From: (b) (6)
To: Dang, Quynh H. (Fed)
Cc: Apon, Daniel C. (Fed); Perlner, Ray A. (Fed); Moody, Dustin (Fed); internal-pqc; Dang, Thinh H. (Fed)
Subject: Re: Kyber"s response discussion tomorrow ?
Date: Friday, June 5, 2020 10:52:50 AM

As I mentioned before, the lattice schemes aren't in a vacuum. Again, this discussion makes me think of GeMSS and how conservative, careful and respectful of our CFP they are.

Not only are they assuming a cost model that is unrealistically conservative on time, it requires memory on the order of the security level, just in case. While I think that it was a bad decision on their part to assume this cost model (in time complexity), they could easily have gone the other direction as some of the lattice candidates have and have said that memory access allows them to require fewer than 2^143 operations. It is entirely doable to tweak the parameters so that all of the other attacks cost more than 2^143 bit operations while the direct attack costs far less but requires memory much higher than the fairly ridiculously low value of 2^89 (only about a billion times the amount currently stored by google... what does that suggest when we are considering the time-span of a few decades?).

GeMSS is also being very respectful of our security requirements in other ways (I think, unless they are overlooking optimizations... of course some of their team are the most famous in the world for efficiency in computer algebra, so...). For example, GeMSS uses the Fiestel-Patarin construction to bootstrap security against hash collisions. Basically they multiply the relatively low cost of hash collision on their small strings across several hash values. This requires the scheme to use its public key for some number of repetitions for each signature. The complexity of the collision attack should clearly include (probably somehow amortized) the cost of the hash algorithm and the cost of the evaluation of the public key multiple times to get all of those collisions. The team doesn't do this, though, they want the number of hash calls and public key evaluations to be greater than 2^128. It looks to me like keeping their exact parameters and using one fewer iteration still is (barely) above 2^143 bit operations (using about 2^80 memory or so with [van Oorshot-Wiener,99]). The point is, they are setting the cost of calls to 1 for collision search on purpose though it hurts performance (probably by about 33% or so).

Now the last choice above is I think just a bad one, but I understand why they think that way. (If you were to only consider the cost of one collision search iteration, meaning one selection of signature and message to try to form a collision, as costing as much as a single call to SHA-3, then the complexity would come out to 2^141.5. In reality, though, even ignoring the cost of evaluating the trapdoor function, you still even in this unrealistic scenario require at least 3 evaluations of SHA-3, so the cost is still over 2^143.) But still, it is very clear to me that this scheme is affected adversely by the choice to respect the CFP while (at least in principal if not in practice) the schemes that choose to consider their views and interpret their own requirements on cost are bending the rules.

We have no reason for the playing field to be perfectly level, but there should be some justification that we can agree to. Otherwise we are comparing apples to chimpanzees. So I think that we should actually have a fairly hard stance on this issue. It may not be the case, but it gives the appearance that schemes like Dilithium are getting an unfair advantage by ignoring requirements against schemes like GeMSS that are going beyond what we asked. (Not that these would be competing in any possible world.)

Cheers,
Daniel

On Fri, Jun 5, 2020 at 10:44 AM Dang, Quynh H. (Fed) <quynh.dang@nist.gov> wrote:

> And of course, there is at least one downside with the second option too: Dan and other people could complain that we changed our security model in a significant way in the final round from the call for proposal.
>
> So, it could be best to ask the Kyber team to increase their noise and we'd stick with the current cost model.
>
> Quynh.
>
> ---
>
> **From:** Dang, Quynh H. (Fed) <quynh.dang@nist.gov>
> **Sent:** Friday, June 5, 2020 9:35 AM
> **To:** Apon, Daniel C. (Fed) <daniel.apon@nist.gov>; Perlner, Ray A. (Fed) <ray.perlner@nist.gov>; Daniel Smith (b) (6)
> **Cc:** Moody, Dustin (Fed) <dustin.moody@nist.gov>; internal-pqc <internal-pqc@nist.gov>; Dang, Thinh H. (Fed) <thinh.dang@nist.gov>
> **Subject:** Re: Kyber's response discussion tomorrow ?
>
> Hi Daniel (one of my favs),
>
> I agree with you.
>
> However, the issue Dan raised is that Kyber-512 does not have a proof that it achieves the security level 1 that we specified in the call.
>
> Kyber-512 might be or could be a couple/few bits of security lower than AES-128 as we specified: not counting communication/memory access costs.
>
> We asked people to specify parameter sets in a conservation ranges to minimize the risk that their parameter sets do not meet the security targets/levels.
>
> There are at least 2 options to handle this issue.
>
> 1) We'd say that a couple of bits off are ok. The problem with this is that we did not say that in the call and this would be contradictory to asking people to be conservative in specifying their parameter sets in the call.
>
> 2) We'd say that: When we made a call for proposals, there was no consensus cost model for communication/memory access that we were aware of. Therefore, we did not count them in specifying the security levels. And, then describe our new position about these

costs.

Quynh.

From: Apon, Daniel C. (Fed) <daniel.apon@nist.gov>
Sent: Friday, June 5, 2020 3:41 AM
To: Perlner, Ray A. (Fed) <ray.perlner@nist.gov>; Daniel Smith (b) (6); Dang, Quynh H. (Fed) <quynh.dang@nist.gov>
Cc: Moody, Dustin (Fed) <dustin.moody@nist.gov>; internal-pqc <internal-pqc@nist.gov>; Dang, Thinh H. (Fed) <thinh.dang@nist.gov>
Subject: RE: Kyber's response discussion tomorrow ?

I felt that the key point in the Kyber Team's response was

"We agree that … 141 [is] smaller than 143, but at the moment we do not consider this to be a sufficient reason to modify the Kyber-512 parameter set.
…
The additional memory requirement of this attack strongly suggests that Kyber-512 is more secure than AES-128 in any realistic cost model.
…
A planar sheet of terabyte micro-SD cards the size of New York City (all five boroughs, 800 km^2 ~ 2^49.5 mm^2) would hold 2^89 bits.
"

I still feel we should do our own internal analysis at the start of the 3rd Round.

I'm utterly opposed to letting DJB's eleventh-hour protestations influence absolutely anything whatsoever.

--Daniel

From: Perlner, Ray A. (Fed) <ray.perlner@nist.gov>
Sent: Thursday, June 4, 2020 3:02 PM
To: Daniel Smith (b) (6); Dang, Quynh H. (Fed) <quynh.dang@nist.gov>
Cc: Moody, Dustin (Fed) <dustin.moody@nist.gov>; internal-pqc <internal-pqc@nist.gov>; Dang, Thinh H. (Fed) <thinh.dang@nist.gov>
Subject: RE: Kyber's response discussion tomorrow ?

Here are my current thoughts on the matter:

I am open to the idea of using a more realistic model of computation than the basic gate model. However, a lot of the ideas in I’ve seen in the literature seem too pessimistic (as in they reckon attacks as being harder than they should be. – at least in the long term)

DJB’s favored model, for example, assumes the computation must be implemented by only nearest neighbor interactions in a 2 dimensional grid. This has some justification, in that trying to violate these assumptions clearly costs more than the basic gate model assumes, but

1. Today’s Supercomputers generally use a meaningfully 3 dimensional arrangement of processors (although the processors themselves are 2 dimensional)
2. Long distance connections needing high performance are implemented by fiber optic cables, and sending a bit through a kilometer of fiber optic cable, while more expensive than sending the bit across a single AND gate, clearly costs less than sending it through a kilometer of densely packed AND gates (which is how DJB’s favored model would treat it.)

NTRU’s “local” model seems in practice to be even more extreme, simply ignoring any algorithm that hasn’t explicitly been implemented locally

Hard limits on the total memory size have also been proposed. I think the smallest numbers I could really convince myself were commensurate with an adversary actually capable of threatening the appropriate security level were 2^100 for levels 1 and 2, 2^150 for levels 3 and 4 , and 2^180 for level 5.

One could perhaps adjust the RAM model to cost random access queries to a memory of size N at N^(1/3) in terms of depth and (log(N))^2 in terms of gate count and require all other gates to be local. (I think I might actually be ok with that, keeping in mind that if the whole thing can be implemented locally, you don’t need to make RAM queries, no matter how large the computation is.)

The other worry though is that things like memory cost are much more susceptible to being optimized away by incremental improvements, which the first iteration of a new attack rarely includes. But there are a lot of smart lattice people, so maybe I can be convinced they’ve thought about this stuff enough that there is no room for further improvement. I’m not convinced yet, though.

Ray

From: Daniel Smith (b) (6)
Sent: Thursday, June 4, 2020 2:23 PM
To: Dang, Quynh H. (Fed) <quynh.dang@nist.gov>

Cc: Moody, Dustin (Fed) <dustin.moody@nist.gov>; internal-pqc <internal-pqc@nist.gov>; Dang, Thinh H. (Fed) <thinh.dang@nist.gov>
Subject: Re: Kyber's response discussion tomorrow ?

Hmmm...

They are calling us out explicitly to offer our position on this. It is a muddy issue in my mind.

I have a bit of a problem with saying, "We are secure because of other stuff that we can't measure really well." For other areas we have been requiring them to ignore memory costs even when that makes a difference for them.

A clear example comes to mind: GeMSS. For GeMSS they had a quite exhaustive analysis of known techniques applied to GeMSS. They quite conservatively used analyses and coefficients that are unrealistic even with zero cost of memory and memory access (which is why confusingly they chose to report some of the numbers as lower than the security bounds when actually they should be fine). When you consider the hidden polynomial factors or actual coefficients, the least costly attack (and the one they are basing the parameters on) is the direct algebraic attack. They are being super conservative and choosing a linear algebra exponent of 2 for dense linear algebra (I think that we can't use sparse techniques here because of the number of solutions (or the density after fixing variables)), but if we take memory into account, then the complexity is altogether different. If our metric is New York City, then this scheme should benefit fairly significantly.

On a historical note, Ray and I argued fairly extensively about this memory issue when we were drafting the CFP. I recall having discussions about the physical feasibility of converting Jupiter into atomic scale memory that violates causality with the speed of its access (sending replies and being set to different values before being asked to) leading up to the release of this document. The issue as I recall was allowing the community to address some complexity issues that had not been pinned down yet at the time and for the community to come to a consensus on how to address these things. Still, we need to have some standard metric for comparisons between schemes.

I think that it is entirely reasonable to address memory and memory access in a cost model. A problem occurs when we lack justification and when we lack consistency in how we apply restrictions in these analyses. Ray and I were arguing on the level of Jupiter and breaking the laws of physics, whereas Kyber is arguing on the level of the 5 boroughs.

> I would be open to allowing teams to specify their cost model addressing memory (in communication with us and with clear justification and theoretical support), and to adjust parameters accordingly. This would need to take place extremely quickly, though, to not make analysis placed on a moving target.
>
> The easiest way to handle the situation is exactly the opposite, though. That is to let the teams do what they are doing and then judge them by our own metrics. The downside of this approach is that there is plenty of room for bias and plenty of reason for skepticism in our choices if any parts of our community think that we are cutting corners unreasonably.
>
> If we chose to allow memory access cost as part of the complexity analysis, there will be consequences. We may have to communicate with each team explicitly, but I think we should make it clear (if we go that route) that they should analyze the memory concerns with strong justification for **minimal** cost models that they can then incorporate. We also need to assess the feasibility of these models and the appropriateness of the bounds they suggest.
>
> I think that we have plenty to talk about, but we'll follow your lead, Dustin.
>
> Cheers,
>
> Daniel
>
> On Thu, Jun 4, 2020 at 1:47 PM Dang, Quynh H. (Fed) <quynh.dang@nist.gov> wrote:
>
> > I think so. If more people think that a talk tomorrow would be good, then I would ask you to consider that.
> >
> > ---
> >
> > **From:** Moody, Dustin (Fed) <dustin.moody@nist.gov>
> > **Sent:** Thursday, June 4, 2020 1:41 PM
> > **To:** Dang, Quynh H. (Fed) <quynh.dang@nist.gov>
> > **Cc:** internal-pqc <internal-pqc@nist.gov>; Daniel Smith (b) (6) ; Dang, Thinh H. (Fed) <thinh.dang@nist.gov>
> > **Subject:** Re: Kyber's response discussion tomorrow ?
> >
> > I think we can discuss via email.

> I don't think we need to have a meeting tomorrow. Maybe on Tuesday.
>
> Let me know if you think otherwise.
>
> Dustin
>
> ---
>
> **From:** Dang, Quynh H. (Fed) <quynh.dang@nist.gov>
> **Sent:** Thursday, June 4, 2020 1:34 PM
> **To:** Moody, Dustin (Fed) <dustin.moody@nist.gov>
> **Cc:** internal-pqc <internal-pqc@nist.gov>; Daniel Smith (b) (6) ; Dang, Thinh H. (Fed) <thinh.dang@nist.gov>
> **Subject:** Kyber's response discussion tomorrow ?
>
> Hi Dustin,
>
> Are we going to discuss Kyber's response tomorrow at 10 ?
>
> Quynh.