From: pqc-forum@list.nist.gov on behalf of Mikhail Kudinov <mkudinov@qapp.tech>
Sent: Thursday, July 23, 2020 11:10 AM
To: pqc-forum
Subject: [pqc-forum] ROUND 3 OFFICIAL COMMENT: SPHINCS+

Dear all,

In this comment, we would like to point out a flaw of existing security proofs of the SPHINCS+ hash-based scheme. Particularly, we would like to pay attention to security proofs of the underlying WOTS+ scheme with preimage resistance (PRE) requirement replaced by second preimage resistance (SPR) + "at least two preimages for every image" requirements [see eq. (14) in Round 2 submission] or decisional second preimage resistance (DSPR) + SPR requirements [see Bernstein et al. "The SPHINCS+ signature framework" 2019].

Both of these approaches are based on the claim that in the case where the given image has several preimages under some cryptographic hash function, the original preimage is information-theoretically hidden among all preimages (see "Case 2" in the Proof of Theorem 2 in [Hülsing et al. "Mitigating Multi-Target Attacks in Hash-based Signatures" 2016] and "SM-DSPR success probability" in the proof of Claim 23 in [Bernstein et al. "The SPHINCS+ signature framework" 2019]). Though this claim is quite reasonable in the case of a single hash function query, the situation becomes much more complicated when one deals with a chain of hash functions like in the WOTS+ scheme.

Let h_i with i=1,...,w-1 be a hash function used to obtain a value at i'th level of the WOTS+ scheme from the one at (i-1)'th level. That is pk_j = h_{w-1}(h_{w-2}(… h_1(sk_j) …)), where sk_j and pk_j are elements of secret and public key respectively and w is a Winternitz parameter (commonly w = 16). Here we assume that all bitmasks are included in h_i. Let IMG_i be an image set of h_i, and let PREIMG_i(y) be a set of all preimages for given y taken from IMG_i. The proposed security proofs are based either on assumption that for each y one has |PREIMG_i(y)| > 1, or that it is computationally hard to recognize whether |PREIMG_i(y)| = 1 or not. The latter is called a DSPR property [D.J. Bernstein, A. Hülsing "Decisional second-preimage resistance: When does SPR imply PRE?" 2019].

Consider the set WOTS_IMG_i = h_i(h_{i-1}(... h_1({0,1}^n) ...)) that is an image set of the whole WOTS+ chain up to level i from a set of all possible secret keys {0,1}^n (n is a security parameter, typically equals to 256). One can reasonably expect that for a secure hash function built in the chain functions and i > 1, |WOTS_IMG_i| < |IMG_i| because of collisions at levels 1, … , i-1. Let WOTS_PREIMG_i(y) be a set of preimages of y under h_i belonging to WOTS_IMG_{i-1}. Having a Challenger's signature, a WOTS+-breaking adversary is able to choose a position in the chain where |WOTS_PREIMG_i(y)| = 1, even though |PREIMG_i(y)| > 1 for some known element y in the WOTS+ structure. In the result, the adversary manages to forge a signature avoiding breaking SPR property (because the forgery consists of the same element used by the Challenger), and by choosing elements having |PREIMG_i(y)| > 1 or |PREIMG_i(y)| = 1 with a proper probability, avoiding breaking DSPR property. Thus the reduction proof fails.

We note that the security proof of the original SPHINCS scheme [Bernstein et al. "SPHINCS: practical stateless hash-based signature" 2015] which is based on PRE+SPR+undetectability (UD) assumptions does not have this flaw, though shows lower security level for the same scheme parameters. We also note that the updated detailed security proof of the WOTS+ scheme based on PRE+SPR+UD assumptions can be found in https://arxiv.org/abs/2002.07419.

With kind regards,
Mikhail Kudinov, Evgeniy Kiktenko, Aleksey Fedorov
Russian Quantum Center (www.rqc.ru) and QApp (www.qapp.tech)
--

From:	pqc-forum@list.nist.gov on behalf of Andreas Hülsing <ietf@huelsing.net>
Sent:	Friday, July 31, 2020 7:27 AM
To:	Mikhail Kudinov; pqc-forum
Cc:	contact@sphincs.org
Subject:	Re: [pqc-forum] ROUND 3 OFFICIAL COMMENT: SPHINCS+

Dear Mikhail, Evgeniy, and Aleksey, dear all,

Due to the summer holidays, the following is not agreed on with the whole SPHINCS+ team and hence should be considered my personal opinion.

Thank you for evaluating our proposal and pointing us to this mistake in the proof. You are right that our reasoning that the input used to compute the image is information theoretically hidden is incorrect in the context of hash chains. It should be noted that

a) this a mere mistake in the proof and does not imply an attack, and

b) in any case, the non-tight proof still applies (as you also state).

The state of affairs is as follows. Looking into the existing proof, there seems not to be any quick fix for this issue that corrects the existing proof. However, this motivated another look at the whole thing and there seems to be an easier proof (following the non-tight proof) that works specifically for SPHINCS+. The very rough outline is as follows: The critical part of the tight proof was written for stateful schemes where WOTS is used to sign adversarialy chosen messages. However, in SPHINCS and SPHINCS+ the messages signed using WOTS are fully controlled by the honest user (or the reduction). That means that we only require security under known-message attacks for which case the non-tight proof becomes tight as the reduction does not have to guess.

We will work on an update after the vacation time is over. We will also continue to look into actual fixes for the existing proof as the above idea only works for stateful schemes when modelling the message digest function as a (quantum-accessible) random oracle, or making additional, stronger hardness assumptions like the existence of chameleon hash functions.

Best wishes,

Andreas

On 23-07-2020 17:09, Mikhail Kudinov wrote:

> Dear all,
>
> In this comment, we would like to point out a flaw of existing security proofs of the SPHINCS+ hash-based scheme. Particularly, we would like to pay attention to security proofs of the underlying WOTS+ scheme with preimage resistance (PRE) requirement replaced by second preimage resistance (SPR) + “at least two preimages for every image” requirements [see eq. (14) in Round 2 submission] or decisional second preimage resistance (DSPR) + SPR requirements [see Bernstein et al. “The SPHINCS+ signature framework” 2019].
>
> Both of these approaches are based on the claim that in the case where the given image has several preimages under some cryptographic hash function, the original preimage is information-theoretically hidden among all preimages (see "Case 2" in the Proof of Theorem 2 in [Hülsing et al. "Mitigating Multi-