
From: Ward Beullens <Ward.Beullens@esat.kuleuven.be>
Sent: Monday, January 15, 2018 4:55 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: OFFICIAL COMMENT: WalnutDSA
Attachments: walnutsa.pdf

Dear all,

The paper [1] describes an attack on an earlier version of WalnutDSA. In response, WalnutDSA was adapted to block this attack by using two braids as private key instead of one. The attached pdf describes a way to circumvent this adaptation, and shows how to use the method of [1] to attack the adapted scheme.

I have communicated with the designers, who have confirmed that the attack works. However, the new attack has the same limitation as the attack of [1] that the forged signatures are much longer than signatures made with the legitimate signing algorithm. WalnutDSA implicitly imposes a bound on the length of the signatures by requiring that the length of a signature (in terms of the number of artin generators) is encoded in a two byte value. The forgeries of [1] are much longer than 2^{16} , so the WalnutDSA scheme is not (yet) broken.

The attack paper [1] mentions "If an efficient algorithm to compute short factorizations exists, the increase in parameters q and N needed to achieve a sufficient level of security would then make WalnutDSA unsuitable for embedded devices." This statement now also holds for the version of WalnutDSA that is submitted to NIST. Moreover, the security of Walnut also depend on the hardness of this problem:

Given a long braid s and a pair (M, σ) , find a short (e.g. shorter than 2^{16}) braids s' such that $(M, \sigma) * s = (M, \sigma) * s'$.

On a separate note: I believe the security proof to be flawed. (I did not discuss this with the designers)

The reduction uses a EUF-CMA adversary to solve their hard problem (the REM problem). However, the EUF-CMA adversary can make signing queries, and the security proof does not describe a method to respond to these queries. Therefore, it seems like the proof only works for a key-only attack, where no signing queries have to be answered.

Moreover, the proof does not work in the quantum random oracle model, makes some assumptions and imposes restrictions on the attacker, so even for key-only attacks the practical value of the security proof seems very limited.

Kind regards,
Ward

[1] Hart, D., Kim, D., Micheli, G., Perez, G. P., Petit, C., & Quek, Y.
A Practical Cryptanalysis of WalnutDSATM.

WalnutDSA

Ward Beullens

January 15, 2018

1 Definitions

We use the notation of the WalnutDSA white paper[1]. We fix a set of T values throughout the document, and we denote by $*$ the action of B_N on $GL(N, \mathbb{F}_q) \times S_n$ of E-multiplication with respect to T .

Definition 1. The **hash** of a message $m \in \{0, 1\}^*$ is $E(H(m)) \in B_N$, which is a pure braid by construction.

Definition 2. For a braid $s \in B_N$ we define $P(s)$ to be equal to

$$(\mathbb{1}_N, e) * s \in GL(N, \mathbb{F}_q) \times S_N.$$

Definition 3. A signature $s \in B_N$ is valid for a message with hash g for the public key $(P(w), P(w'))$ if and only if

$$P(w) * s = P(g) * w',$$

where $*$ stands for E-multiplication with respect to some T -values.

Remark. The above definition of what it means for a signature to be valid is not completely identical to the definition of the WalnutDSA white paper, where only the matrix part of the above equation is required to hold. However, for all signatures that are produced by the legitimate Walnut signing algorithm, the permutation part of the above equation also holds. Therefore we can assume this stricter definition.

2 Two properties of WalnutDSA

The Walnut digital signature algorithm exhibits two interesting properties. Together, these properties generalize theorem 4 of [2]:

Theorem 1. Suppose that $w, w', s, g \in B_N$ are braids such that s is a valid signature for a messages with hash g for the public key $(P(w), P(w'))$. Then we have that s^{-1} is a valid signature for any message with hash g^{-1} for the public key $(P(w'), P(w))$.

Proof. By definition we have

$$P(w) * s = P(g) * w'.$$

Acting on this by s^{-1} from the left and using the definition of P we get

$$\begin{aligned} (\mathbb{1}_N, e) * w &= (\mathbb{1}_N, e) * g * w' * s^{-1} \\ &= (CB(g)_{\downarrow T}, e) * w' * s^{-1}, \end{aligned}$$

where the second equality uses the definition of E-multiplication and the fact that g is a pure braid. Multiplying the matrix part of this equality by $(CB(g)_{\downarrow T})^{-1}$ from the left (which is compatible with $*$) we get

$$(CB(g^{-1})_{\downarrow T}, e) * w = (\mathbb{1}_N, e) * w' * s^{-1},$$

or equivalently

$$P(g^{-1}) * w = P(w') * s^{-1}$$

which shows that s^{-1} is a valid signature for any message with hash g^{-1} for the public key $(P(w'), P(w))$. \square

Theorem 2. *Suppose that $w, w', w'', s_1, s_2, g_1, g_2 \in B_N$ are braids such that s_1 is a signature for a message with hash g_1 that is valid for the public key $(P(w), P(w'))$, and such that s_2 is a signature for a message with hash g_2 that is valid for the public key $(P(w'), P(w''))$. Then we have that $s_1 s_2$ is a valid signature for any message with hash $g_1 g_2$ for the public key $(P(w), P(w''))$.*

Proof. using the fact that s_1 is a valid signature for g_1 , and using the definition of E-multiplication we get

$$\begin{aligned} P(w) * s_1 * s_2 &= P(g_1) * w' * s_2 \\ &= (CB(g_1)_{\downarrow T} \cdot CB(w')_{\downarrow T} \cdot \sigma_{w'} (CB(s_2))_{\downarrow T}, \sigma_{w'} \circ \sigma_{s_2}), \end{aligned}$$

where $\sigma_{w'}$ denotes the permutation of w' . Using that s_2 is a signature for g_2 for the public key $(P(w'), P(w''))$ we can continue

$$\begin{aligned} P(w) * s_1 * s_2 &= (CB(g_1)_{\downarrow T} \cdot CB(g_2)_{\downarrow T} \cdot CB(w'')_{\downarrow T}, \sigma_{w''}) \\ &= P(g_1 g_2) * w'', \end{aligned}$$

which shows that $s_1 s_2$ is a valid signature for any message with hash $g_1 g_2$ for the public key $(P(w), P(w''))$. \square

3 Attack on WalnutDSA

Our attack uses the factorization attack of Hart et al. [2] on an earlier version of WalnutDSA as a black box. The earlier version of WalnutDSA uses only one braid as a private key. Hence, their attack works in the case $w = w'$.

Assumption. *Given a long enough list of signatures s_1, \dots, s_k for messages with hashes g_1, \dots, g_k respectively that are valid for a public key $(P(w), P(w))$, and a target hash g , the attack of [2] can forge a signature s that is valid for any message with hash g for the same public key.*

3.1 Description of the attack

We will show how an attacker can use a number of valid signature-hash pairs $(s_1, g_1), \dots, (s_k, g_k)$ that are valid for a public key $(P(w), P(w'))$ to forge a signature for any message that is valid for the same public key.

For any $i, j \in \{1, \dots, k\}$ we have that $s_i s_j^{-1}$ is a valid signature for any message with hash $g_i g_j^{-1}$ for the public key $(P(w), P(w'))$. This follows easily from theorems 1 and 2. Therefore, by collecting k signed messages for $(P(w), P(w'))$, the attacker can obtain $k(k-1)/2$ message-hash pairs for $(P(w), P(w'))$. This fact, in combination with the attack of [2] allows the attacker to forge a signature for any hash value g that is valid under the public key $(P(w), P(w'))$.

Let g be the hash of a message that the attacker wants to forge a signature for. As per the previous paragraph, the attacker can obtain a signature s for the hash $g g_1^{-1}$, which is valid for the public key $(P(w), P(w'))$. Theorem 2 then says that $s s_1$ is a valid signature for any message with hash $g g_1^{-1} g_1 = g$ for the public key $(P(w), P(w'))$, so this is a forgery that the attacker was looking for.

3.2 Limitations of the attack

The new attack has the same limitations as the attack of [2], and the countermeasures proposed in [2] should suffice to also block the new attack. The new attack produces forged signatures that are roughly twice as long as the forgeries of [2], but it is easier to collect many signature-hash pairs, which can help to find shorter forgeries.

References

- [1] Iris Anshel, Derek Atkins, Dorian Goldfeld, and Paul E Gunnells. Walnutdsa (tm): A quantum resistant group theoretic digital signature algorithm. *IACR Cryptology ePrint Archive*, 2017:58, 2017.
- [2] Daniel Hart, DoHoon Kim, Giacomo Micheli, Guillermo Pascual Perez, Christophe Petit, and Yuxuan Quek. A practical cryptanalysis of walnutdsatm.

From: Derek Atkins <datkins@securerf.com>
Sent: Tuesday, January 16, 2018 4:56 PM
To: ward.beullens@student.kuleuven.be; pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: WalnutDSA

Dear Ward, All,

As pointed out by Ward Beullens, it is possible with his method to generate extremely long forged signatures estimated at 2^{30} Artin generators long (which is $\sim 2^{32}$ bits). The WalnutDSA specification as submitted (which Ward Beullens notes), limits signature lengths. We have already tested the Dehornoy method, as suggested by both Ward Beullens and [1], to shorten forged signatures of this form and the result was only a 10% reduction, not nearly enough to reach valid signature lengths (even a 90% reduction wouldn't come close). Therefore, WalnutDSA, as specified in the NIST submission, remains a viable method.

In regards to the security proof, in our paper we define a forger to be a randomized algorithm that can make hash queries to a random oracle and signature queries to a simulator that does not know $\text{Priv}(S)$ but can simulate an honest signer. In our paper, we defined a signature query to be the message and the public key of the signer. The response to the query is a valid signature.

The forger is trying to generate a signature for a message that hasn't been queried before while the simulator is trying to break the hard problem of reversing E-multiplication. It is assumed that the simulator can simulate proper responses to the forger's signature queries in a way that is indistinguishable from query responses by a true signer (who knows the private key).

It is clearly stated in our paper that we assume the signatures produced by the forger are of the same type as those produced by a true signer, i.e., lie in a certain double coset of the braid group. This is a restrictive assumption which occurs because of the non-commutativity of the braid group and does not arise in security proofs for cryptosystems which make use of cyclic groups, for example. We also remark in our paper that Kobitz and Menezes [2] point out that "although it is a common approach in modern security proofs to restrict the capabilities of the adversary, it is important to show that certain classes of attacks can be ruled out."

Therefore, we submit that our security proof is not flawed in light of the perspective put forth by Kobitz and Menezes.

-- The WalnutDSA Team

[1] D. Hart; D. Kim; G. Micheli; G. Pascual Perez; C. Petit; Y. Quek, A Practical Cryptanalysis of WalnutDSA, preprint 2017.

[2] N. Kobitz; A. Menezes, Another look at "provable security," J. Cryptol. 20, 3–37 (2007).

On Mon, 2018-01-15 at 10:57 +0100, Ward Beullens wrote:

Dear all,

The paper [1] describes an attack on an earlier version of WalnutDSA. In

From: Ward Beullens <ward.beullens@student.kuleuven.be>
Sent: Thursday, January 18, 2018 11:28 AM
To: Derek Atkins; pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: WalnutDSA

Dear Derek, All,

I agree that WalnutDSA is not broken yet, but I don't think the attack is insignificant either. After all, the WalnutDSA team decided it was worthwhile to modify the original scheme in an attempt to block this attack, even at the cost of doubling the public and private key sizes. (Will this modification be kept, now that it is clear that it does not really block this attack?)

About the security proof: I don't think my concern was addressed. The proof uses a EUF-CMA forger, but does not give a way to answer the signing queries of this forger. You say "we define a forger to be a randomized algorithm that can make hash queries to a random oracle and signature queries to a simulator that does not know $\text{Priv}(S)$ but can simulate an honest signer."

This seems very peculiar to me, do you assume the existence of such a simulator? If this assumption were true, then WalnutDSA is not secure.

The security proof mentions "The Forger F is defined to be a randomized algorithm that on input a message $m \in \{0, 1\}^*$, a signer's public key $\text{Pub}(S)$, and a coin ρ , outputs a 4-tuple (m, h, g_ρ, s) , where $h = H(m)$ and $g_\rho = G_\rho(V)$ and $V \leftarrow \text{Cloak}$, $s \leftarrow \text{DC}_{m,V,H,G}$. It is assumed that the probability that (m, h, g_ρ, s) is a valid WalnutDSA-I signature is non-negligible."

This looks like a universal (the message can be freely chosen), key-only attack (there is no mention of signing queries). How does this relate to the claim that Walnut is EUF-CMA secure?

Kind regards,
Ward

Op 16/01/2018 om 22:55 schreef Derek Atkins:

Dear Ward, All,

As pointed out by Ward Beullens, it is possible with his method to generate extremely long forged signatures estimated at 2^{30} Artin generators long (which is $\sim 2^{32}$ bits). The WalnutDSA specification as submitted (which Ward Beullens notes), limits signature lengths. We have already tested the Dehornoy method, as suggested by both Ward Beullens and [1], to shorten forged signatures of this form and the result was only a 10% reduction, not nearly enough to reach valid signature lengths (even a 90% reduction wouldn't come close). Therefore, WalnutDSA, as specified in the NIST submission, remains a viable method.

In regards to the security proof, in our paper we define a forger to be a randomized algorithm that can make hash queries to a random oracle and signature queries to a simulator that does not know $\text{Priv}(S)$ but can simulate an honest signer. In our paper, we defined a signature query to be the message and the public key of the signer. The response to the query is a valid signature.

From: Blackburn, S <S.Blackburn@rhul.ac.uk>
Sent: Monday, January 22, 2018 5:55 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: OFFICIAL COMMENT: WalnutDSA
Attachments: Walnut_private_key_2.pdf

Dear All,

There are fewer than 2^{512} possibilities for each half of the public key in WalnutDSA. This makes the scheme vulnerable to 'square root' attacks (using an approach similar to Pollard-rho). It is recommended that parameter sizes should be increased so this approach is no longer feasible.

More details are given in the attached pdf, which is also available at the following link:

https://www.dropbox.com/s/dyje618qfb4zvsj/Walnut_private_key_2.pdf?dl=0

Yours,

Simon Blackburn

Simon R. Blackburn
Professor of Pure Mathematics
Department of Mathematics
Royal Holloway University of London
Egham, Surrey TW20 0EX, United Kingdom
Tel: (+44) (0)1784 443422
E-mail: S.Blackburn@rhul.ac.uk
Web: <http://www.ma.rhul.ac.uk/sblackburn>

Recovering a Private Key in WalnutDSA

Simon R. Blackburn
Department of Mathematics
Royal Holloway University of London
Egham, Surrey TW20 0EX, United Kingdom
e-mail: `s.blackburn@rhul.ac.uk`.

January 22, 2018

Abstract

There are fewer than 2^{512} possibilities for each half of the public key in WalnutDSA. This makes the scheme vulnerable to ‘square root’ attacks (using an approach similar to Pollard-rho). It is recommended that parameter sizes should be increased so this approach is no longer feasible.

1 Introduction

WalnutDSA(TM) is a digital signature scheme proposed by SecureRF for the NIST Post-quantum standardization call, using techniques from braid group theory. Parameters for two security levels, at 128- and 256-bits, have been defined, as well as 40-bit test parameters. This note describes an attack that recovers an equivalent private key of a signature from the corresponding public key. Analysis is given to support the claim that the attack recovers an equivalent private key for 256-bit parameters using significantly fewer than 2^{256} operations (one reasonable estimate being as little as 2^{168} operations). The 128-bit parameters are possibly also affected, with one estimate for deriving a private key being as little as 2^{108} operations.

The current version of WalnutDSA is specified by Anshel, Atkins, Goldfeld and Gunnells in [1, 3]. There has been a cryptanalysis due to Hart, Kim,

Micheli, Perez, Petit and Quek [5] of the original proposed scheme [2], which caused WalnutDSA to be significantly revised by introducing two braids (rather than the original one) as the private key. A recent note due to Beullens [4] shows that the new version of WalnutDSA remains vulnerable to the attack of Hart *et al.*

We will use the notation in [1, 3] without comment, and we will assume knowledge of these papers.

2 Recovering the private key

Let X be the set of pairs (M, σ) , where M is an $N \times N$ matrix over \mathbb{F}_q , and where $\sigma \in S_N$ is a permutation of $\{1, 2, \dots, N\}$. We write $1 \in X$ for the element (I, e) , where I is the identity $N \times N$ matrix, and where e is the identity permutation. The braid group B_N acts on X (on the right) via E-multiplication: for $x \in X$ and $g \in B_N$, we write $x * g \in X$ for the E-multiplication of x and g . We write Ω for the orbit of 1 under this action, so $\Omega = 1 * B_N$.

The private key consists of two N -string braids $S, S' \in B_N$, each a product of Artin generators of length ℓ . The public key consists of the pair of elements $(1 * S, 1 * S') \in \Omega^2$.

The appendix describes an algorithm (using standard techniques) that takes as input an element $\omega \in \Omega$, and outputs an element $g \in B_N$ such that $1 * g = \omega$. The algorithm runs in $\sqrt{|\Omega|}$ operations, and uses insignificant memory. We will apply this algorithm twice, first with $\omega = 1 * S$ and secondly with $\omega = 1 * S'$, to recover an equivalent private key $(\tilde{S}, \tilde{S}') \in B_N \times B_N$ for WalnutDSA. Each braid in the equivalent private key is a product of length 2λ in the Artin generators for some integer λ . We choose λ to be as small as possible so that the elements $1g$ are approximately uniformly distributed in Ω when g is a length λ word in the Artin generators. An accurate value for λ seems hard to determine, but we give a heuristic estimate for λ using the following argument. A necessary condition for λ is that the number of braids of length λ must be at least $|\Omega|$. Using the approximations in Anshel *et al.* [3] for the number of braids of length λ , we set λ to be the smallest integer such that

$$|\Omega| < (2^\lambda/\lambda)(N-1) \binom{\lambda-2+N}{N-1}.$$

The table below gives the parameters of the attack for 40-, 128- and 256-bit security levels (SL). The size of Ω is given as a range of possible values, computed as follows. The upper bound is $N! q^{N(N-1)}$, as there are $N!$ choices for the permutation σ and $q^{N(N-1)}$ choices for an $N \times N$ matrix over \mathbb{F}_q with last row $(0, 0, \dots, 0, 1)$. The lower bound of $N! q^{N(N-3)}$ uses the estimate for the number of choices for M that is used by the authors of the scheme [3, Section 11] in their security analysis.

SL	N	q	$ \Omega $	$\sqrt{ \Omega }$	ℓ	2λ
40	8	2^5	$[2^{215.3}, 2^{295.3}]$	$[2^{108}, 2^{148}]$	132	$[360, 514]$
128	8	2^5	$[2^{215.3}, 2^{295.3}]$	$[2^{108}, 2^{148}]$	132	$[360, 514]$
256	8	2^8	$[2^{335.3}, 2^{463.3}]$	$[2^{168}, 2^{232}]$	287	$[592, 842]$

3 Conclusions

The attack presented here recovers an equivalent 256-bit private key using significantly fewer than 2^{256} operations (between 2^{168} and 2^{232} operations). It is quite possible that a 128-bit key might also be recovered in fewer than 2^{128} operations using this approach (as little as 2^{108} operations is one estimate). The 40-bit keys (for testing) are unaffected by this attack.

Experiments might provide more precise length estimates, but it seems likely that the equivalent private key will be comparable in length to the true private key (at most a factor of 4 longer if we use the estimates of 2λ above). This is likely to lead to forged signatures of about the right length (particularly if we reduce the length of cloaking elements, which we can do as we no longer care to keep keys secret). (Previous attacks [4, 5] produce signatures of length approximately 2^{25} . These signatures are significantly longer than signatures produced with the private key, and so would fall foul of the limit on a signature length proposed in the specification.)

Ward Beullens (who kindly looked at a draft of this note) points out that the running time of the attack can be improved by a factor of $\sqrt{N!}$, at the expense of possibly slightly longer forgeries, by modifying the algorithm in the appendix to take the application setting into account. The algorithm in the appendix should be modified to choose $g_1, g_2 \in G$ so that $1g_1 \in \Omega$ and $\omega g_2^{-1} \in \Omega$ always have a trivial permutations associated with them.

A minor additional improvement to the algorithm in the appendix can be made by mandating that the determinant of the matrices associated with $1g_1 \in \Omega$ and ωg_2^{-1} are both 1.

Recommendation: The attack presented here is exponential, and so can be avoided by increasing parameter sizes. To prevent this attack for k -bit security level, parameters should be chosen so that $q^{N(N-3)-1} > 2^{2k}$. So revising parameters at both 128-bit and 256-bit security levels is recommended.

The authors of WalnutDSA have looked a draft of this note, and confirmed that they believe the attack to be correct. They suggest increasing N from 8 to 10 to prevent this attack.

References

- [1] Iris Anshel, Derek Atkins, Dorian Goldfeld and Paul E. Gunnells, ‘The Walnut digital signature algorithm (TM) specification’, *NIST post-quantum cryptography standardization project*. <https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions>.
- [2] Iris Anshel, Derek Atkins, Dorian Goldfeld and Paul E. Gunnells, ‘WalnutDSA (TM): A quantum-resistant digital signature algorithm’, *Cryptology ePrint archive*, Report 2017/1160, 18 September 2017. <https://eprint.iacr.org/2017/058>.
- [3] Iris Anshel, Derek Atkins, Dorian Goldfeld and Paul E. Gunnells, ‘WalnutDSA (TM): A quantum-resistant digital signature algorithm’, *Cryptology ePrint archive*, Report 2017/1160, 30 November 2017. <https://eprint.iacr.org/2017/058>.
- [4] Ward Beullens, ‘WalnutDSA’, *NIST Official Comment*, 15 January 2018.
- [5] Daniel Hart, DoHoon Kim, Giacomo Micheli, Guillermo Pascual Perez, Christophe Petit and Yuxuan Quek, ‘A practical cryptanalysis of WalnutDSA (TM)’, *Cryptology ePrint archive*, Report 2017/1160, 29 November 2017. <https://eprint.iacr.org/2017/1160>.

A An algorithm from computational group theory

Let Ω be a finite set, and let G be a group that acts transitively on Ω (on the right). Let $1 \in \Omega$ be fixed. Suppose we are given $\omega \in \Omega$, and we wish to

find $g \in G$ such that $1g = \omega$. We present an algorithm (using standard cryptographic techniques akin to the Pollard-rho discrete logarithm algorithm) that solves this problem using $O(\sqrt{|\Omega|})$ operations (in expectation).

First, choose a pseudorandom number generator r that takes a seed $x \in \Omega$ and outputs a pair $r(x) = (g, a)$, where $g \in G$ and $a \in \{0, 1\}$. (In our application, we would fix a positive integer λ , and set g to be a random product of Artin generators of length λ . So $g = \prod_{i=1}^{\lambda} b_{n_i}^{\epsilon_i}$, where the integers $n_i \in \{1, 2, \dots, N-1\}$ and $\epsilon_i \in \{-1, 1\}$ are outputs of our pseudorandom number generator. We choose λ large enough so that the elements $1 * g$ are approximately uniform in Ω .)

We then define a function $f : \Omega \rightarrow \Omega$ by

$$f(x) = \begin{cases} 1g & \text{if } r(x) = (g, 0), \\ \omega g^{-1} & \text{if } r(x) = (g, 1). \end{cases}$$

Next, we find a collision for f : elements $x_1, x_2 \in \Omega$ with $f(x_1) = f(x_2)$. This can be done using standard collision-finding techniques (such as Floyd cycle finding) in $O(\sqrt{|\Omega|})$ expected time, and constant memory.

Now, $r(x_1) = (g_1, a_1)$ and $r(x_2) = (g_2, a_2)$ for some group elements $g_1, g_2 \in G$ and bits $a_1, a_2 \in \{0, 1\}$. We repeat this process using different pseudorandom generators (and so different functions f) until we have a collision with $a_1 \neq a_2$. The event that $a_1 \neq a_2$ occurs with probability about $1/2$, provided that the distribution of the elements $1g_1 \in \Omega$ and $\omega g_2^{-1} \in \Omega$ are each approximately uniformly distributed in Ω . So the expected number of functions f we need to consider is constant (just 2). So we can find a collision of this restricted form in $O(\sqrt{|\Omega|})$ expected time.

Without loss of generality, suppose $a_1 = 0$ and $a_2 = 1$. By the definition of f , we have $1g_1 = \omega g_2^{-1}$, and so $g = g_1 g_2$ is the element we are seeking.

As a final comment, we note that the following algorithm (which uses $O(\sqrt{|\Omega|})$ memory) is an alternative to the algorithm above: First, choose $O(\sqrt{|\Omega|})$ elements $g_1 \in G$ at random. For each element, compute and store the pair $(1g_1, g_1)$. Store the pairs in a data structure so that the pair with a given first coordinate (if it exists) can be efficiently retrieved. Secondly, choose elements $g_2 \in G$ at random. For each element, compute ωg_2^{-1} . Repeat until ωg_2^{-1} is equal to the first coordinate of a pair generated earlier. Then we have found g_1 and g_2 such that $\omega g_2^{-1} = 1g_1$, and so we may return $g = g_1 g_2$.

From: Derek Atkins <datkins@securerf.com>
Sent: Monday, January 22, 2018 12:12 PM
To: pqc-forum@list.nist.gov; pqc-comments
Subject: OFFICIAL COMMENT: WalnutDSA typographic error

Dear All,

We have found a typographic error in the WalnutDSA specification in section 5.1 on page 5. In the document we incorrectly repeated the use of variable, so where the spec says "We assume the hash output M is 4l bits long" this should be "4d bits", meaning we assume that the hash length is a multiple of 4 bits.

We are sorry for the error.

The WalnutDSA Team

--

Derek Atkins
Chief Technology Officer
SecureRF Corporation

Office: 203.227.3151 x1343
Direct: 617.623.3745
Mobile: 617.290.5355
Email: DAtkins@SecureRF.com

This email message may contain confidential, proprietary and / or legally privileged information and intended only for the use of the intended recipient(s) and others specifically authorized. Any disclosure, dissemination, copying, distribution or use of the information contained in this email message, including any attachments, to or by anyone other than the intended recipient is strictly prohibited. If you received this in error, please immediately advise the sender by reply email or at the telephone number above, and then delete, shred, or otherwise dispose of this message.

From: Derek Atkins <datkins@securerf.com>
Sent: Monday, January 22, 2018 10:41 PM
To: S.Blackburn@rhul.ac.uk; pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: WalnutDSA

Dear Simon,

Thank you for pointing this out. A very small parameter increase from N=8 to N=10 will fix this with only a small increase to sizes and performance.

Thanks,

The WalnutDSA Team

On Mon, 2018-01-22 at 10:54 +0000, Blackburn, S wrote:

Dear All,

There are fewer than 2^{512} possibilities for each half of the public key in WalnutDSA. This makes the scheme vulnerable to 'square root' attacks (using an approach similar to Pollard-rho). It is recommended that parameter sizes should be increased so this approach is no longer feasible.

More details are given in the attached pdf, which is also available at the following link:

https://www.dropbox.com/s/dyje618qfb4zvsj/Walnut_private_key_2.pdf?dl=0

Yours,

Simon Blackburn

Simon R. Blackburn
Professor of Pure Mathematics
Department of Mathematics
Royal Holloway University of London
Egham, Surrey TW20 0EX, United Kingdom
Tel: (+44) (0)1784 443422
E-mail: S.Blackburn@rhul.ac.uk
Web: <http://www.ma.rhul.ac.uk/sblackburn>

--

Derek Atkins
Chief Technology Officer
SecureRF Corporation

Office: 203.227.3151 x1343
Direct: 617.623.3745
Mobile: 617.290.5355
Email: DAtkins@SecureRF.com

From: Ward Beullens <ward.beullens@student.kuleuven.be>
Sent: Tuesday, January 23, 2018 11:28 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: [pqc-forum] OFFICIAL COMMENT: WalnutDSA

Dear all,

In order to verify a WalnutDSA signature-message pair (s,m) one first converts the message m into a braid, and then one checks if this braid satisfies some condition involving the signature s .

The way that messages are converted to braids consists of 2 steps:

- 1) Compute a 512-bit hash digest of the message h . (For the 256-bit security parameters)
- 2) Convert the hash to a braid with the encoder algorithm described in section 7.

The problem with this approach is that the second step is not injective (e.g $0x0044$, $0x4400$, $0x0404$, $0x0440$ are all mapped to the same braid g_1^6), and that the number of braids that are encodings of messages is significantly less than 2^{512} . This opens up the possibility of an EUF-CMA attack. An attacker finds two messages m_1 and m_2 that are converted to the same braid, asks the signing algorithm for a valid signature s for m_1 , and returns the valid signature-message pair (s,m_2) .

Using a recurrence relation I calculated that the number of different encoding of the 2^{512} possible hash values is at most $2^{483.67}$, so collision finding takes roughly 2^{242} computations, slightly less than the target of 2^{256} .

The solution to the problem is to use a hash function with a longer output, or to use an encoding function which is injective. I believe both methods will lead to a slight increase to the size of the signatures.

Kind regards,
Ward

--

You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
Visit this group at <https://groups.google.com/a/list.nist.gov/group/pqc-forum/>.

From: Derek Atkins <datkins@securerf.com>
Sent: Tuesday, January 23, 2018 11:48 AM
To: ward.beullens@student.kuleuven.be; pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: WalnutDSA

Dear Ward, all,

Regarding your first comment, the structure of the security proof we give is modeled on that of "Security Proofs and Signature Schemes" by Pointcheval and Stern (EUROCRYPT 1996). One has to have access to valid signatures for chosen messages. In their paper, the authors give two models for this (cf. Fig 4 in their paper). In the first model the attacker has access to a signer Sigma with a private key who must make queries to a hash function f in the process of producing valid signatures. In the second, the signer Sigma is replaced by a simulator S . The simulator plays the same role; it produces signatures on message inputs, but does not do so by querying the hash function f using the private key. We will add another explicit reference to this paper in ours. The point of the proof is that if one assumes the existence of such a simulator, then one could use it to break the hard problem.

For your second comment, the forger uses the simulator to produce the signatures. This was said explicitly at the beginning of the security proof.

If you need further clarification of the proof we can discuss it further offline.

The WalnutDSA Team

On Thu, 2018-01-18 at 17:28 +0100, Ward Beullens wrote:

Dear Derek, All,

I agree that WalnutDSA is not broken yet, but I don't think the attack is insignificant either. After all, the WalnutDSA team decided it was worthwhile to modify the original scheme in an attempt to block this attack, even at the cost of doubling the public and private key sizes. (Will this modification be kept, now that it is clear that it does not really block this attack?)

About the security proof: I don't think my concern was addressed. The proof uses a EUF-CMA forger, but does not give a way to answer the signing queries of this forger. You say "we define a forger to be a randomized algorithm that can make hash queries to a random oracle and signature queries to a simulator that does not know $\text{Priv}(S)$ but can simulate an honest signer."

This seems very peculiar to me, do you assume the existence of such a simulator? If this assumption were true, then WalnutDSA is not secure.

The security proof mentions "The Forger F is defined to be a randomized algorithm that on input a message $m \in \{0, 1\}^*$, a signers public key $\text{Pub}(S)$, and a coin ρ , outputs a 4-tuple (m, h, g_ρ, s) , where $h = H(m)$ and $g_\rho = G_\rho(V)$ and $V \leftarrow \text{Cloak}$, $s \leftarrow \text{DC}_m, V, H, G$. It is assumed that the probability that (m, h, g_ρ, s) is a valid WalnutDSA-I signature is non-negligible."

This looks like a universal (the message can be freely chosen), key-only attack (there is no mention of signing queries). How does this relate to the claim that Walnut is EUF-CMA secure?

Kind regards,
Ward

From: Derek Atkins <datkins@securef.com>
Sent: Tuesday, January 23, 2018 11:50 AM
To: ward.beullens@student.kuleuven.be; pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: WalnutDSA

Dear Ward, All,

Thank you for bringing this to our attention. The encoder can easily be fixed by changing to a 2-bit encoder where each 2-bit input maps directly to a generator g_i . This minor change would make the encoder properly injective.

Thanks,

The WalnutDSA Team

On Tue, 2018-01-23 at 17:28 +0100, Ward Beullens wrote:

Dear all,

In order to verify a WalnutDSA signature-message pair (s,m) one first converts the message m into a braid, and then one checks if this braid satisfies some condition involving the signature s .

The way that messages are converted to braids consists of 2 steps:

- 1) Compute a 512-bit hash digest of the message h . (For the 256-bit security parameters)
- 2) Convert the hash to a braid with the encoder algorithm described in section 7.

The problem with this approach is that the second step is not injective (e.g $0x0044$, $0x4400$, $0x0404$, $0x0440$ are all mapped to the same braid g_1^6), and that the number of braids that are encodings of messages is significantly less than 2^{512} . This opens up the possibility of an EUF-CMA attack. An attacker finds two messages m_1 and m_2 that are converted to the same braid, asks the signing algorithm for a valid signature s for m_1 , and returns the valid signature-message pair (s,m_2) .

Using a recurrence relation I calculated that the number of different encoding of the 2^{512} possible hash values is at most $2^{483.67}$, so collision finding takes roughly 2^{242} computations, slightly less than the target of 2^{256} .

The solution to the problem is to use a hash function with a longer output, or to use an encoding function which is injective. I believe both methods will lead to a slight increase to the size of the signatures.

Kind regards,
Ward

--

Derek Atkins
Chief Technology Officer
SecureRF Corporation

Office: 203.227.3151 x1343

From: Ward Beullens <ward.beullens@student.kuleuven.be>
Sent: Tuesday, January 23, 2018 4:15 PM
To: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: WalnutDSA

Dear WalnutDSA team,

I had a look at the paper you reference. An entire page of the security proof is devoted to showing how the simulator S is constructed (i.e. the proof of Lemma 8). This part is missing from the WalnutDSA security proof and this is basically the concern that I raised.

There is a different way to see that something is wrong with the security proof. Consider these 2 parameters: the length of the output of the hash function that is used and the length limit that has to be imposed on the signatures. If these parameters are chosen poorly, the scheme is not secure, so in that case the security proof should fail somewhere. However, these 2 parameters are not mentioned in the security proof, nor in the formulation of the hard problem that EUF-CMA security is supposedly reduced to.

If the security proof is correct, can you please pinpoint where the security proof fails if

- 1) The length of the hash function that is used is too small, and
- 2) The verification algorithm does not impose a length limit on the signatures.

Kind regards,
Ward

Op 23/01/2018 om 17:47 schreef Derek Atkins:

Dear Ward, all,

Regarding your first comment, the structure of the security proof we give is modeled on that of "Security Proofs and Signature Schemes" by Pointcheval and Stern (EUROCRYPT 1996). One has to have access to valid signatures for chosen messages. In their paper, the authors give two models for this (cf. Fig 4 in their paper). In the first model the attacker has access to a signer Σ with a private key who must make queries to a hash function f in the process of producing valid signatures. In the second, the signer Σ is replaced by a simulator S . The simulator plays the same role; it produces signatures on message inputs, but does not do so by querying the hash function f using the private key. We will add another explicit reference to this paper in ours. The point of the proof is that if one assumes the existence of such a simulator, then one could use it to break the hard problem.

For your second comment, the forger uses the simulator to produce the signatures. This was said explicitly at the beginning of the security proof.

If you need further clarification of the proof we can discuss it further offline.

The WalnutDSA Team

On Thu, 2018-01-18 at 17:28 +0100, Ward Beullens wrote:

Dear Derek, All,

I agree that WalnutDSA is not broken yet, but I don't think the attack is insignificant either. After all, the WalnutDSA team decided it was worthwhile to modify the original

From: Ward Beullens <ward@beullens.com>
Sent: Thursday, February 01, 2018 10:21 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: OFFICIAL COMMENT: WalnutDSA
Attachments: WalnutDSAScript.sage.sage

Dear All,

There is an attack that breaks EUF-CMA security of the walnutDSA scheme with 2^{56} operations for the 256 security bit parameters, and 2^{35} operations for the 128 bit security level. I have communicated with the designers, they have confirmed that the attack seems to work, and they proposed countermeasures.

My previous attack exploited the fact that the range of E composed with H is smaller than 2^{512} , in order to do a collision attack. Now I have found that this range, after the composition with P gets much smaller. So, similar to my previous attack, one can easily find two messages m_1 and m_2 such that $P(E(H(m_1))) = P(E(H(m_2)))$, breaking EUF-CMA security.

Using sage, (I attached the script) I picked a sample of random hashes h_i and computed the matrix part of $P(E(h_i))$ for these hashes. Then I calculated that these matrices span a subspace of dimension only 14. This means that $P(E(H(m)))$ can reach at most $(2^8)^{14} = 2^{112}$ values, so a collision search costs roughly 2^{56} steps, much less than the security level of 2^{256} that was aimed for.

For the 128 bit security parameters the collision finding would require roughly $\sqrt{(2^5)^{14}} = 2^{35}$ E-multiplications, so it should be feasible to demonstrate this attack in practice.

Kind regards,
Ward

```

import random
import itertools

N = 8;
B = BraidedGroup(N, 'b');
K = GF(2^8);
MS = MatrixSpace(K, N, N);

B.inject_variables();

S = Set();

def RandomNonzeroElement():
    while True:
        r = K.random_element()
        if r != 0:
            return r;

def PBG(i, j):
    return B([a for a in range(j-1, i, -1)] + [i, i] + [-a for a in range(i+1, j)]);

# Choosing T values
Tvals = [RandomNonzeroElement() for i in range(0, N)];
Tvals[0] = K.one()
Tvals[1] = K.one()

CBOne = (MS.identity_matrix(), Permutations(N).identity())

#One step of E-multiplication
def EmulStep(x, g):
    M, p = x;
    M = copy(M);
    p = copy(p);

    if (g>0) :
        a = Tvals[p(g)-1];
        b = -a;
        c = 1;
    else :
        a = 1;
        b = -Tvals[p(-g+1)-1]^(-1);
        c = -b;

    g = abs(g);

    if(g>1):
        M.set_column(g-2, M.column(g-2) + a*M.column(g-1));

    M.set_column(g, M.column(g) + c*M.column(g-1));
    M.set_column(g-1, b*M.column(g-1));
    return (M, B([g]).permutation()*p);

#E-multiplication
def Emul(x, b):
    return reduce(EmulStep, [x]+list(b.Tietze()));

#The number of generators per hash value
Length = 256;
#The number of hash values
Samples = 40;

#The generators of the free group

```

```

WalnutDSAScript.sage.sage
G = [ PBG(1, N), PBG(3, N) , PBG(5, N) , PBG(7, N) ]
Bi gMatri x = matri x(K, N*N, Sampl es);
for i i n range(0, Sampl es):
    x = CBOne;
    for j i n range(0, Length):
        x = Emul (x, random. choi ce(G));
    m, p = x
    Bi gMatri x. set_col umn(i, m. l i st())
    print("progress = "+str(100. 0*(i)/Sampl es)+"%")

print("Di mensi on of span of P(E(h_i)) i s :")
print(Bi gMatri x. rank())

```

From: Derek Atkins <datkins@securerf.com>
Sent: Friday, February 02, 2018 1:33 PM
To: ward@beullens.com; pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: WalnutDSA
Attachments: GeneralWalnutEncoder.sage

Dear Ward, All,

Thank you for your comment. You already had pointed out an issue with the encoder that we agreed with and that we would address with one of several available options. Moreover, we acknowledged this related comment which points out that the "still-unchanged" encoding as given in the specification resulted in a too-small dimensional (14) subspace generated by the images of the public keys. As we noted in our most recent response to you, the introduction of an alternate encoder (see following suggestion) will address both of these issues.

There are many different ways to injectively encode messages into the free subgroup of the pure braid group, and we can easily opt for a method that takes this into account. One possible way that is similar to what we already do is to break the hashed message into 128 (or 256) 2-bit blocks and to choose a fixed finite ordering S of certain 4-tuples of generators. Then as one traverses the blocks one selects a generator from the current tuple from S ; at each successive block one takes the next tuple from S , and then cycles back to the beginning when the end of S is reached. More precisely, suppose we work in B_{12} . One could take S to be the list

[579B,468A,3579,2468,1357,2468,3579,468A]

and if the hashed message has 2-bit blocks

BLK1, BLK2, ... BLK128

then one does the following:

use BLK1 to select from g_5, g_7, g_9, g_B

use BLK2 to select from g_4, g_6, g_8, g_A

...

use BLK5 to select from g_1, g_3, g_5, g_7

use BLK6 to select from g_2, g_4, g_6, g_8

...

use BLK8 to select from g_4, g_6, g_8, g_A

use BLK9 to select from g_5, g_7, g_9, g_B

...

and so on. It is easy to verify that one obtains a large subspace this way (see attached Sage program). For example, in B_{12} one finds a 102-dimensional subspace generated by the public keys. Over F_{256} this subspace contains 2^{816} elements, which should be more than sufficient. Moreover, having the entries repeat in this manner not only remains injective, continues to hold a dimension of 102 in B_{12} , but also results in a shorter encoding which means a shorter signature.

The WalnutDSA Team

On Thu, 2018-02-01 at 16:21 +0100, Ward Beullens wrote:

Dear All,

There is an attack that breaks EUF-CMA security of the walnutDSA scheme

```

## Compute the dimensionality of the WalnutDSA Encoder Matrix

import random
import itertools

# Note; If you change N, you should change the definitions of G
# down near the end.
N = 12;
B = Brai dGroup(N, ' b' );
K = GF(2^8);
MS = Matri xSpace(K, N, N);

B. i nject_vari abl es();

S = Set();

def RandomNonzeroEl ement():
    while True:
        r = K. random_el ement()
        if r != 0:
            return r;

def myPerm(b):
    p = b. permutati on();
    return p*Permutati ons(4). i denti ty();

def i sPure(a):
    #pri nt(a, a. permutati on(). cycl e_stri ng());
    return a. permutati on(). cycl e_stri ng() == " ()";

def getAl l Brai ds(l):
    def i sReduced(a):
        for i in range(0, len(a)-1):
            if a[i] == -a[i+1]:
                return False;
        return True;

    def toBrai d(a):
        return B(a);

    def Al l Brai ds(l):
        S = [i for i in range(1, N)] + [-i for i in range(1, N)];
        return

    i tertool s. i map(toBrai d, i tertool s. i fi l ter(i sReduced, i tertool s. product(*([S]*l))));

    S = set();
    for i in range(0, l+1):
        S. update(Al l Brai ds(i))
    return S;

def getAl l PureBrai ds(l):
    l = l/2;
    D = di ct()
    for b in getAl l Brai ds(l):
        p = myPerm(b);
        if p in D:
            D[p]. add(b);
        el se:
            D[p] = set([b]);

    PB = set();
    for p in D:
        pi nv = p. i nverse();

```

General WalnutEncoder.sage

```

    for a in D[p]:
        for b in D[pi nv]:
            PB.add(a*b);
            PB.add(b*a);
    return PB;

def PBG(i , j):
    return B([a for a in range(j -1, i, -1) ] + [ i , i ] + [ -a for a in range(i+1, j)
]);

def getPureBraidGenerators():
    S = set();
    for i in range(1, N):
        for j in range(i+1, N):
            S.add(PBG(i , j))
    return S;

# Choosing T values
Tvals = [RandomNonzeroElement() for i in range(0, N)];
Tvals[0] = K.one()
Tvals[1] = K.one()

CBOne = (MS.identity_matrix(), Permutations(N).identity())

def Emul Step(x, g):
    M, p = x;
    M = copy(M);
    p = copy(p);

    if (g>0) :
        a = Tvals[p(g)-1];
        b = -a;
        c = 1;
    else :
        a = 1;
        b = - Tvals[p(-g+1)-1]^(-1);
        c = -b;

    g = abs(g);

    if(g>1):
        M.set_column(g-2 , M.column(g-2) + a* M.column(g-1));

    M.set_column(g , M.column(g) + c*M.column(g-1));
    M.set_column(g-1 , b*M.column(g-1));
    return (M, B([g]).permutation()*p);

def Emul (x, b):
    return reduce(Emul Step, [x]+list(b.Tietze()));

#The number of generators per hash value
Length = 256; #512/2
#The number of hash values (this must be greater than the expected dimension)
Samples = 120;

#The generators of the free group
# For B12: 579B, 468A, 3579, 2468, 1357
G = [ [PBG(5, N), PBG(7, N), PBG(9, N), PBG(11, N)],
      [PBG(4, N), PBG(6, N), PBG(8, N), PBG(10, N)],
      [PBG(3, N), PBG(5, N), PBG(7, N), PBG(9, N)],
      [PBG(2, N), PBG(4, N), PBG(6, N), PBG(8, N)],
      [PBG(1, N), PBG(3, N), PBG(5, N), PBG(7, N)],
      [PBG(2, N), PBG(4, N), PBG(6, N), PBG(8, N)],

```



```

                                General WalnutEncoder.sage
[PBG(3, N), PBG(5, N), PBG(7, N), PBG(9, N)],
[PBG(4, N), PBG(6, N), PBG(8, N), PBG(10, N)] ]

Bi gMatri x = matri x(K, N*N, Sampl es);

#for i in range(0, 10):
#    print("Choosi ng: " + str(random. choi ce(G[i %l en(G)])))

for i in range(0, Sampl es):
    x = CBOne;
    for j in range(0, Length):
        x = Emul (x, random. choi ce(G[j %l en(G)]));
    m, p = x
    Bi gMatri x. set_col umn(i, m. l i st())
    print("progress = "+str(100. 0*(i)/Sampl es)+"%")

print("Di mensi on of span of P(E(h_i)) is :")
print(Bi gMatri x. rank())

```

From: Derek Atkins <datkins@securerf.com>
Sent: Wednesday, April 04, 2018 9:06 AM
To: pqc-forum@list.nist.gov
Subject: [pqc-forum] OFFICIAL COMMENT: WalnutDSA -- new exponential attack

All,

Researchers Blackburn and Beullens have notified us about a new exponential attack against WalnutDSA. Although they have not provided a final version of their paper that they are now posting, we believe it will show that the resulting attack, prior to our initial response, performs in $q^{(N - 5/2)}$ time, which is what they said in a draft paper and we verified with the code they have provided. They further claimed they can reduce this running time to $q^{(N/2 - 1)}$, but this assumes $t_1=t_2=1$. We have not been able to verify this theoretical run time.

WalnutDSA has multiple parameters that can be adjusted to reach a desired security level and each parameter affects performance in a different way. We have only studied this attack for a very brief period of time, but assuming their worst-case analysis of $q^{(N/2 - 1)}$ is correct, we would propose that the N and q parameters be increased to N=11, q=M31 for 128-bit security and N=11, q=M61 for 256-bit security (where Mx is the Mersenne Prime $2^x - 1$). These parameter-only changes block their attack (this was confirmed by them).

An additional benefit of these new parameters, specifically changing to a prime field, is that we can easily create cloaking elements that no longer require $t_1=t_2=1$. By taking generators to the fourth power (instead of the 2nd) the t-value t_1 can be randomly chosen while $t_2 = -t_1^{-1}$. This further complicates this attack and increases its run time by a $\sqrt{60*q}$ factor, which then runs in $\sqrt{60}*q^{((N-1)/2)}$ time (confirmed by them). This allows us to reduce N to 10 and still yield a 2^{142} and 2^{277} security level.

Using B10M31, SUPERCOP reports that signature generation is 161 million cycles and signature validation is 175285 cycles, without leveraging AVX. We expect some improvements including when we integrate AVX, and a similar increase for the 256-bit parameters.

Details about this attack are forthcoming from the researchers, but these updated parameters effectively address it.

Thanks,

The WalnutDSA Team
--
Derek Atkins
Chief Technology Officer
SecureRF Corporation

Office: 203.227.3151 x1343
Direct: 617.623.3745
Mobile: 617.290.5355
Email: DAtkins@SecureRF.com

This email message may contain confidential, proprietary and / or legally privileged information and intended only for the use of the intended recipient(s) and others specifically authorized. Any disclosure, dissemination, copying, distribution or use of the information contained in this email message, including any attachments, to or by anyone other than the intended recipient is strictly prohibited. If you received this in error, please immediately advise the sender by reply email or at the telephone number above, and then delete, shred, or otherwise dispose of this message.

--
You received this message because you are subscribed to the Google Groups "pqc-forum" group.

From: Ward Beullens <ward.beullens@student.kuleuven.be>
Sent: Wednesday, April 04, 2018 11:57 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: OFFICIAL COMMENT: WalnutDSA

Dear All,

As was announced by Derek, we have found another practical attack on the Walnut signature scheme. The attack is described in detail the paper "Practical attacks against the Walnut digital signature scheme", together with the attacks that were previously announced on the forum.

The paper is now available at ePrint: <https://eprint.iacr.org/2018/318>

The code for the various attacks is available here : <https://github.com/WardBeullens/Nutcracker>

From the conclusion of our paper: "The security of the parameter sets submitted to the NIST PQC project is completely broken by the attacks.

We show that it is possible to forge signatures or compute equivalent secret keys in under a second for 128-bit security parameters. Even for 256-bit security parameters this takes less than a minute. Updating the parameters to resist the best known attack (see Sect. 5) would significantly increase the public key and signature sizes. This would make the scheme more difficult to implement on the low-resource processors that SecureRF is targeting and destroy the size advantages of Walnut over other post-quantum signature schemes such as lattice-based, hash-based and multivariate signature schemes. We note also that these latter schemes have been subject to much more scrutiny, which improves our confidence in their security."

We want to address some misconceptions that might arise after reading Derek's message to the forum.

1. In case of doubt, we should make it clear that we shared the full code for our attacks, and earlier drafts of our paper, with SecureRF. We sent the final version of the code (modulo cleaning up the code a bit) to SecureRF on 24th March, and answered any questions they asked on their experiments with it. The posted version of our paper only differs from the last draft we shared in the suggested choice of parameters at the end of Section 5: we informed SecureRF of this change before posting.
2. "They further claimed they can reduce this running time to $q^{(N/2 - 1)}$, but this assumes $t_1=t_2=1$." The Walnut Specification document mandates that the first two t-values, t_1 and t_2 , are fixed to 1. There is no "assumption" here.
3. "We have not been able to verify this theoretical run time." The run time is not theoretical. The code available on GitHub (and shared with SecureRF) runs with the run time we would expect from our analysis.
4. "These parameter-only changes block their attack (this was confirmed by them)." We have not confirmed that these parameters block our attack. However, we believe this to be the case, and we have proposed similar sized parameters in our paper.

Kind regards,
Ward and Simon

From: Derek Atkins <datkins@securerf.com>
Sent: Sunday, April 08, 2018 4:34 PM
To: ward.beullens@esat.kuleuven.be; pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: WalnutDSA

Dear Ward, All,

I'd like to apologize for the delay in responding; I've been on vacation this week and generally away from my email.

On Wed, 2018-04-04 at 17:57 +0200, Ward Beullens wrote:

From the conclusion of our paper: "The security of the parameter sets submitted to the NIST PQC project is completely broken by the attacks. We show that it is possible to forge signatures or compute equivalent secret keys in under a second for 128-bit security parameters. Even for 256-bit security parameters this takes less than a minute. Updating the parameters to resist the best known attack (see Sect. 5) would significantly increase the public key and signature sizes. This would make the scheme more difficult to implement on the low-resource processors that SecureRF is targeting and destroy the size advantages of Walnut over other post-quantum signature schemes such as lattice-based, hash-based and multivariate signature schemes. We note also that these latter schemes have been subject to much more scrutiny, which improves our confidence in their security."

In the passage above, and elsewhere in this paper, the authors move away from presenting their attack on Walnut and rather focus it on the company supporting this work. We strongly believe that this subjective editorializing does not belong in this paper or the NIST process. In many places in this paper, this style makes it difficult to discern their actual analysis of Walnut. Our understanding is that NIST is looking to identify candidates for a Post-Quantum era and it has not included minimum or maximum implementation guidelines. Although we do not yet agree with the authors on what the final public key and signature sizes will be, we believe that for now, the reader should focus on the fact that we do agree with the authors that "Updating the parameters to resist the best known attack..." defeats their analysis.

In regards to suitability, we are still studying this paper, additional responses, and its impact on the performance of WalnutDSA. Using parameters that we have already shared with the authors, and acknowledged in return emails, the performance still delivers the fastest verification times of all the methods submitted to NIST, and the signature generation time is not greatly impacted. Furthermore, as a result of parameter-only changes, our very-small code size for implementation has not been affected.

We want to address some misconceptions that might arise after reading Derek's message to the forum.

1. In case of doubt, we should make it clear that we shared the full code for our attacks, and earlier drafts of our paper, with SecureRF. We sent the final version of the code (modulo cleaning up the code a bit) to SecureRF on 24th March, and answered any questions they asked on their experiments with it. The posted version of our paper only differs from the last draft we shared in the suggested choice of parameters at the end of Section 5: we informed SecureRF of this change before posting.

As we stated in our post and confirmed by Ward, we did not receive a final draft of the paper. The last version we received was March 23, 12 days before the final paper was archived, and we pointed this out because many items and issues were provided to the authors which either needed further review, were not published, or were already resolved. For example, we had already addressed the encoder issue repeated in this paper in our January 23, 2018 response, where we suggested using a 2-bit encoder, and subsequently on February 2, 2018, where we suggested a rotating encoder that utilizes the full space. Earlier, on January 22, 2018, we suggested increasing N to 10. Each of these suggestions were subsequently confirmed by the authors, but omitted from the paper.

2. "They further claimed they can reduce this running time to $q^{(N/2 - 1)}$, but this assumes $t_1=t_2=1$." The Walnut Specification document mandates that the first two t -values, t_1 and t_2 , are fixed to 1. There is no "assumption" here.

It would appear this statement is being used to mix two issues – and we are partly responsible. Our submission to NIST assumes that $t_1=t_2=1$ and, in order to comply with NIST's process, there are to be no changes in responding to attacks during this initial period. However, outside of this process, we can make any number of suggestions based on input received. Almost immediately after the authors sent their first outline of the paper, we asked them about the case when $t_1, t_2 \neq 1$. Separate from this response, and already established, we can block the attack for the case of $t_1=t_2=1$ by choosing $N=11, q=M31$ and $N=11, q=M61$ which are parameter-only changes. So their comment is correct, and we should have separated our comment on your claimed reduced runtime from this assumption and the likely impact of $t_1, t_2 \neq 1$, as you have done now by separating your #2 from #3 comments here.

3. "We have not been able to verify this theoretical run time." The run time is not theoretical. The code available on GitHub (and shared with SecureRF) runs with the run time we would expect from our analysis.

The latest code you have provided to us and on GitHub has not yet supported your claimed run time of $q^{(N/2-1)}$, so it is still theoretical. In running the latest version of the code, it failed to build enough data after nearly a week of running to verify its growth. Your paper suggests that the attack can get stuck, and indeed even running with small parameters like $N=3$ or $N=4$ and $q=7$, the attack was getting stuck more often than not, sometimes taking ~ 12 hours to complete a single run with a single set of parameters. We did point this out to you before you published your paper.

4. "These parameter-only changes block their attack (this was confirmed by them)." We have not confirmed that these parameters block our attack. However, we believe this to be the case, and we have proposed similar sized parameters in our paper.

This statement seems to contradict your previous statement. If the runtime is not theoretical then simply plugging in the proposed parameters should be sufficient to confirm they block your attack. In regards to our comment on your confirming these parameters, our initial discussion was for the case of $t_1, t_2 \neq 1$, and you responded by proposing $N=10, q=2^{32}$, and $N=10, q=2^{64}$ (for $t_1=t_2=1$). From your April 2 email, one can easily infer it also applies to the case when $t_1=t_2=1$ provided $N=11, q=M31$ and $N=11, q=M61$. Most importantly, we confirm that parameter-only changes block your attack and we appreciate your acknowledging this to be the case.

Thanks,

The WalnutDSA Team

--

Derek Atkins
Chief Technology Officer
SecureRF Corporation

From: Jacob Alperin-Sheriff <jacobmas@gmail.com>
Sent: Sunday, April 08, 2018 5:29 PM
To: Derek Atkins
Cc: ward.beullens@esat.kuleuven.be; pqc-comments; pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: WalnutDSA

Given the lack of dispute that Beullens algorithm breaks 128 in < second, 256 in less than a minute, I find it very very hard to believe that a version of Walnut that resists the Beullens algorithm solely via parameter changes would have faster verification times than the fastest lattice and multivariate schemes, given that the completely broken 256 bit parameters already don't seem any faster than the fastest (or at most just barely so).

On Sun, Apr 8, 2018, 4:33 PM Derek Atkins <datkins@securerf.com> wrote:

Dear Ward, All,

I'd like to apologize for the delay in responding; I've been on vacation this week and generally away from my email.

On Wed, 2018-04-04 at 17:57 +0200, Ward Beullens wrote:

From the conclusion of our paper: "The security of the parameter sets submitted to the NIST PQC project is completely broken by the attacks. We show that it is possible to forge signatures or compute equivalent secret keys in under a second for 128-bit security parameters. Even for 256-bit security parameters this takes less than a minute. Updating the parameters to resist the best known attack (see Sect. 5) would significantly increase the public key and signature sizes. This would make the scheme more difficult to implement on the low-resource processors that SecureRF is targeting and destroy the size advantages of Walnut over other post-quantum signature schemes such as lattice-based, hash-based and multivariate signature schemes. We note also that these latter schemes have been subject to much more scrutiny, which improves our confidence in their security."

In the passage above, and elsewhere in this paper, the authors move away from presenting their attack on Walnut and rather focus it on the company supporting this work. We strongly believe that this subjective editorializing does not belong in this paper or the NIST process. In many places in this paper, this style makes it difficult to discern their actual analysis of Walnut. Our understanding is that NIST is looking to identify candidates for a Post-Quantum era and it has not included minimum or maximum implementation guidelines. Although we do not yet agree with the authors on what the final public key and signature sizes will be, we believe that for now, the reader should focus on the fact that we do agree with the authors that "Updating the parameters to resist the best known attack..." defeats their analysis.

In regards to suitability, we are still studying this paper, additional responses, and its impact on the performance of WalnutDSA. Using parameters that we have already shared with the authors, and acknowledged in return emails, the performance still delivers the fastest verification times of all the methods submitted to NIST, and the signature generation time is not greatly impacted. Furthermore, as a result of parameter-only changes, our very-small code size for implementation has not been affected.

We want to address some misconceptions that might arise after reading Derek's message to the forum.

1. In case of doubt, we should make it clear that we shared the full code for our attacks, and earlier drafts of our paper, with SecureRF. We sent the final version of the code (modulo cleaning up the code a bit)

From: Derek Atkins <datkins@securerf.com>
Sent: Sunday, April 08, 2018 7:10 PM
To: jacobmas@gmail.com
Cc: pqc-forum@list.nist.gov; ward.beullens@esat.kuleuven.be; pqc-comments
Subject: Re: [pqc-forum] OFFICIAL COMMENT: WalnutDSA

Dear Jacob,

Thank you for your email.

Our SUPERCOP runs clocked the verification using the updated 128-bit parameters at 175,285 cycles. If you are attending the conference, I would be happy to meet and share the SUPERCOP runs with you.

We will also be posting the code with the updated parameters later this week.

-derek

On Sun, 2018-04-08 at 21:28 +0000, Jacob Alperin-Sheriff wrote:

Given the lack of dispute that Beullens algorithm breaks 128 in < second, 256 in less than a minute, I find it very very hard to believe that a version of Walnut that resists the Beullens algorithm solely via parameter changes would have faster verification times than the fastest lattice and multivariate schemes, given that the completely broken 256 bit parameters already don't seem any faster than the fastest (or at most just barely so).

--

Derek Atkins
Chief Technology Officer
SecureRF Corporation

Office: 203.227.3151 x1343
Direct: 617.623.3745
Mobile: 617.290.5355
Email: DAtkins@SecureRF.com

This email message may contain confidential, proprietary and / or legally privileged information and intended only for the use of the intended recipient(s) and others specifically authorized. Any disclosure, dissemination, copying, distribution or use of the information contained in this email message, including any attachments, to or by anyone other than the intended recipient is strictly prohibited. If you received this in error, please immediately advise the sender by reply email or at the telephone number above, and then delete, shred, or otherwise dispose of this message.

From: Anton Menshov <menshov.a.v@gmail.com>
Sent: Friday, May 04, 2018 12:32 PM
To: pqc-comments
Subject: OFFICIAL COMMENT: WalnutDSA -- new practical attack

Dear All,

We recently came up with a new practical attack on WalnutDSA, which works with braids only and which success doesn't depend on the base finite field.

The attack is described here - <https://eprint.iacr.org/2018/393> .

Kind regards,
Anton, Sasha, and Matvei.

From: Derek Atkins <datkins@securerf.com>
Sent: Tuesday, May 15, 2018 7:04 AM
To: pqc-forum@list.nist.gov; menshov.a.v@gmail.com
Subject: Re: [pqc-forum] Fwd: OFFICIAL COMMENT: WalnutDSA -- new practical attack

Dear All,

We were made aware of this attack against WalnutDSA only when it was posted at the end of last week.

The paper does not contain any complexity analysis, and the code they supplied did not work initially (we have since created a working copy) all of which is making the evaluation and testing of our refutation take a bit longer than we would normally hope.

When we complete our analysis on the complexity of this attack (shortly), we will provide the details of how it is defeated.

Thanks,

The WalnutDSA team

--

Derek Atkins
Chief Technology Officer
SecureRF Corporation

Office: 203.227.3151 x1343
Direct: 617.623.3745
Mobile: 617.290.5355
Email: DAtkins@SecureRF.com

This email message may contain confidential, proprietary and / or legally privileged information and intended only for the use of the intended recipient(s) and others specifically authorized. Any disclosure, dissemination, copying, distribution or use of the information contained in this email message, including any attachments, to or by anyone other than the intended recipient is strictly prohibited. If you received this in error, please immediately advise the sender by reply email or at the telephone number above, and then delete, shred, or otherwise dispose of this message.

--

You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
Visit this group at <https://groups.google.com/a/list.nist.gov/group/pqc-forum/>.

From: Derek Atkins <datkins@securerf.com>
Sent: Wednesday, May 23, 2018 3:39 PM
To: pqc-forum@list.nist.gov; menshov.a.v@gmail.com
Subject: Re: [pqc-forum] Fwd: OFFICIAL COMMENT: WalnutDSA -- new practical attack

Dear All,

The WalnutDSA team has had a chance to analyze this attack and run their test code. We have found that this attack shows that our hard-coded parameter of 3 cloaking elements was too small, and that by increasing that number we can block this attack.

This attack requires two invariants in order to succeed:

- 1) The cloaking elements must be conjugates, and
- 2) It has to know the permutation of the cloaking elements

As we increase the number of cloaking elements in the WalnutDSA signature, we can also invalidate these two invariants. Specifically, when we add these additional cloaking elements we can do so at randomly chosen points in the signature, including inside the cloaking elements v , $v1$, and $v2$ and inside the private braids. When we insert these "concealed cloaking elements" in v , $v1$, and $v2$, then those elements are no longer conjugates themselves (breaking invariant #1). Further, these concealed cloaking elements have an unknown permutation, breaking invariant #2. To counteract this, they would require $N!$ searches to guess their permutations.

If we add K additional concealed cloaking elements in this manner, then it would require $(N!)^K$ additional work to run this attack. Therefore, if we choose K such that $(N!)^K > 2^{(2 * \text{SecurityLevel})}$ we know there is sufficient work for an attacker that we again reach our desired security goal.

An added bonus of making this parameter change is that we can now take L to 0, which winds up shrinking our signatures by over 30%, resulting in a significant size and performance boost!

Thanks,

The WalnutDSA Team

--
Derek Atkins
Chief Technology Officer
SecureRF Corporation

Office: 203.227.3151 x1343
Direct: 617.623.3745
Mobile: 617.290.5355
Email: DAtkins@SecureRF.com

This email message may contain confidential, proprietary and / or legally privileged information and intended only for the use of the intended recipient(s) and others specifically authorized. Any disclosure, dissemination, copying, distribution or use of the information contained in this email message, including any attachments, to or by anyone other than the intended recipient is strictly prohibited. If you received this in error, please immediately advise the sender by reply email or at the telephone number above, and then delete, shred, or otherwise dispose of this message.

--

From: Anton Menshov <menshov.a.v@gmail.com>
Sent: Monday, May 28, 2018 1:21 AM
To: pqc-comments; pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] Fwd: OFFICIAL COMMENT: WalnutDSA -- new practical attack

Dear All,

We would like to note that based on Remark 2.2 of our paper a reader may think that we need to know the permutation to locate critical letters. But in fact in our attack we only need to know $\sigma_w^{-1}(a)$ and $\sigma_w^{-1}(b)$ (which is known since $P(w)$, a , and b are public) to locate all critical letters inside a signature.

So the observation #2 together with the estimation $(N!)^{\text{Kappa}}$ is not correct.

Furthermore, even insertions of cloaking elements at randomly chosen points in a signature wouldn't prevent an attacker to locate all critical letters.

Taking $L=0$ is a good idea, long conjugators make cloaking elements visible.

Anton.

On Wed, May 23, 2018 at 3:39 PM Derek Atkins <datkins@securerf.com> wrote:

Dear All,

The WalnutDSA team has had a chance to analyze this attack and run their test code. We have found that this attack shows that our hard-coded parameter of 3 cloaking elements was too small, and that by increasing that number we can block this attack.

This attack requires two invariants in order to succeed:

- 1) The cloaking elements must be conjugates, and
- 2) It has to know the permutation of the cloaking elements

As we increase the number of cloaking elements in the WalnutDSA signature, we can also invalidate these two invariants. Specifically, when we add these additional cloaking elements we can do so at randomly chosen points in the signature, including inside the cloaking elements v , $v1$, and $v2$ and inside the private braids. When we insert these "concealed cloaking elements" in v , $v1$, and $v2$, then those elements are no longer conjugates themselves (breaking invariant #1). Further, these concealed cloaking elements have an unknown permutation, breaking invariant #2. To counteract this, they would require $N!$ searches to guess their permutations.

If we add Kappa additional concealed cloaking elements in this manner, then it would require $(N!)^{\text{Kappa}}$ additional work to run this attack. Therefore, if we choose Kappa such that $(N!)^{\text{Kappa}} > 2^{(2 * \text{SecurityLevel})}$ we know there is sufficient work for an attacker that we again reach our desired security goal.

An added bonus of making this parameter change is that we can now take L to 0, which winds up shrinking our signatures by over 30%, resulting in a significant size and performance boost!

Thanks,

The WalnutDSA Team

--

From: Simon-Philipp Merz <simon-philipp.merz@new.ox.ac.uk>
Sent: Monday, May 28, 2018 6:36 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: OFFICIAL COMMENT: WalnutDSA -- quantum collision search

Dear All,

following the attack of Simon Blackburn and Ward Beullens, we want to point out that the tight bound of quantum query complexity for finding 2-collisions of random functions has been revealed to be $O(M^{\{1/3\}})$, where M is the size of a codomain [1]. Consequently, we have a speedup compared to classical collision finding algorithms.

Applying this result to the attack of Simon Blackburn and Ward Beullens, the newest parameters of WalnutDSA do not attain the bounds claimed.

As we do not necessarily have T -values equal to 1 anymore we can't permute them towards the end. Thus we may assume that the T -values are chosen randomly. Then, the complexity of the attack to reverse the E-Multiplication using the finer subgroups method is dominated by a collision search in the quotient of size $q^{\{N-1\}}$ in the first step. Using the notation from Simons and Wards paper, we need to find k suitable collisions to get the braids c_1, \dots, c_k in order to complete the attack. In the implementation by Ward this parameter k has been chosen to equal 60.

Consequently, the query complexity of the first step using quantum collision search would be $\sqrt{k} \cdot q^{\{(N-1)/3\}}$.

For the 128-bit parameters $N=10$, $q=M31$ and $k=60$ this equals roughly $2^{\{96\}}$ instead of $2^{\{142\}}$ and for the 256-bit parameters $N=10$, $q=M61$ and $k=60$ this yields roughly $2^{\{186\}}$ instead of $2^{\{277\}}$.

Kind regards,
Simon-Philipp Merz

[1] Brassard, Gilles, Peter Høyer, and Alain Tapp. "Quantum cryptanalysis of hash and claw-free functions." Latin American Symposium on Theoretical Informatics. Springer, Berlin, Heidelberg, 1998.

From: Derek Atkins <datkins@securerf.com>
Sent: Tuesday, May 29, 2018 8:38 PM
To: simon-philipp.merz@new.ox.ac.uk; pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] OFFICIAL COMMENT: WalnutDSA

Dear All,

On Tue, 2018-05-29 at 13:12 +0000, Simon-Philipp Merz wrote:

Consequently, the query complexity of the first step using quantum collision search would be $\sqrt{k} \cdot q^{\{(N-1)/3\}}$.

For the suggested 128-bit parameters $N=10$, $q=M31$ and $k=60$ this equals roughly $2^{\{96\}}$ instead of $2^{\{142\}}$ and for the 256-bit parameters $N=10$, $q=M61$ and $k=60$ this yields roughly $2^{\{186\}}$ instead of $2^{\{277\}}$.

Thank you for this analysis using a quantum computer for collisions. I would argue that this attack, while interesting, is still less effective than Grover. So given a quantum computer, we believe it would still be more effective to use Grover than this attack.

Considering in Category 1 only the classical security needs to reach 2^{128} (quantum security is technically only 2^{64}), even under this attack it reaffirms the category 1 and category 5 levels for the proposed parameters.

Thank you,

The WalnutDSA Team

--

Derek Atkins
Chief Technology Officer
SecureRF Corporation

Office: 203.227.3151 x1343
Direct: 617.623.3745
Mobile: 617.290.5355
Email: DAtkins@SecureRF.com

This email message may contain confidential, proprietary and / or legally privileged information and intended only for the use of the intended recipient(s) and others specifically authorized. Any disclosure, dissemination, copying, distribution or use of the information contained in this email message, including any attachments, to or by anyone other than the intended recipient is strictly prohibited. If you received this in error, please immediately advise the sender by reply email or at the telephone number above, and then delete, shred, or otherwise dispose of this message.

From: Derek Atkins <datkins@securerf.com>
Sent: Tuesday, May 29, 2018 8:44 PM
To: pqc-forum@list.nist.gov; menshov.a.v@gmail.com; pqc-comments
Subject: Re: [pqc-forum] Fwd: OFFICIAL COMMENT: WalnutDSA -- new practical attack

Dear Anton, All,

On Mon, 2018-05-28 at 01:20 -0400, Anton Menshov wrote:

We would like to note that based on Remark 2.2 of our paper a reader may think that we need to know the permutation to locate critical letters. But in fact in our attack we only need to know $\sigma_w^{-1}(a)$ and $\sigma_w^{-1}(b)$ (which is known since $P(w)$, a , and b are public) to locate all critical letters inside a signature.

You seem to assume that these additional cloaking elements would be inserted consecutively within the signature. That is an incorrect assumption.

While it is true that a and b are public, and while it is true that $P(w)$ is public for v_1 , v_2 , and v_3 , $P(w)$ is NOT public for the additional concealed cloaking elements. Also, you will not know $\sigma_w^{-1}(a)$ or $\sigma_w^{-1}(b)$. Moreover, by inserting these concealed cloaking elements, we also ensure that v_1 , v_2 , and v_3 are not conjugates, which your attack requires to succeed.

Originally the cloaking element is created by generating w with a known σ_0 and mapping that to a and b ($P(w)$); then we turn that into a conjugate: $w b_i^4 w^{-1}$. However when we add the concealed cloaking elements that is no longer the case. We generate w as before, but then we break it into subwords w_l and w_r . Next, we create a concealed cloaking element for permutation $\sigma_0 . \sigma_w$ -- let's call it CE -- and then we insert it into the left side: w_l CE w_r . We do the same thing on the right side, split (the original) w^{-1} into subwords w_{2_l} and w_{2_r} and generate a cloaking element for the (unknown) permutation $\sigma_0 . \sigma_w . \sigma_{w_{2_1}} \Rightarrow$ CE2. Then we generate the full cloaking element w_l CE w_r b_i^4 w_{2_l} CE2 w_{2_r} .

You will note several things:

- 1) this is no longer a conjugate
- 2) While you know $P(w)$, you do not know $P(CE)$ or $P(CE2)$
- 3) There are now 3 sets of fourth-power entries (the main one, another in CE, and a third in CE2), which could all be at different b_i
- 4) For your attack to work as you propose, you would need to remove CE and CE2 first, which requires finding (or guessing) $P(CE)$ and $P(CE2)$ -- for which there are $N!$ choices each.
- 5) We have implemented this change in your code, and it indeed blocks your attack.

If you have further questions we suggest you first contact us offline to discuss, and if you feel anything material is identified, you can always post it to the forum

The WalnutDSA Team

--
Derek Atkins
Chief Technology Officer
SecureRF Corporation

From: Simon-Philipp Merz <simon-philipp.merz@gmx.de>
Sent: Friday, November 30, 2018 9:15 AM
To: pqc-comments
Cc: pqc-forum@list.nist.gov
Subject: OFFICIAL COMMENT: WalnutDSA

Dear WalnutDSA team, All,

we have found an attack on the most recent instantiation of the Walnut digital signature algorithm which attacks its "obfuscation procedure". The WalnutDSA team acknowledged our attack of which we informed them in August. The paper explaining the attack is now available at ePrint: <https://eprint.iacr.org/2018/1142>

In braid group cryptography public braids often contain secret braids as a factor and it is hoped that rewriting the product of braid words hides its individual factors. We provide experimental evidence that this is in general not the case and make an observation that can be exploited to decompose products in braid groups of the form ABC when only B is known.

This decomposition allows a universal forgery attack on Walnut signatures for the 128-bit and 256-bit parameters within seconds, given a single message-signature pair. The attack worked on 99.8% of 128-bit and 100% of 256-bit signatures in our experiments.

Moreover, we propose a potential countermeasure for our attack that allows to thwart it at the cost of slower signature generation and increased signature size.

Kind regards,
Christophe and Simon