Damien's original messages; thanks to Mike for summarizing so well.

In particular, Damien's comments refer to the average-case Ring-LWE problem, and explicitly put worst-case lattice problems (and associated hardness theorems) out of scope. Yet for mysterious reasons, subsequent posts made some disputable claims about these topics, so I'd like to make a few remarks.

(Since worst-case hardness theorems for Ring-LWE appear to be of no consequence to the remaining NIST submissions, I won't comment further on the topic in this forum.)

On Sat, May 4, 2019 at 3:00 PM D. J. Bernstein <djb@cr.yp.to> wrote:

> The Lyubashevsky--Peikert--Regev paper

>

>

> <https://eprint.iacr.org/2012/230>

>

> was supposedly "proving" that Ring-LWE "enjoys very strong hardness
> guarantees", namely a theorem relating Ring-LWE to an approximate
> Ideal-SVP problem. However, a closer look shows that this "very strong
> hardness guarantee" provides little assurance for security reviewers:

> ...

>

> * A few cryptanalysts have been attacking the approximate Ideal-SVP
> problem in the last few years, and have found ways to exploit the
> structure of the problem---especially the extra automorphisms
> provided by the usual fields---to do surprising levels of damage.

>

> At some point the claim of a "very strong hardness guarantee" has to
> succumb to the facts: this "guarantee" starts from hypotheses that
> haven't been holding up well against cryptanalysis, and then draws
> security conclusions about irrelevant cryptosystem parameters.

And again later:

> The underlying "worst-case problems on ideal lattices" have not been
> holding up well to cryptanalysis.

Certainly we should be dealing in facts! However, the facts don't support the conclusion that the worst-case problems considered in LPR "haven't been holding up well against cryptanalysis." Indeed, the problems of interest remain completely unaffected by recent cryptanalysis.

Fact 1: the relevant worst-case problems from LPR (which underlie the LPR cryptosystem, and most other cryptographic constructions based on Ring-LWE) are to quantumly approximate Ideal-SVP in cyclotomic rings to within some **polynomial** factors in the dimension n .

(Of course one can consider Ring-LWE parameters that correspond to Ideal-SVP for larger approximation factors, e.g., quasi-polynomial or even slightly subexponential. But applications using the former are rare and "exotic," like Fully Homomorphic Encryption, and I'm not aware of any that require the latter. Lattice-based crypto mainly lies in the realm of (near-)polynomial factors.)

Fact 2: none of the (excellent and exciting) cryptanalytic work on approx-Ideal-SVP comes close to making a dent in the above problems.

Specifically, there are no known speedups for poly-approx Ideal-SVP (in any proposed ring) versus poly-approx SVP in general lattices, apart from long-known minor speedups arising from rotational symmetries. (Such speedups exist for all proposed rings, and for NTRU and Ring-LWE problems as well.)

Fact 3: recent works have given quantum algorithms that approximate Ideal-SVP to within subexponential $2^{O(\sqrt{n})}$ factors in polynomial time, and also offer time-approximation tradeoffs---but again, they don't provide any speedup for polynomial approximation factors.

Moreover, the latest algorithms don't require any automorphisms or subfields, though they do need "advice" about the number field that can be obtained by expensive preprocessing. See <https://eprint.iacr.org/2019/215> for the state of the art, and <https://eprint.iacr.org/2019/234> to understand the substantial concrete factors hidden by the $O(\sqrt{n})$ notation in the cyclotomic case.

(I feel compelled to point out a sense of *deja vu* here: Bernstein has previously missed and/or glossed over the major difference between polynomial approximation factors and huge, cryptographically irrelevant $2^{O(\sqrt{n})}$ factors. Here's an example from almost three years ago---and it's far from the only one: <https://groups.google.com/forum/#!msg/cryptanalytic-algorithms/y-wAnhmGslo/VCDFFcXnAgAJ>)

- > What the LPR paper actually says that it is (1) "introducing an algebraic variant of LWE" and (2) "proving that it too enjoys very strong hardness guarantees. Specifically, we show that the ring-LWE distribution is pseudorandom, assuming that worst-case problems on ideal lattices are hard for polynomial-time quantum algorithms."

- > A closer look shows that these "guarantees" are for irrelevant
- > cryptosystem parameters. (I think
- > <https://eprint.iacr.org/2016/360> deserves the primary credit for this
- > observation.)

No, the fact that worst-case hardness theorems don't (yet) yield practical concrete parameters was already recognized many years earlier, e.g., in the 2009 survey <https://cseweb.ucsd.edu/~daniele/papers/PostQuantum.pdf> by Micciancio and Regev (in a book edited by Bernstein). Here's the relevant quote on page 3:

"Second, in principle the worst-case security guarantee can help us in choosing concrete parameters for the cryptosystem, although in practice this leads to what seems like overly conservative estimates, and as we shall see later, one often sets the parameters based on the best known attacks."

Sincerely yours in cryptography,
Chris

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Thursday, April 16, 2020 2:39 AM
To: pqc-comments
Cc: pqc-forum
Subject: ROUND 2 OFFICIAL COMMENT: NTRU Prime
Attachments: signature.asc

This is an overview of three families of recent cycle-count improvements for NTRU Prime.

1. For Haswell: NTRU Prime already had fast code online but now has even faster code online. This is the first announcement. For dimension 761, the new Haswell (hiphop) speeds from SUPERCOP are as follows:

ntrulpr761: 44092 keygen, 72372 enc, 85908 dec
sntrup761: 810148 keygen, 48780 enc, 59120 dec

Code sharing across sizes has also produced a bunch of speedups in code online for the smaller and larger sizes, 653 and 857, although in those cases it's clear that some extra per-size code would save more time.

IND-CCA2 security allows keys to be reused any number of times, so the per-ciphertext cycles are only enc+dec. Only two other NISTPQC submissions offer IND-CCA2 options as fast as sntrup761 in Haswell enc+dec performance. (For example, enc+dec is 147408 for kyber768, 107900 for sntrup761, and 73776 for kyber90s768. The other is NTRU.)

2. For Cortex-M4: Bo-Yin Yang has just announced the results of serious optimization of sntrup761 and ntrulpr761, demonstrating the suitability of NTRU Prime for this platform. Previous reports of NTRU Prime performance on this platform used the reference NTRU Prime code, which has always prioritized readability; the new code is much faster.

3. Regarding performance comparisons between ntruhrss701 (CECPQ2) and sntrup761, and claims that sntrup761 keygen is too slow for applications that constantly generate one-time keys (such as, again, CECPQ2): There's an important announcement coming up momentarily. To limit the sizes of "OFFICIAL COMMENT"s, this will be sent only to pqc-forum@list.nist.gov, but it is included here by reference.

---Dan