From:	Perlner, Ray (Fed)
To:	Liu, Yi-Kai (Fed); Chen, Lily (Fed); Peralta, Rene C. (Fed); Moody, Dustin (Fed); Alperin-Sheriff, Jacob (Fed); Bassham, Lawrence E. (Fed); Daniel Smith-Tone; Jordan, Stephen P (Fed); Miller, Carl A. (Fed); Smith-Tone, Daniel C. (Fed)
Subject:	RE: First cut at a summary of our thinking on security strengths for the forum.
Date:	Wednesday, October 26, 2016 4:30:55 PM

I think your additional suggested question is a good one.

I'm also happy to add a qualifying word or two to " To demonstrate what these security strength categories will mean for submitters in practice, consider the following scenario:"
Maybe something like " To demonstrate what NIST expects these security strength categories to mean for submitters in practice, consider the following hypothetical scenario:"
On the other hand, while I do think "Groverize" is glossing over a lot of complexity, like rebalancing, I'm not sure it's necessary to go into that much detail in a forum post. I don't think the post as written says anything which rules out that such complexities may be involved.

-----Original Message-----
From: Liu, Yi-Kai (Fed)
Sent: Wednesday, October 26, 2016 4:17 PM
To: Perlner, Ray (Fed) <ray.perlner@nist.gov>; Chen, Lily (Fed) <lily.chen@nist.gov>; Peralta, Rene (Fed) <rene.peralta@nist.gov>; Moody, Dustin (Fed) <dustin.moody@nist.gov>; Alperin-Sheriff, Jacob (Fed) <jacob.alperin-sheriff@nist.gov>; Bassham, Lawrence E (Fed) <lawrence.bassham@nist.gov>; Daniel Smith-Tone <daniel-c.smith@louisville.edu>; Jordan, Stephen P (Fed) <stephen.jordan@nist.gov>; Miller, Carl A. (Fed) <carl.miller@nist.gov>; Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov>
Subject: Re: First cut at a summary of our thinking on security strengths for the forum.

Hi Ray,

Sorry for not replying earlier! I just wanted to add two comments:

- I think we should ask people who are planning to submit a cryptosystem to us, whether they feel that they will be able to do the analysis we are requesting? Do they need more guidance? Do they have sufficient expertise?

- I'm a little worried that the advice about setting a certain classical security level, and then seeing how much the attacks can be "Groverized," might be too simplistic. Maybe we should say clearly that this is just a hypothetical example, and in a real attack, it's going to be more complicated. (For instance, once you use Grover search to speed up some steps in the attack, you can then change some of the parameters to re-balance the amount of work in the other steps of the attack.)

Cheers,

--Yi-Kai

_____
From: Perlner, Ray (Fed)
Sent: Wednesday, October 26, 2016 2:48:19 PM
To: Chen, Lily (Fed); Peralta, Rene (Fed); Moody, Dustin (Fed); Alperin-Sheriff, Jacob (Fed); Bassham, Lawrence E (Fed); Daniel Smith-Tone; Jordan, Stephen P (Fed); Liu, Yi-Kai (Fed); Miller, Carl A. (Fed); Smith-Tone, Daniel (Fed)
Subject: RE: First cut at a summary of our thinking on security strengths for the forum.

Note that the previous message was intended as a revised version of the pqc-forum post soliciting feedback on security strengths.

From: Perlner, Ray (Fed)

Sent: Wednesday, October 26, 2016 2:46 PM
To: Chen, Lily (Fed) <lily.chen@nist.gov>; Peralta, Rene (Fed) <rene.peralta@nist.gov>; Moody, Dustin (Fed) <dustin.moody@nist.gov>; Alperin-Sheriff, Jacob (Fed) <jacob.alperin-sheriff@nist.gov>; Bassham, Lawrence E (Fed) <lawrence.bassham@nist.gov>; 'Daniel Smith-Tone' <daniel-c.smith@louisville.edu>; Jordan, Stephen P (Fed) <stephen.jordan@nist.gov>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov>; Miller, Carl A. (Fed) <carl.miller@nist.gov>; Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov>
Subject: RE: First cut at a summary of our thinking on security strengths for the forum.

We got a lot of comments on our target security strengths section in the draft call for proposals. As a result, we feel it is appropriate to request some feedback before committing to an approach for measuring security in our final CFP. Here is a summary of our past and current thinking:
Previously we defined 5 security levels giving a classical and a quantum security strength. However, we were defining quantum security a little oddly, and we think this may have led to some misunderstandings. Our goal in defining these levels was to capture the practical cost in time and dollars of breaking a scheme with the listed security strength. The biggest factors we felt were not adequately captured by existing metrics for quantum security are:

1) The difficulty of parallelizing variants of Grover's algorithm and

2) The relative cost of classical vs quantum gates.

However, since most quoted figures in the literature for quantum "bits of security" don't take these things into account, we feel it was a mistake to use that language to describe what we were asking for.
Our current plan shares much with the previous approach. We still think it's reasonable to categorize submitted parameter sets into 5 rough security strength categories, where categories 1,3, and 5 are at least as hard to break as AES128, AES192, and AES256, respectively, and categories 2 and 4 are at least as hard to break as SHA256 and SHA384 respectively. However, we don't necessarily think that quantum security can really be captured by a single number: The practical cost of an attack will be parametrized at least by the maximum circuit depth that can be permitted by real world quantum gate times, and the relative cost of classical and quantum gates. So instead, our approach would be to say that, for any reasonable assumptions, regarding maximum circuit depth and relative quantum/classical cost, attacks against the schemes in a given security strength category should not be cheaper than attacks against the defining algorithm (e.g. something in security strength category 4 should be no easier to break, given any reasonable assumption, than SHA384.) For reference, we'd consider a maximum depth ranging from 2^40 to 2^90 logical gates, and a relative quantum/classical cost ranging from 1 to 2^40 to be reasonable.
We also wish to clarify that we do not, and did not intend to require that submitters provide different parameter sets for all 5 security levels. In our view, a parameter set with a higher security strength automatically satisfies the requirements for any of the lower security strengths. Our current suggestion is that submitters provide at least one parameter set meeting or exceeding security strength 4 or 5, and as many additional parameter sets as the submitter believes are necessary to take advantage of any security/performance tradeoffs offered by the design approach.
To demonstrate what these security strength categories will mean for submitters in practice, consider the following scenario: Assume you are submitting a postquantum algorithm where there is only one tunable parameter, corresponding to classical security, and there are no quantum attacks, aside from generic ones (i.e. those based on Grover's algorithm and related stuff like amplitude amplification and quantum walks.)

Then, in order to meet security strengths 1,3,5, you simply need to set classical security to equal 128, 192, 256 bits respectively.

Meeting security strength 2 will require some amount of classical security between 128 and 192 bits, and meeting security strength 4 will require some amount of classical security between 192 and 256 bits. Whether the required amount of classical security is at the low or high end of this range will depend upon how well the classical attacks against your scheme Groverize (i.e. how effective are the generic techniques for decreasing the cost of the classical attacks, using quantum computers.) If they Groverize well, you will need a classical security strength on the high end of the range, and if they Groverize poorly, you will need a classical security strength on the low end of the range.

One possible change we may consider making to the current approach would be to eliminate the security strengths based on hash functions. This would simplify the security analysis somewhat, by effectively making generic quantum cryptanalysis irrelevant to our evaluation criteria. However, it would leave us with no way to give credit to

algorithms, if the classical attacks against them are hard to Groverize. A number of commenters suggested making a change in the opposite direction. Some even suggested going so far as to treat an algorithm with 128 bits of classical security and no quantum speedup, as being equivalently strong to a 256-bit block cipher, since both have "128 bits of quantum security." We don't think this is reasonable. We can come up with plausible computation models where something with 192 bits of classical security and no quantum speedup might be as hard to break as AES 256 (and we can come up with plausible models where nothing with less than 256 bits of classical security is as hard to break as AES256) but we can't come up with a reasonable justification for treating something with much less than 192 bits of classical security as being as strong as AES 256.

From: Perlner, Ray (Fed)
Sent: Wednesday, October 26, 2016 10:15 AM
To: Chen, Lily (Fed) <lily.chen@nist.gov<mailto:lily.chen@nist.gov>>; Peralta, Rene (Fed) <rene.peralta@nist.gov<mailto:rene.peralta@nist.gov>>; Moody, Dustin (Fed) <dustin.moody@nist.gov<mailto:dustin.moody@nist.gov>>; Alperin-Sheriff, Jacob (Fed) <jacob.alperin-sheriff@nist.gov<mailto:jacob.alperin-sheriff@nist.gov>>; Bassham, Lawrence E (Fed) <lawrence.bassham@nist.gov<mailto:lawrence.bassham@nist.gov>>; Daniel Smith-Tone <daniel-c.smith@louisville.edu<mailto:daniel-c.smith@louisville.edu>>; Jordan, Stephen P (Fed) <stephen.jordan@nist.gov<mailto:stephen.jordan@nist.gov>>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov<mailto:yi-kai.liu@nist.gov>>; Miller, Carl A. (Fed) <carl.miller@nist.gov<mailto:carl.miller@nist.gov>>; Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov<mailto:daniel.smith@nist.gov>>
Subject: RE: First cut at a summary of our thinking on security strengths for the forum.

Here's how I would put it: Assume you have a postquantum algorithm where there is only one tunable parameter, corresponding to classical security, and there are no quantum attacks, aside from generic ones (i.e. those based on Grover's algorithm and related stuff like amplitude amplification and quantum walks.)

Then, in order to meet security strengths 1,3,5, you simply need to set classical security to equal 128, 192, 256 bits respectively. The effectiveness of Groverization is irrelevant.

Meeting security strength 2 will require some amount of classical security between 128 and 192 bits, and meeting security strength 4 will require some amount of classical security between 192 and 256 bits. Whether the required amount of classical security is at the low or high end of this range will depend upon how well the classical attacks against your scheme Groverize (i.e. how effective are the generic techniques for speeding up a classical attack with quantum computers.) If they Groverize well, you will need a classical security strength on the high end of the range, and if they Groverize poorly, you will need a classical security strength on the low end of the range.

Does this clarify things?
Should I stick some version of the above text into the message sent to the pqc-forum?

From: Chen, Lily (Fed)
Sent: Wednesday, October 26, 2016 3:23 AM
To: Perlner, Ray (Fed) <ray.perlner@nist.gov<mailto:ray.perlner@nist.gov>>; Peralta, Rene (Fed) <rene.peralta@nist.gov<mailto:rene.peralta@nist.gov>>; Moody, Dustin (Fed) <dustin.moody@nist.gov<mailto:dustin.moody@nist.gov>>; Alperin-Sheriff, Jacob (Fed) <jacob.alperin-sheriff@nist.gov<mailto:jacob.alperin-sheriff@nist.gov>>; Bassham, Lawrence E (Fed) <lawrence.bassham@nist.gov<mailto:lawrence.bassham@nist.gov>>; Daniel Smith-Tone <daniel-c.smith@louisville.edu<mailto:daniel-c.smith@louisville.edu>>; Jordan, Stephen P (Fed) <stephen.jordan@nist.gov<mailto:stephen.jordan@nist.gov>>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov<mailto:yi-kai.liu@nist.gov>>; Miller, Carl A. (Fed) <carl.miller@nist.gov<mailto:carl.miller@nist.gov>>; Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov<mailto:daniel.smith@nist.gov>>
Subject: Re: First cut at a summary of our thinking on security strengths for the forum.

I think this is absolutely a good approach, i.e. to reach pqc community about the topic. I like Ray's write up. Here are one point I think we might want to be more clear.

About whether to map the classical/quantum strength to hash functions, I think we can be more straightforward. My understanding is that we are separating algorithms based on whether and how a quantum speed up on classical attack

is effective (so called Groverizer). If my understanding is correct, then the current last paragraph reflects this meaning, however, a little hard to read. It also would be helpful to tell why we feel it is important to distinguish these two situations. Perhaps, we just want to know the best Groverizer can do on the classical attacks. Would it be possible if we simply say that we request algorithms with n bits classical security (n = 128, 192, and 256) and no quantum attacks, generic or Groverizing, can do better than Grover algorithm on AES with n bits key?

Please notice that the security strength levels in the CFP is for submitters. No single family can satisfy both, i.e. resistant to Groverizer and not resistant to Groverizer. It will not be realistic to map to all five levels. A family can properly map either to levels (1, 3, 5) or to levels (2, 4). However, when we select, we can distinguish these two situations. We can say that we are in favor of the algorithms, on which Groverizer is less effective. But in a long run, the Groverizer resistance will be only one of the many factors we are considering.

Attached please see a copy of ETSI report. quantum security is discussed in Section 5.3.2 in general and Sections 6.5, 7.5, 8.5, 9.4, and 10.5 for each family.

Lily

From: Perlner, Ray (Fed)
Sent: Tuesday, October 25, 2016 4:14:36 PM
To: Peralta, Rene (Fed); Moody, Dustin (Fed); Alperin-Sheriff, Jacob (Fed); Bassham, Lawrence E (Fed); Chen, Lily (Fed); Daniel Smith-Tone; Jordan, Stephen P (Fed); Liu, Yi-Kai (Fed); Miller, Carl A. (Fed); Smith-Tone, Daniel (Fed)
Subject: RE: First cut at a summary of our thinking on security strengths for the forum.

Er. I forgot to add: That advice only holds assuming there's no classical attack cheaper than 2**(2k) work. If there is, the classical attack determines the security level.

From: Perlner, Ray (Fed)
Sent: Tuesday, October 25, 2016 4:10 PM
To: Peralta, Rene (Fed) <rene.peralta@nist.gov<mailto:rene.peralta@nist.gov>>; Moody, Dustin (Fed) <dustin.moody@nist.gov<mailto:dustin.moody@nist.gov>>; Alperin-Sheriff, Jacob (Fed) <jacob.alperin-sheriff@nist.gov<mailto:jacob.alperin-sheriff@nist.gov>>; Bassham, Lawrence E (Fed) <lawrence.bassham@nist.gov<mailto:lawrence.bassham@nist.gov>>; Chen, Lily (Fed) <lily.chen@nist.gov<mailto:lily.chen@nist.gov>>; Daniel Smith-Tone <daniel-c.smith@louisville.edu<mailto:daniel-c.smith@louisville.edu>>; Jordan, Stephen P (Fed) <stephen.jordan@nist.gov<mailto:stephen.jordan@nist.gov>>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov<mailto:yi-kai.liu@nist.gov>>; Miller, Carl A. (Fed) <carl.miller@nist.gov<mailto:carl.miller@nist.gov>>; Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov<mailto:daniel.smith@nist.gov>>
Subject: RE: First cut at a summary of our thinking on security strengths for the forum.

Actually, the circuit size is proportional to 2**(2k)/ depth. So, you only can get the size down to 2**k if you can accommodate a circuit with depth 2**k (often not a realistic assumption.) That said, our guidance in that case should be quite simple. If 2k>128, you meet security strength 1, if 2k>192 you meet security strength 3, and if 2k > 256, you meet security strength 5. (although I guess you need to be careful about the constants. If the function you're trying to invert is much cheaper than AES to compute, then you need k to be a bit larger.)

From: Peralta, Rene (Fed)
Sent: Tuesday, October 25, 2016 4:02 PM
To: Moody, Dustin (Fed) <dustin.moody@nist.gov<mailto:dustin.moody@nist.gov>>; Alperin-Sheriff, Jacob (Fed) <jacob.alperin-sheriff@nist.gov<mailto:jacob.alperin-sheriff@nist.gov>>; Bassham, Lawrence E (Fed) <lawrence.bassham@nist.gov<mailto:lawrence.bassham@nist.gov>>; Chen, Lily (Fed) <lily.chen@nist.gov<mailto:lily.chen@nist.gov>>; Daniel Smith-Tone <daniel-c.smith@louisville.edu<mailto:daniel-c.smith@louisville.edu>>; Jordan, Stephen P (Fed) <stephen.jordan@nist.gov<mailto:stephen.jordan@nist.gov>>; Liu, Yi-Kai (Fed) <yi-kai.liu@nist.gov<mailto:yi-kai.liu@nist.gov>>; Miller, Carl A. (Fed) <carl.miller@nist.gov<mailto:carl.miller@nist.gov>>; Perlner, Ray (Fed) <ray.perlner@nist.gov<mailto:ray.perlner@nist.gov>>; Smith-Tone, Daniel (Fed) <daniel.smith@nist.gov<mailto:daniel.smith@nist.gov>>; Peralta, Rene (Fed)

<rene.peralta@nist.gov<mailto:rene.peralta@nist.gov>>
Subject: Re: First cut at a summary of our thinking on security strengths for the forum.

Can we give specific guidance to a submitter who wants to claim that the best a quantum computer can do to break her submission is to use Grover's algorithm on a space of size 2**(2k)? I don't know how to translate this into quantum circuit depth (I think the quantum circuit size is about 2**k).

Rene.

_____
From: Moody, Dustin (Fed)
Sent: Tuesday, October 25, 2016 3:38 PM
To: Alperin-Sheriff, Jacob (Fed); Bassham, Lawrence E (Fed); Chen, Lily (Fed); Daniel Smith-Tone; Jordan, Stephen P (Fed); Liu, Yi-Kai (Fed); Miller, Carl A. (Fed); Moody, Dustin (Fed); Peralta, Rene (Fed); Perlner, Ray (Fed); Smith-Tone, Daniel (Fed)
Subject: FW: First cut at a summary of our thinking on security strengths for the forum.

Ray has written up a post that we can submit to our pqc-forum to elicit feedback on our way of dealing with quantum security (see below). Let us know if you have any suggestions or comments. We would like to put this on the pqc-forum by the end of the week. Thanks!

Dustin

From: Perlner, Ray (Fed)
Sent: Tuesday, October 25, 2016 3:34 PM
To: Moody, Dustin (Fed) <dustin.moody@nist.gov<mailto:dustin.moody@nist.gov>>
Subject: First cut at a summary of our thinking on security strengths for the forum.

We got a lot of comments on our target security strengths section in the draft call for proposals. As a result, we feel it is appropriate to request some feedback before committing to an approach for measuring security in our final CFP. Here is a summary of our past and current thinking:

Previously we defined 5 security levels giving a classical and a quantum security strength. However, we were defining quantum security a little oddly, and we think this may have led to some misunderstandings. Our goal in defining these levels was to capture the practical cost in time and dollars of breaking a scheme with the listed security strength. The biggest factors we felt were not adequately captured by existing metrics for quantum security are

1) The difficulty of parallelizing variants of Grover's algorithm and

2) The relative cost of quantum vs classical gates.

However, since most quoted figures in the literature for quantum "bits of security" don't take these things into account, we feel it was a mistake to use that language to describe what we were asking for.

Our current plan shares much with the previous approach. We still think it's reasonable to categorize submitted parameter sets into 5 rough security strength categories, where categories 1,3, and 5 are at least as hard to break as AES128, AES192, and AES256, respectively, and categories 2 and 4 are at least as hard to break as SHA256 and SHA384 respectively. However, we don't necessarily think that quantum security can really be captured by a single number: The practical cost of an attack will be parametrized at least by the maximum circuit depth that can be

permitted by real world quantum gate times, and the relative cost of classical and quantum gates. So instead, our approach would be to say that, for any reasonable assumptions, regarding maximum circuit depth and relative quantum/classical cost, attacks against the schemes in a given security strength category should not be cheaper than attacks against the defining algorithm (e.g. something in security strength category 4 should be no easier to break, given any reasonable assumption, than SHA384.) For reference, we'd consider a maximum depth ranging from 2^40 to 2^90 logical gates, and a relative quantum/classical cost ranging from 1 to 2^40 to be reasonable.

We also wish to clarify that we do not, and did not intend to require that submitters provide different parameter sets for all 5 security levels. In our view, a parameter set with a higher security strength automatically satisfies the requirements for any of the lower security strengths. Our current suggestion is that submitters provide at least one parameter set meeting or exceeding security strength 4 or 5, and as many additional parameter sets as the submitter believes are necessary to take advantage of any security/performance tradeoffs offered by the design approach.

One possible change we may consider making to the current approach would be to eliminate the security strengths based on hash functions. This would simplify the security analysis somewhat, by effectively making generic quantum cryptanalysis irrelevant to our evaluation criteria. However, it would leave us with no way to give credit to algorithms, if the classical attacks against them are hard to Groverize. A number of commenters suggested making a change in the opposite direction. Some even suggested going so far as to treat an algorithm with 128 bits of classical security and no quantum speedup, as being equivalently strong to a 256-bit block cipher, since both have "128 bits of quantum security." We don't think this is reasonable. We can come up with plausible computation models where something with 192 bits of classical security and no quantum speedup might be as hard to break as AES 256 (and we can come up with plausible models where nothing with less than 256 bits of classical security is as hard to break as AES256) but we can't come up with a reasonable justification for treating something with much less than 192 bits of classical security as being as strong as AES 256.

Does the current approach seem sound?
What (if any) changes would you suggest?