From: makoto saitou <drmmakoto@gmail.com>
Sent: Wednesday, September 9, 2020 10:01 PM
To: pqc-comments
Cc: pqc-forum
Subject: ROUND 3 OFFICIAL COMMENT: Picnic
Attachments: Screenmemo_2020-09-10-10-02-38.png

Dear Sirs,
We are pleased to DESIGN the Digital US Dollar Project using Your Picnic for deploying Our ReEncryption Keychain instead of BIockchain.
Sincerely yours,
Makoto Saito in Japan

From: Robert Ransom <rransom.8774@gmail.com>
Sent: Monday, September 28, 2020 5:30 PM
To: pqc-comments
Cc: pqc-forum
Subject: ROUND 3 OFFICIAL COMMENT: Picnic

The U.S. patents listed as covering the MQDSS submission, US 8,522,033 and US 8,959,355, cover various categories of identification schemes and signature schemes where the public key is the result of applying “a multi-order multi-variable polynomial” to the secret key.

(US 8,522,033 contains typographical errors in which “y-f(s)=0” is written as “y=f(s)=0”; that error does not occur in US 8,959,355.)

LowMC is a multi-variable, multi-order polynomial, and all of the Picnic, Picnic2, and Picnic3 signature schemes and identification schemes appear to be covered by various claims of both patents.

Robert Ransom

From: D. J. Bernstein <djb@cr.yp.to>
Sent: Tuesday, May 4, 2021 2:46 AM
To: pqc-comments
Cc: pqc-forum
Subject: ROUND 3 OFFICIAL COMMENT: Picnic
Attachments: signature.asc

https://eprint.iacr.org/2021/578 claims, regarding "recently proposed instances of the Picnic signature scheme", that "2 out of 3 new instances do not achieve their claimed security level".

The comparison at the three security levels is to 2^128, 2^192, and 2^256 bit operations. Comparing to 2^143, 2^207, 2^272 would strengthen the claim to "3 out of 3". (The confusion here would have been avoided by the recommendation in https://blog.cr.yp.to/20161030-pqnist.html to avoid discretizing security levels in the first place.)

However, I would be opposed to NISTPQC taking this paper's claim as a reason to penalize Picnic. For example, at the middle level, the attack using 2^188 bit operations is bottlenecked by constant random access to 2^164 bits of memory. The cost of randomly accessing a bit within N bits of memory is approximately the cost of sqrt(N)/2^5 bit operations (see Section 6.6 of https://ntruprime.cr.yp.to/nist/ntruprime-20201007.pdf), in this case 2^77 bit operations, evidently multiplied by at least 2^164; overall at least 2^241, i.e., at least 2^34 times more bit operations than the 2^207 target. This gap is too big to be closed by a switch from a two-dimensional memory structure to an imaginary cubical memory structure, never mind the question of whether the expected lifetime of NISTPQC is long enough to consider such memory volumes.

To be clear, I'm concerned about the number of new pieces being pulled together in Picnic. This includes concern about the possibility of Picnic being vulnerable to further improvements in this line of papers.
But we've seen larger recent improvements in lattice attacks, including improvements that (based on the best estimates available, with many question marks, which is also worrisome) reduced security levels of all lattice submissions, and still NIST hasn't penalized those submissions.

---Dan

From: Greg Zaverucha <gregz@microsoft.com>
Sent: Tuesday, May 4, 2021 1:53 PM
To: pqc-comments
Cc: pqc-forum
Subject: ROUND 3 OFFICIAL COMMENT: Picnic

This a response to the recent paper by Itai Dinur that provides cryptanalysis of LowMC [1]. We have reviewed the work and discussed it, and found it's high-quality and improves the state of the art for LowMC cryptanalysis. It provides attacks on the parameter sets described in the LowMC Cryptanalysis Challenge [2], a cryptanalysis challenge we set up to build confidence in the LowMC parameter sets used by Picnic. As set out by the challenge, the attack Dinur presents recovers a LowMC private key from a plaintext/ciphertext pair, which corresponds to recovering a Picnic private key from the public key.

Judging the impact of the attack on Picnic3's classification in the NIST security levels is difficult. The time cost of the attack is slightly less than the cost of a brute force key recovery attack on AES, but requires a very large amount of memory. For example, at L1 the attack requires 2^130 time (bit operations) and 2^112 memory (bits) and at L5 2^245 time and 2^219 memory. It seems unlikely to us that this would be cheaper than a brute force attack, however, there is no generally accepted way to argue this more formally. (See Dan Bernstein's email to this list on 5/4/2021 for one way.)

We plan to specify additional parameters that use the LowMC parameters with a partial S-box layer (used in all Picnic instances before Round 3) and a change to the MPC parameters as in Picnic3. We can describe these as Picnic2 but with N=16 parties instead of N=64 (reverting many of the Picnic2->Picnic3 changes). The performance of these parameters will not be quite as good as Picnic3, however the security of LowMC is better understood, making them a more conservative option. Benchmarks for these parameters are given in [3], Table 4 (the boldface rows with N=16). At L1, the signatures are 1.1x larger, 2x slower to create and 1.3x slower to verify. For the Picnic parameter sets based on the ZKB++ proof system (e.g., Picnic-L1-FS, Picnic-L1-full) both options for LowMC are already present.

The Picnic Team

[1] Itai Dinur. Cryptanalytic Applications of the Polynomial Method for Solving Multivariate Equation Systems over GF(2). To appear at EUROCRYPT 2021. ePrint 2021/578 https://eprint.iacr.org/2021/578
[2] LowMC Crypanalysis Challenge. https://lowmcchallenge.github.io/
[3] Daniel Kales and Greg Zaverucha. Improving the Performance of the Picnic Signature Scheme. TCHES 2020. https://eprint.iacr.org/2020/427