From: oscar.garcia-morchon@philips.com
Sent: Saturday, August 04, 2018 11:00 AM
To: pqc-forum
Subject: [pqc-forum] OFFICIAL COMMENT: Round5 = Round2 + Hila5

Dear all,

with this email we want to officially announce Round5, the merging of the Round2 and Hila5 proposals.

Technically, the merged proposal is based on Round2 combined with Hila5's techniques such as error correction to bring down the failure probability so that Round5 can achieve better bandwidth and CPU performance. The Round5 parameters are adapted to fully comply with NIST security levels.

We also want to announce that in addition to the members of the original Hila5 and Round2 proposals, Thijs Laarhoven (TU Eindhoven) has also joined the Round5 team.

All information is available at https://round5.org/ In the next weeks, we will post further updates.

In the website you can already find two papers. The first paper describes Round5 - including algorithms, security analysis and parameters. The second paper details an optimized implementation - due to Markku J.O. Saarinen - on Cortex M4. This implementation is available at https://github.com/round5/r5nd_tiny. The papers show that Round5 is a leading lattice-based candidate in terms of security, bandwidth and CPU performance.

Best regards,

Oscar, Markku, and the rest of the Round5 team.

From: Markku-Juhani O. Saarinen <mjos.crypto@gmail.com>
Sent: Saturday, August 04, 2018 12:29 PM
To: pqc-forum
Subject: [pqc-forum] Re: OFFICIAL COMMENT: Round5 = Round2 + Hila5

Hello All,

As announced by Oscar, my "Hila5" project has merged with the "Round2" effort to produce **Round5**, and this has been acknowledged by NIST (we gave them an early heads-up about our merger plans). I will no longer work on Hila5, so all analytic efforts should be directed to the new candidate. Round5 does inherit some design features from Hila5, perhaps most importantly the constant-time error correction code XEf.

Oscar will be the principal submitter of the new merged proposal (or tweak), so my opinions and communications should be taken as unofficial. However, I'd like to fill in some details:

We have worked hard on the new Round5 proposal. At least in my opinion it is clearly superior to either of the old ones; it has the best message sizes of lattice-based candidates, and also leading performance characteristics of *all* candidates.

This is largely because the flexibility of the design allowed a fine-tuned parameter search to meet the post-quantum and classical security levels while optimizing parameters, primarily for bandwidth (public key and ciphertext message sizes). The team has done a fantastic job at it.

* Hayo Baan (Philips, NL)
* Sauvik Bhattacharya (Philips, NL)
* Oscar Garcia-Morchon (Philips, NL)
* Thijs Laarhoven (TU/e, NL)
* Markku-Juhani O. Saarinen (PQShield, UK)
* Ronald Rietman (Philips, NL)
* Ludo Tolhuizen (Philips, NL)
* Jose Luis Torre Arce (Philips, NL)
* Zhenfei Zhang (OnboardSecurity, US)

The official homepage is at https://round5.org/. I also intend to keep some unofficial implementation-related materials at https://mjos.fi/round5/

While we prepare for the official NIST tweak proposal package, we are putting two preprints out at this time:

On parameter selection and security analysis:

* “Round5: Compact and Fast Post-Quantum Public-Key Encryption” by S. Bhattacharya, O. Garcia-Morchon, T. Laarhoven, R. Rietman, M.-J. Saarinen, L. Tolhuizen, and Z. Zhang. https://round5.org/doc/round5paper.pdf

Implementation aspects (of the ring variant) are discussed in:

* “Shorter Messages and Faster Post-Quantum Encryption with Round5 on Cortex M” by M.-J. Saarinen, S. Bhattacharya, O. Garcia-Morchon, R. Rietman, L. Tolhuizen, and Z.

Zhang. https://round5.org/doc/r5m4text.pdf

I'm still working on optimizing implementations. However, the Cortex M4 implementation discussed in the latter paper is already available at https://github.com/round5/r5nd_tiny

Again, please note that all parameter selections etc before the actual NIST submission are still subject to change.

Cheers,
- markku

Dr. Markku-Juhani O. Saarinen <mjos@iki.fi>

From: Leo Ducas <leo.ducas1@gmail.com>
Sent: Tuesday, August 07, 2018 2:07 AM
To: pqc-forum
Subject: [pqc-forum] Re: OFFICIAL COMMENT: Round5 = Round2 + Hila5

Dear authors,

I note that the failure analysis assumes that ``bit failures occur independently'', but I'm unconvinced it would be the case, especially in the ring setting. I have searched for solution to this issue for a long time, and still don't know how to properly address this issue theoretically.

May I suggest to resort to experimental analysis to test how close or not to independent these failure events are, at least in a regime where failures are statistically measurable ?

Best regards
-- Leo Ducas

Le samedi 4 août 2018 18:29:10 UTC+2, Markku-Juhani O. Saarinen a écrit :

> Hello All,
>
> As announced by Oscar, my "Hila5" project has merged with the "Round2" effort to produce **Round5**, and this has been acknowledged by NIST (we gave them an early heads-up about our merger plans). I will no longer work on Hila5, so all analytic efforts should be directed to the new candidate. Round5 does inherit some design features from Hila5, perhaps most importantly the constant-time error correction code XEf.
>
> Oscar will be the principal submitter of the new merged proposal (or tweak), so my opinions and communications should be taken as unofficial. However, I'd like to fill in some details:
>
> We have worked hard on the new Round5 proposal. At least in my opinion it is clearly superior to either of the old ones; it has the best message sizes of lattice-based candidates, and also leading performance characteristics of *all* candidates.
>
> This is largely because the flexibility of the design allowed a fine-tuned parameter search to meet the post-quantum and classical security levels while optimizing parameters, primarily for bandwidth (public key and ciphertext message sizes). The team has done a fantastic job at it.
>
> * Hayo Baan (Philips, NL)
> * Sauvik Bhattacharya (Philips, NL)
> * Oscar Garcia-Morchon (Philips, NL)
> * Thijs Laarhoven (TU/e, NL)
> * Markku-Juhani O. Saarinen (PQShield, UK)
> * Ronald Rietman (Philips, NL)
> * Ludo Tolhuizen (Philips, NL)
> * Jose Luis Torre Arce (Philips, NL)
> * Zhenfei Zhang (OnboardSecurity, US)

From: Markku-Juhani O. Saarinen <mjos.crypto@gmail.com>
Sent: Tuesday, August 07, 2018 6:44 AM
To: pqc-forum
Subject: [pqc-forum] Re: OFFICIAL COMMENT: Round5 = Round2 + Hila5

On Tuesday, August 7, 2018 at 7:06:42 AM UTC+1, Leo Ducas wrote:

> I note that the failure analysis assumes that ``bit failures occur independently'',
> but I'm unconvinced it would be the case, especially in the ring setting. I have
> searched for solution to this issue for a long time, and still don't know how to
> properly address this issue theoretically.

Hi Leo,

Same thing here, it's a known open issue. Also I find that my inconclusive results based on probability convolutions (Section 3.2, https://eprint.iacr.org/2017/424) are now being cited in other works (for example upcoming SAC paper https://eprint.iacr.org/2018/150).

Another, more easily avoidable problem is that some authors also appear to completely ignore the side-channel aspect of error correction codes, not fully realising that retrofitting side-channel resistance to error correction codes of certain type can incur a significant cost. Let me remind that this is one of the additional evaluation criteria (Section 4.A.6 of the Call).

In symmetric cryptographic design, which is arguably more mature, such a design omission would be unthinkable. However on some probability distributions one unfortunately still has to rely on heuristic statistical arguments there too.

> May I suggest to resort to experimental analysis to test how close or not to
> independent these failure events are, at least in a regime where failures are
> statistically measurable ?

I have done this, and it pretty much confirmed the analysis performed by the team. However affirmative computations of this type are rarely reported with analytic work, and the computations were not performed on the exact final parameters.

For example: One larger experiment was performed on an experimental set of parameters, but sufficiently large to be in similar scale to the ones we are using. I ran a version with n=d=700, h=196, q=2^14, p=2^8, t=2^4 with mu=700 message bits for m=2570540000 (2^31.26) keygen-encrypt-decrypt iterations, which produced *one* message with a single bit error. This puts the per-bit error rate at c=1/(mu*m) = 2^-40.71 and was consistent with the bounds we had analyzed for those parameters. Note that not all mu bits are actually used to carry a message and error correction codes were being used etc.

The computation took 30 CPU hours on a 16-node cluster and ran until that single error was found. It did provide evidence that errors do not happen radically more often than we expected, but of course statistical conclusions cannot be drawn from a single bit error.

Best Regards,
- markku

Dr. Markku-Juhani O. Saarinen <mjos@iki.fi>

From: Markku-Juhani O. Saarinen <mjos.crypto@gmail.com>
Sent: Tuesday, August 07, 2018 7:29 AM
To: pqc-forum
Subject: [pqc-forum] Re: OFFICIAL COMMENT: Round5 = Round2 + Hila5

Hi Leo,

Should add this is of course not only related to our submission, and I was answering only in a personal capacity (I'm getting slack from team for giving an initial reply without even consulting them -- they have more detailed answer to give than the few experiments I did.)

However we will probably instantiate and report a more robust experimental regime to answer the question of independence specifically in our case.

Cheers,
- markku

On Tuesday, August 7, 2018 at 11:43:45 AM UTC+1, Markku-Juhani O. Saarinen wrote:

> On Tuesday, August 7, 2018 at 7:06:42 AM UTC+1, Leo Ducas wrote:
>
> > I note that the failure analysis assumes that ``bit failures occur independently'',
> > but I'm unconvinced it would be the case, especially in the ring setting. I have
> > searched for solution to this issue for a long time, and still don't know how to
> > properly address this issue theoretically.
>
> Hi Leo,
>
> Same thing here, it's a known open issue. Also I find that my inconclusive results based on probability convolutions (Section 3.2, https://eprint.iacr.org/2017/424) are now being cited in other works (for example upcoming SAC paper https://eprint.iacr.org/2018/150).
>
> Another, more easily avoidable problem is that some authors also appear to completely ignore the side-channel aspect of error correction codes, not fully realising that retrofitting side-channel resistance to error correction codes of certain type can incur a significant cost. Let me remind that this is one of the additional evaluation criteria (Section 4.A.6 of the Call).
>
> In symmetric cryptographic design, which is arguably more mature, such a design omission would be unthinkable. However on some probability distributions one unfortunately still has to rely on heuristic statistical arguments there too.
>
> > May I suggest to resort to experimental analysis to test how close or not to
> > independent these failure events are, at least in a regime where failures are
> > statistically measurable ?
>
> I have done this, and it pretty much confirmed the analysis performed by the team. However affirmative computations of this type are rarely reported with analytic work, and the computations were not performed on the exact final parameters.
>
> For example: One larger experiment was performed on an experimental set of parameters, but sufficiently large to be in similar scale to the ones we are using. I ran a version with n=d=700, h=196, q=2^14, p=2^8, t=2^4 with mu=700 message bits for m=2570540000 (2^31.26) keygen-encrypt-decrypt iterations, which produced

From: Leo Ducas <leo.ducas1@gmail.com>
Sent: Tuesday, August 07, 2018 2:29 PM
To: pqc-forum
Subject: [pqc-forum] Re: OFFICIAL COMMENT: Round5 = Round2 + Hila5

Thanks for the pointers and your early experiments.

Beyond testing the full scheme itself, I think what would be interesting is to test independence for intermediate results. Indeed, I suspect that the product terms es' and e's have strong correlation, but they start to fade off a bit when summing es'+e's, and in the end, this may be mostly drown out by the rounding error term, whose coordinates are fully independent.

Best regards
-- Leo Ducas

Le mardi 7 août 2018 13:29:25 UTC+2, Markku-Juhani O. Saarinen a écrit :

> Hi Leo,
>
> Should add this is of course not only related to our submission, and I was answering only in a personal capacity (I'm getting slack from team for giving an initial reply without even consulting them -- they have more detailed answer to give than the few experiments I did.)
>
> However we will probably instantiate and report a more robust experimental regime to answer the question of independence specifically in our case.
>
> Cheers,
> - markku
>
> On Tuesday, August 7, 2018 at 11:43:45 AM UTC+1, Markku-Juhani O. Saarinen wrote:
>
> > On Tuesday, August 7, 2018 at 7:06:42 AM UTC+1, Leo Ducas wrote:
> >
> > > I note that the failure analysis assumes that ``bit failures occur independently'',
> > > but I'm unconvinced it would be the case, especially in the ring setting. I have
> > > searched for solution to this issue for a long time, and still don't know how to
> > > properly address this issue theoretically.
> >
> > Hi Leo,
> >
> > Same thing here, it's a known open issue. Also I find that my inconclusive results based on probability convolutions (Section 3.2, https://eprint.iacr.org/2017/424) are now being cited in other works (for example upcoming SAC paper https://eprint.iacr.org/2018/150).
> >
> > Another, more easily avoidable problem is that some authors also appear to completely ignore the side-channel aspect of error correction codes, not fully realising that retrofitting side-channel resistance to error correction codes of certain type can incur a significant cost. Let me remind that this is one of the additional evaluation criteria (Section 4.A.6 of the Call).
> >
> > In symmetric cryptographic design, which is arguably more mature, such a design omission would be unthinkable. However on some probability distributions one unfortunately still has to rely on heuristic statistical arguments there too.