From: Moody, Dustin (Fed)
To: (b) (6)
Subject: RE: notes for the meeting tomorrow
Date: Monday, December 18, 2017 12:24:00 PM

Thanks – I'll put them into the Master List

From: Daniel Smith (b) (6)
Sent: Monday, December 18, 2017 12:23 PM
To: Moody, Dustin (Fed) <dustin.moody@nist.gov>
Cc: Chen, Lily (Fed) <lily.chen@nist.gov>
Subject: notes for the meeting tomorrow

Hi,

Here are a collection of notes on the schemes I reviewed. Instead of focusing on the ones that are in the "definitely need to discuss" queue, I'm writing on all of the ones I had. Please pick and choose comments accordingly.

CRYSTALS-Kyber: Cool name. No specific comment. Seems complete.

DAGS: The suggested 256-bit secure parameters are either broken or not functional depending on your perspective. The 256-bit parameters for the KEM only provide 160 or 192 bits of entropy for shared keys, so it is not capable of establishing shared keys of length at least 256-bits. Please recheck all other KEMs to make sure that they are not making the same error. (The submitters have acknowledged it.) It is an easy fix in this case, so I recommend acceptance. There is another issue in that in the statement on advantages and disadvantages it is claimed that there are no decoding errors, however, the decapsulation algorithms explicitly uses "bottom" to indicate decoding errors. Something is not clear here. I also see no other place indicating the error rate nor any justification for the no decoding errors comment. This is a serious mistake which should be clarified. I still recommend erring on the side of acceptance.

DME: This is another one that I can break the 256-bit secure parameters. The submitters are not familiar with standard references in multivariate crypto. By using the field equations I can upper bound the complexity of running F4 on the 256-bit parameters using generic methods to achieve an attack of complexity about 2^205. We notified the submitters of this and they said that this is plausible but they don't know the science (obviously paraphrasing). They also have other false claims in their security analysis, such as the solving degree being around q. Simply not close to true. In fact, their most secure parameters with q=2^24 has a solving degree somewhere around 18. Still, I don't think that the algebraic attack is what will ultimately break this one. I don't have a full attack, though, so I recommend acceptance as complete and we can see how it is broken in a few months.

Emblem: My notes say that I only noticed category one parameters.

Gui: They haven't quite followed the rules on the software submission. We really must accept this one, though. This is the oldest efficient unbroken post-quantum signature I know of. We should keep pushing them relentlessly for adherence to the rules, but the search would certainly be incomplete without this. Besides, the science is actually correct in this one.

Kerus: I can break this one by hand for all parameters... literally. I'll attach the comment I put in the checklist. In pass one, the matrices T1=YA, and T2=BY^-1X are sent openly. In pass two, the matrices T3=DYAC^-1 and T4=CBY^-1XH are sent openly. In pass three, the matrix T5=DXH is sent openly. Since A and B are centro-symmetric, the product P=A^-1B^-1 is as well. Notice that P satisfies the linear equation T3PT4=T5 by the commutativity property of centro-symmetric matrices. So we can solve for P. Then T1PT2=X, the shared "secret." The scheme is totally bogus. It should be a definite reject.

LAC: Seems okay to me.

LEDAkem: The submitters backtracked from a IND-CCA claim to an IND-CPA claim, essentially admitting that they were wrong on that one. I'm not sure about this one, but I would err on the side of acceptance. It is likely to be eliminated early in it's current state anyway. If they can make it better and we accept changes later in the process... great. I doubt it will survive, though.

NTRUKEM: Seems fine.

NTRUPRIME: The proposal advertises two schemes SNTRUP and NTRULP. It has no timing data that I see for the second. I think then that it is complete for SNTRUP and incomplete for NTRULP (as two submissions).
PQRSA: Doesn't mention amount of volatile memory used, which is likely large, unlike other submissions. Probably need some data on that, and I don't remember it being provided.

RaCoSS: No specific comment. Seems alright.

Rainbow: Again, it seems that there is an issue with Intel specific implementation issues. I'm not sure if these guys are trying to pull something or they though it that across platforms referred to across Intel platforms. Again, we should hold them accountable, but definitely accept. There is no multivariate scheme I know of with more theoretical support than rainbow. It is on a very solid footing.

SPHINCS+: No specific comment. Seems okay.

STRPI and TPSig: The submitters responded to our query about the linearity of decryption/signing by stating that these functions ARE linear, but "only for the secret key holder." That is just nonsense. It is linear or it is not. y=3x is still linear even if I don't look at it. One could interpolate the affine function by generating a large number of plaintext/ciphertext pairs and solving for unknown coefficients. The resulting map is a more efficient decryption technique than provided and the inverse is a more efficient encryption algorithm, though we might as well use the identity function. The scheme is not secure. On the other hand, the size of the scheme is such that I cannot accurately say that the scheme does "not incorporate major components believed to be insecure against quantum computers." Mostly since large enough quantum computers have yet to be built…mostly. This is a definite reject. The idea that a function is linear for some people and not others is laughable.

That's all I got. Thanks!

Cheers,
Daniel