

Loop-Abort Faults on Lattice-Based Fiat-Shamir and Hash-and-Sign Signatures

Thomas Espitau¹, Pierre-Alain Fouque², Benoît Gérard³,
and Mehdi Tibouchi⁴(✉)

¹ École normale supérieure de Cachan & Sorbonne Universités,
UPMC Univ Paris 06, LIP6, Paris, France
tespitau@ens-cachan.fr

² Institut Universitaire de France & IRISA & Université de Rennes I, Rennes, France
pierre-alain.fouque@univ-rennes1.fr

³ DGA.MI and IRISA, Rennes, France
benoit.gerard@irisa.fr

⁴ NTT Secure Platform Laboratories, Tokyo, Japan
tibouchi.mehdi@lab.ntt.co.jp

Abstract. Although postquantum cryptography is of growing practical concern, not many works have been devoted to implementation security issues related to postquantum schemes.

In this paper, we look in particular at fault attacks against implementations of lattice-based signature schemes, looking both at Fiat-Shamir type constructions (particularly BLISS, but also GLP, PASSSing and Ring-TESLA) and at hash-and-sign schemes (particularly the GPV-based scheme of Ducas–Prest–Lyubashevsky). These schemes include essentially all practical lattice-based signatures, and achieve the best efficiency to date in both software and hardware. We present several fault attacks against those schemes yielding a full key recovery with only a few or even a single faulty signature, and discuss possible countermeasures to protect against these attacks.

Keywords: Fault attacks · Digital signatures · Postquantum cryptography · Lattices · BLISS · GPV

1 Introduction

Lattice-based cryptography. Recent progress in quantum computation [7], the NSA advisory memorandum recommending the transition away from Suite B and to postquantum cryptography [1], as well as the announcement of the NIST standardization process for postquantum cryptography [6] all suggest that research on postquantum schemes, which is already plentiful but mostly focused on theoretical constructions and asymptotic security, should increasingly take into account real world implementation issues.

Among all postquantum directions, lattice-based cryptography occupies a position of particular interest, as it relies on well-studied problems and comes

with uniquely strong security guarantees, such as worst-case to average-case reductions [35]. A number of works have also focused on improving the performance of lattice-based schemes, and actual implementation results suggest that properly optimized schemes may be competitive with, or even outperform, classical factoring- and discrete logarithm-based cryptography.

The literature on the underlying number-theoretic problems of lattice-based cryptography is extensive (even though concrete bit security is not nearly as well understood as for factoring and discrete logarithms; in addition, ring-based schemes have recently been subjected to new families of attacks that might eventually reduce their security, especially in the postquantum setting). On the other hand, there is currently a distinct lack of cryptanalytic results on the *physical* security of implementations of lattice-based schemes (or in fact, postquantum schemes in general! [39]). It is well-known that physical attacks, particularly against public-key schemes, are often simpler, easier to mount and more devastating than attacks targeting underlying hardness assumptions: it is often the case that a few bits of leakage or a few fault injections can reveal an entire secret key (the well-known attacks from [3, 5] are typical examples). We therefore deem it important to investigate how fault attacks may be leveraged to recover secret keys in the lattice-based setting, particularly against signature schemes as signatures are probably the most likely primitive to be deployed in a setting where fault attacks are relevant, and have also received the most attention in terms of efficient implementations both in hardware and software.

Practical implementations of lattice-based signatures. Efficient signature schemes are typically proved secure in the random oracle model, and can be roughly divided in two families: the hash-and-sign family (which includes schemes like FDH and PSS), as well as signatures based on identification schemes, using the Fiat-Shamir heuristic or a variant thereof. Efficient lattice-based signatures can also be divided along those lines, as observed for example in the survey of practical lattice-based digital signature schemes presented by O’Neill and Güneysu at the NIST workshop on postquantum cryptography [23, 24].

The Fiat-Shamir family is the most developed, with a number of schemes coming with concrete implementations in software, and occasionally in hardware as well. Most schemes in that family follow Lyubashevsky’s “Fiat-Shamir with aborts” paradigm [26], which uses rejection sampling to ensure that the underlying identification scheme achieves honest-verifier zero-knowledge. Among lattice-based schemes, the exemplar in that family is Lyubashevsky’s scheme from EUROCRYPT 2012 [27]. It is, however, of limited efficiency, and had to be optimized to yield practical implementations. This was first carried out by Güneysu et al., who described an optimized hardware implementation of it at CHES 2012 [20], and then to a larger extent by Ducas et al. in their scheme BLISS [9], which includes a number of theoretical improvements and is the top-performing lattice-based signature. It was also implemented in hardware by Pöppelmann et al. [36]. Other schemes in that family include Hoffstein et al.’s PASSSign [22], which

incorporates ideas from NTRU, and Akleyek et al.’s Ring-TESLA [2], which boasts a tight security reduction.

On the hash-and-sign side, there were a number of early proposals with heuristic security (and no actual security proofs), particularly GGH [18] and NTRUSign [21], but despite several attempts to patch them¹ they turned out to be insecure. A principled, provable approach to designing lattice-based hash-and-sign signatures was first described by Gentry et al. in [16], based on discrete Gaussian sampling over lattices. The resulting scheme, GPV, is rather inefficient, even when using faster techniques for lattice Gaussian sampling [30]. However, Ducas et al. [11] later showed how it could be optimized and instantiated over NTRU lattices to achieve a relatively efficient scheme with particularly short signature size. The DLP scheme is somewhat slower than BLISS in software, but still a good contender for practical lattice-based signatures, and seemingly the only one in the hash-and-sign family.

Our contributions. In this work, we initiate the study of fault attacks against lattice-based signature schemes, and obtain attacks against all the practical schemes mentioned above.

As noted previously, early lattice-based signature schemes with heuristic security have been broken using standard attacks [15, 17, 32] but recent constructions including [9, 11, 16, 26, 27] are provably secure, and cryptanalysis therefore requires a more powerful attack model. In this work we consider fault attacks.

We present two attacks, both using a similar type of faults which allows the attacker to cause a loop inside the signature generation algorithm to abort early. Successful loop-abort faults have been described many times in the literature, including against DSA [31] and pairing computations [34], and in our attacks they can be used to recover information about the private signing key. The underlying mathematical techniques used to actually recover the key, however, are quite different in the two attacks.

Our first attack applies to the schemes in the Fiat-Shamir family: we describe it against BLISS [9, 36], and show how it extends to GLP [20], PASSSign [22] and Ring-TESLA [2]. In that attack, we inject a fault in the loop that generates the random “commitment value” \mathbf{y} of the sigma protocol associated with the Fiat-Shamir signature scheme. That commitment value is a random polynomial generated coefficient by coefficient, and an early loop abort causes it to have abnormally low degree, so that the protocol is no longer zero-knowledge. In fact, this will usually leak enough information that *a single faulty signature is enough to recover the entire signing key*. More specifically, we show that the faulty signature can be used to construct a point that is very close to a vector in a suitable integer lattice of moderate dimension, and such that the difference is essentially (a subset of) the signing key, which can thus be recovered using lattice reduction.

¹ There is a provably secure scheme due to Melchor et al. [29] that claims to “seal the leak on NTRUSign”, but it actually turns the construction into a Fiat-Shamir type scheme, using rejection sampling à la Lyubashevsky.

Our second attack targets the GPV-based hash-and-sign signature scheme of Ducas et al. [11]. In that case, we consider early loop abort faults against the discrete Gaussian sampling in the secret trapdoor lattice used in signature generation. The early loop abort causes the signature to be a linear combination of the last few rows of the secret lattice. A few faulty signatures can then be used to recover the span of those rows, and using the special structure of the lattice, we can then use lattice reduction to find one of the rows up to sign, which is enough to completely reconstruct the secret key. In practice, if we can cause loop aborts after up to m iterations, we find that $m + 2$ *faulty signatures are enough for full key recovery* with high probability.

Both of our attacks are supported by extensive simulations in Sage [38], whose source code is provided in the full version of this paper [13].

We also take a close look at the concrete software and hardware implementations of the schemes above, and discuss the concrete feasibility of injecting the required loop-abort faults in practice. We find the attacks to be highly realistic. Finally, we discuss several possible countermeasures to protect against our attacks.

Related work. To the best of our knowledge, the first previous work on fault attacks against lattice-based signatures, and in particular the only one mentioned in the survey of Taha and Eisenbarth [39], is the fault analysis work of Kamal and Youssef on NTRUSign [25]. It is, however, of limited interest since NTRUSign is known to be broken [12, 32]; it also suffers from a very low probability of success.

Much more recently, a relevant preprint has also been made available online by Bindel et al. [4] concurrently with this work. That paper proposes various fault attacks against the same Fiat-Shamir type schemes that we consider in this paper. Most of the attacks, however, are either in a contrived model (targeting key generation), or require unrealistically many faults and are arguably straightforward (bypassing rejection sampling in signature generation or size/correctness checks in signature verification). One attack described in the paper can be seen as posing a serious threat, namely the one described in [4, Sect. IV-B], but it amounts to a weaker variant of our Fiat-Shamir attack, using simple linear algebra rather than lattice reduction. As a result, it requires several hundred faulty signatures, whereas our attack needs only one.

Another interesting concurrent work is the recent cache attack against BLISS of Bruinderink et al. [19]. It uses cache side-channels to extract information about the coefficients of the commitment polynomial \mathbf{y} , and then lattice reduction to recover the signing key based on that side-channel information. In that sense, it is similar to our Fiat-Shamir attack. However, since the nature of the information to be exploited is quite different than in our setting, the mathematical techniques are also quite different. In particular, again, in contrast with our fault attack, that cache attack requires many signatures for a successful key recovery.

2 Description of the Lattice-Based Signature Schemes We Consider

Notation. For any integer q , we represent the ring \mathbb{Z}_q by $[-q/2, q/2) \cap \mathbb{Z}$. Vectors are considered as column vectors and will be written in bold lower case letters and matrices with upper case letters. By default, we will use the ℓ_2 Euclidean norm, $\|\mathbf{v}\|_2 = (\sum_i v_i^2)^{1/2}$ and ℓ_∞ -norm as $\|\mathbf{v}\|_\infty = \max_i |v_i|$.

The Gaussian distribution with standard deviation $\sigma \in \mathbb{R}$ and center $c \in \mathbb{R}$ at $x \in \mathbb{R}$, is defined by $\rho_{c,\sigma}(x) = \exp(-\frac{(x-c)^2}{2\sigma^2})$ and more generally by $\rho_{\mathbf{c},\sigma}(\mathbf{x}) = \exp(-\frac{(\mathbf{x}-\mathbf{c})^2}{2\sigma^2})$ and when $\mathbf{c} = \mathbf{0}$, by $\rho_\sigma(\mathbf{x})$. The discrete Gaussian distribution over \mathbb{Z} centered at $\mathbf{0}$ is defined by $D_\sigma(x) = \rho_\sigma(x)/\rho_\sigma(\mathbb{Z})$ (or $D_{\mathbb{Z},\sigma}$) and more generally over \mathbb{Z}^m by $D_\sigma^m(\mathbf{x}) = \rho_\sigma(\mathbf{x})/\rho_\sigma(\mathbb{Z}^m)$, where $\rho_\sigma(\mathbb{Z}^m) = \sum_{\mathbf{x} \in \mathbb{Z}^m} \rho_\sigma(\mathbf{x})$.

Description of BLISS. The BLISS signature scheme [9] is possibly the most efficient lattice-based signature scheme so far. It has been implemented in both software [10] and hardware [36], and boasts performance numbers comparable to classical factoring and discrete-logarithm based schemes. BLISS can be seen as a ring-based optimization of the earlier lattice-based scheme of Lyubashevsky [27], sharing the same “Fiat-Shamir with aborts” structure [26]. One can give a simplified description of the scheme as follows: the public key is an NTRU-like ratio of the form $\mathbf{a}_q = \mathbf{s}_2/\mathbf{s}_1 \bmod q$, where the signing key polynomials $\mathbf{s}_1, \mathbf{s}_2 \in \mathcal{R} = \mathbb{Z}[\mathbf{x}]/(\mathbf{x}^n + 1)$ are small and sparse. To sign a message μ , one first generates commitment values $\mathbf{y}_1, \mathbf{y}_2 \in \mathcal{R}$ with normally distributed coefficients, and then computes a hash \mathbf{c} of the message μ together with $\mathbf{u} = -\mathbf{a}_q \mathbf{y}_1 + \mathbf{y}_2 \bmod q$. The signature is then the triple $(\mathbf{c}, \mathbf{z}_1, \mathbf{z}_2)$, with $\mathbf{z}_i = \mathbf{y}_i + \mathbf{s}_i \mathbf{c}$, and there is rejection sampling to ensure that the distribution of \mathbf{z}_i is independent of the secret key. Verification is possible because $\mathbf{u} = -\mathbf{a}_q \mathbf{z}_1 + \mathbf{z}_2 \bmod q$. The real BLISS scheme, described in full in Fig. 1, includes several optimizations on top of the above description. In particular, to improve the repetition rate, it targets a bimodal Gaussian distribution for the \mathbf{z}_i ’s, so there is a random sign flip in their definition. In addition, to reduce key size, the signature element \mathbf{z}_2 is actually transmitted in compressed form \mathbf{z}_2^\dagger , and accordingly the hash input includes only a compressed version of \mathbf{u} . These various optimizations are essentially irrelevant for our purposes.

Description of the GPV-based scheme of Ducas et al. The second signature scheme we consider is the one proposed by Ducas et al. at ASIACRYPT 2014 [11]. It is an optimization using NTRU lattices of the GPV hash-and-sign signature scheme of Gentry et al. [16], and has been implemented in software by Prest [37]. As in GPV, the signing key is a “good” basis of a certain lattice Λ (with short, almost orthogonal vectors), and the public key is a “bad” basis of the same lattice (with longer vectors and a large orthogonality defect). To sign a message μ , one simply hashes it to obtain a vector \mathbf{c} in the ambient space of Λ , and uses the good, secret basis to sample $\mathbf{v} \in \Lambda$ according to a discrete Gaussian

<pre> 1: function KEYGEN() 2: sample $\mathbf{f}, \mathbf{g} \in \mathcal{R} = \mathbb{Z}[\mathbf{x}]/(\mathbf{x}^n + 1)$, uniformly with $\lceil \delta_1 n \rceil$ coefficients in $\{\pm 1\}$, $\lceil \delta_2 n \rceil$ coefficients in $\{\pm 2\}$ and other equal to zero 3: $\mathbf{S} = (\mathbf{s}_1, \mathbf{s}_2)^T \leftarrow (\mathbf{f}, 2\mathbf{g} + 1)^T$ 4: if $N_\kappa(\mathbf{S}) \geq C^2 \cdot 5 \cdot (\lceil \delta_1 n \rceil + 4\lceil \delta_2 n \rceil) \cdot \kappa$ then restart 5: $\mathbf{a}_q = (2\mathbf{g} + 1)/\mathbf{f} \bmod q$ (restart if \mathbf{f} is not invertible) 6: return $(pk = \mathbf{a}_1, sk = \mathbf{S})$ where $\mathbf{a}_1 = 2\mathbf{a}_q \bmod 2q$ 7: end function 1: function VERIFY($\mu, pk = \mathbf{a}_1, (\mathbf{z}_1, \mathbf{z}_2^\dagger, \mathbf{c})$) 2: if $\ (\mathbf{z}_1, 2^d \cdot \mathbf{z}_2^\dagger)\ _2 > B_2$ then reject 3: if $\ (\mathbf{z}_1, 2^d \cdot \mathbf{z}_2^\dagger)\ _\infty > B_\infty$ then reject 4: accept iff $\mathbf{c} = H(\lceil \zeta \cdot \mathbf{a}_1 \cdot \mathbf{z}_1 + \zeta \cdot q \cdot \mathbf{c} \rceil_d + \mathbf{z}_2^\dagger \bmod p, \mu)$ 5: end function </pre>	<pre> 1: function SIGN($\mu, pk = \mathbf{a}_1, sk = \mathbf{S}$) 2: $\mathbf{y}_1 \leftarrow D_{\mathbf{z}, \sigma}^n, \mathbf{y}_2 \leftarrow D_{\mathbf{z}, \sigma}^n$ 3: $\mathbf{u} = \zeta \cdot \mathbf{a}_1 \cdot \mathbf{y}_1 + \mathbf{y}_2 \bmod 2q$ 4: $\mathbf{c} \leftarrow H(\lfloor \mathbf{u} \rfloor_d \bmod p, \mu)$ 5: choose a random bit b 6: $\mathbf{z}_1 \leftarrow \mathbf{y}_1 + (-1)^b \mathbf{s}_1 \mathbf{c}$ 7: $\mathbf{z}_2 \leftarrow \mathbf{y}_2 + (-1)^b \mathbf{s}_2 \mathbf{c}$ 8: rejection sampling: restart to step 2 except with probability $1/(M \exp(-\ \mathbf{Sc}\ /(2\sigma^2)) \cosh(\langle \mathbf{z}, \mathbf{Sc} \rangle / \sigma^2))$ 9: $\mathbf{z}_2^\dagger \leftarrow (\lfloor \mathbf{u} \rfloor_d - \lfloor \mathbf{u} - \mathbf{z}_2 \rfloor_d) \bmod p$ 10: return $(\mathbf{z}_1, \mathbf{z}_2^\dagger, \mathbf{c})$ 11: end function </pre>
---	---

Fig. 1. Description of the BLISS signature scheme. The random oracle H takes its values in the set of polynomials in \mathcal{R} with 0/1 coefficients and Hamming weight exactly κ , for some small constant κ . The value ζ is defined as $\zeta \cdot (q - 2) = 1 \bmod 2q$. The authors of [9] propose four different sets of parameters with security levels at least 128 bits. The interesting parameters for us are: $n = 512$, $q = 12289$, $\sigma \in \{215, 107, 250, 271\}$, $(\delta_1, \delta_2) \in \{(0.3, 0), (0.42, 0.03), (0.45, 0.06)\}$ and $\kappa \in \{23, 30, 39\}$. We refer to the original paper for other parameters and for the definition of notation like N_κ and $\lfloor \cdot \rfloor_d$, as they are not relevant for our attack. The instruction in red (sampling of \mathbf{y}_1) is where we introduce our faults. (Color figure online)

distribution of small variance supported on Λ and centered at \mathbf{c} . That vector \mathbf{v} is the signature; it is, in particular, a lattice point very close to \mathbf{c} . That property can be checked using the bad, public basis, but that basis is too large to sample such close vectors (this, combined with the fact that the discrete Gaussian leaks no information about the secret basis, is what makes it possible to prove security). The actual scheme of Ducas–Lyubashevsky–Prest, described in Fig. 2, uses a lattice of the same form as NTRU: $\Lambda = \{(\mathbf{y}, \mathbf{z}) \in \mathcal{R}^2 \mid \mathbf{y} + \mathbf{z} \cdot \mathbf{h} = 0\}$, where the public key \mathbf{h} is again a ratio $\mathbf{g}/\mathbf{f} \bmod q$ of small, sparse polynomials in $\mathcal{R} = \mathbb{Z}[\mathbf{x}]/(\mathbf{x}^n + 1)$. The use of such a lattice yields a very compact representation of the keys, and makes it possible to compress the signature as well by publishing only the second component of the sampled vector \mathbf{v} . As a result, this hash-and-sign scheme is very space efficient (even more than BLISS). However, the use of lattice Gaussian sampling makes signature generation significantly slower than BLISS at similar security levels.

<pre> 1: function KEYGEN(n, q) 2: $\mathbf{f} \leftarrow D_{\sigma_0}^n, \mathbf{g} \leftarrow D_{\sigma_0}^n$ \triangleright $\sigma_0 = 1.17\sqrt{q/2n}$ 3: if $\ (\mathbf{g}, -\mathbf{f})\ _2 > \sigma$ then restart \triangleright $\sigma = 1.17\sqrt{q}$ 4: if $\ (\frac{q\bar{\mathbf{f}}}{\mathbf{f}\mathbf{f}+\mathbf{g}\bar{\mathbf{g}}}, \frac{q\bar{\mathbf{g}}}{\mathbf{f}\mathbf{f}+\mathbf{g}\bar{\mathbf{g}}})\ _2 > \sigma$ then restart 5: using the extended Euclidean algorithm, compute $\rho_f, \rho_g \in \mathcal{R}$ and $R_f, R_g \in \mathbb{Z}$ s.t. $\rho_f \cdot \mathbf{f} = R_f \bmod \mathbf{x}^n + 1$ and $\rho_g \cdot \mathbf{g} = R_g \bmod \mathbf{x}^n + 1$ 6: if $\gcd(R_f, R_g) \neq 1$ or $\gcd(R_f, q) \neq 1$ then restart 7: using the extended Euclidean algorithm, compute $u, v \in \mathbb{Z}$ s.t. $u \cdot R_f + v \cdot R_g = 1$ 8: $\mathbf{F} \leftarrow qv\rho_g, \mathbf{G} \leftarrow -qu\rho_f$ 9: repeat 10: $\mathbf{k} \leftarrow \lfloor \frac{\mathbf{F}\bar{\mathbf{f}}+\mathbf{G}\bar{\mathbf{f}}}{\mathbf{f}\mathbf{f}+\mathbf{g}\bar{\mathbf{g}}} \rfloor \in \mathcal{R}$ 11: $\mathbf{F} \leftarrow \mathbf{F} - \mathbf{k} \cdot \mathbf{f}, \mathbf{G} \leftarrow \mathbf{G} - \mathbf{k} \cdot \mathbf{g}$ 12: until $\mathbf{k}=0$ 13: $\mathbf{h} \leftarrow \mathbf{g} \cdot \mathbf{f}^{-1} \bmod q$ 14: $\mathbf{B} \leftarrow \begin{pmatrix} \mathbf{M}_{\mathbf{g}} & -\mathbf{M}_{\mathbf{f}} \\ \mathbf{M}_{\mathbf{G}} & -\mathbf{M}_{\mathbf{F}} \end{pmatrix} \in \mathbb{Z}^{2n \times 2n}$ \triangleright short lattice basis 15: return $sk = \mathbf{B}, pk = \mathbf{h}$ 16: end function </pre>	<pre> 1: function GAUSSIANAMPLER($\mathbf{B}, \sigma, \mathbf{c}$) \triangleright we denote by \mathbf{b}_i (resp. $\tilde{\mathbf{b}}_i$) the rows of \mathbf{B} (resp. of its Gram-Schmidt matrix $\tilde{\mathbf{B}}$) 2: $\mathbf{v} \leftarrow \mathbf{0}$ 3: for $i = 2n$ down to 1 do 4: $\mathbf{c}' \leftarrow \langle \mathbf{c}, \tilde{\mathbf{b}}_i \rangle / \ \tilde{\mathbf{b}}_i\ _2^2$ 5: $\sigma' \leftarrow \sigma / \ \tilde{\mathbf{b}}_i\ _2$ 6: $r \leftarrow D_{\mathbb{Z}, \sigma', \mathbf{c}'}$ 7: $\mathbf{c} \leftarrow \mathbf{c} - r\mathbf{b}_i$ and $\mathbf{v} \leftarrow \mathbf{v} + r\mathbf{b}_i$ 8: end for 9: return \mathbf{v} \triangleright \mathbf{v} sampled according to the lattice Gaussian distribution $D_{\Lambda, \sigma, \mathbf{c}}$ 10: end function 1: function SIGN($\mu, sk = \mathbf{B}$) 2: $\mathbf{c} \leftarrow H(\mu) \in \mathbb{Z}_q^n$ 3: $(\mathbf{y}, \mathbf{z}) \leftarrow \text{GAUSSIANAMPLER}(\mathbf{B}, \sigma, (\mathbf{c}, \mathbf{0}))$ \triangleright \mathbf{y}, \mathbf{z} are short and satisfy $\mathbf{y} + \mathbf{z} \cdot \mathbf{h} = \mathbf{c} \bmod q$ 4: return \mathbf{z} 5: end function 1: function VERIFY($\mu, pk = \mathbf{h}, \mathbf{z}$) 2: accept iff $\ \mathbf{z}\ _2 + \ H(\mu) - \mathbf{z} \cdot \mathbf{h}\ _2 \leq \sigma\sqrt{2n}$ 3: end function </pre>
--	---

Fig. 2. Description of the GPV-based signature scheme of Ducas–Lyubashevsky–Prest. The random oracle H takes its values in \mathbb{Z}_q^n . We denote by $\mathbf{f} \mapsto \bar{\mathbf{f}}$ the conjugation involution of $\mathcal{R} = \mathbb{Z}[\mathbf{x}]/(\mathbf{x}^n + 1)$, i.e. for $\mathbf{f} = \sum_{i=0}^{n-1} f_i x^i$, $\bar{\mathbf{f}} = f_0 - \sum_{i=1}^{n-1} f_{n-i} x^i$. $\mathbf{M}_{\mathbf{a}}$ represents the matrix of the multiplication by \mathbf{a} in the polynomial basis of \mathcal{R} , which is anticirculant of dimension n . For 128 bits of security, the authors of [11] recommend the parameters $n = 256$ and $q \approx 2^{10}$. The constant 1.17 is an approximation of $\sqrt{e/2}$. The steps in red (main loop of the Gaussian sampler) is where we introduce our faults. (Color figure online)

3 Attack on Fiat-Shamir Type Lattice-Based Signatures

The first fault attack that we consider targets the lattice-based signature schemes of Fiat-Shamir type, and specifically the generation of the random “commitment” element in the underlying sigma protocols, which is denoted by \mathbf{y} in our descriptions. That element consists of one or several polynomials generated coefficient by coefficient, and the idea of the attack is to introduce a fault in that random sampling to obtain a polynomial of abnormally small degree, in which case signatures will leak information about the private signing key. For simplicity’s sake, we introduce the attack against BLISS in particular, but it works against the other Fiat-Shamir type schemes (GLP, PASSSign and Ring-TESLA) with almost no changes: see the full version of this paper [13] for details.

In BLISS, the commitment element actually consists of two polynomials $(\mathbf{y}_1, \mathbf{y}_2)$, and it suffices to attack \mathbf{y}_1 . Intuitively, \mathbf{y}_1 should mask the secret key element \mathbf{s}_1 in the relation $\mathbf{z}_1 = \pm \mathbf{s}_1 \mathbf{c} + \mathbf{y}_1$, and therefore modifying the distribution of \mathbf{y}_1 should cause some information about \mathbf{s} to leak in signatures. The actual picture in the Fiat-Shamir with aborts paradigm is in fact slightly different (namely, rejection sampling ensures that the distribution of \mathbf{z}_1 is independent of \mathbf{s}_1 , but only does so under the assumption that \mathbf{y}_1 follows the correct distribution), but the end result is the same: perturbing the generation of \mathbf{y}_1 should lead to secret key leakage.

Concretely speaking, in BLISS, $\mathbf{y}_1 \in \mathcal{R}_q$ is a ring element generated according to a discrete Gaussian distribution², and that generation is typically carried out coefficient by coefficient in the polynomial representation. Therefore, if we can use faults to cause an early termination of that generation process, we should obtain signatures in which the element \mathbf{y}_1 is actually a low-degree polynomial. If the degree is low enough, we will see that this reveals the whole secret key right away, from a single faulty signature!

Indeed, suppose that we can obtain a faulty signature obtained by forcing a termination of the loop for sampling \mathbf{y}_1 after the m -th iteration, with $m \ll n$. Then, the resulting polynomial \mathbf{y}_1 is of degree at most $m - 1$. As part of the faulty signature, we get the pair $(\mathbf{c}, \mathbf{z}_1)$ with $\mathbf{z}_1 = (-1)^b \mathbf{s}_1 \mathbf{c} + \mathbf{y}_1$. Without loss of generality, we may assume that $b = 0$ (we will recover the whole secret key only up to sign, but in BLISS, $(\mathbf{s}_1, \mathbf{s}_2)$ and $(-\mathbf{s}_1, -\mathbf{s}_2)$ are clearly equivalent secret keys). Moreover, with high probability, \mathbf{c} is invertible: if we heuristically assume that \mathbf{c} behaves like a random element of the ring from that standpoint, we expect it to be the case with probability about $(1 - 1/q)^n$, which is over 95% for all proposed BLISS parameters. We thus get an equation of the form:

$$\mathbf{c}^{-1} \mathbf{z}_1 - \mathbf{s}_1 \equiv \mathbf{c}^{-1} \mathbf{y}_1 \equiv \sum_{i=0}^{m-1} y_{1,i} \mathbf{c}^{-1} \mathbf{x}^i \pmod{q} \quad (1)$$

Thus, the vector $\mathbf{v} = \mathbf{c}^{-1} \mathbf{z}_1$ is very close to the sublattice of \mathbb{Z}^n generated by $\mathbf{w}_i = \mathbf{c}^{-1} \mathbf{x}^i \pmod{q}$ for $i = 0, \dots, m - 1$ and $q\mathbb{Z}^n$, and the difference should be \mathbf{s}_1 .

The previous lattice is of full rank in \mathbb{Z}^n , so the dimension is too large to apply lattice reduction directly. However, the relation given by Eq. (1) also holds for all subsets of indices. More precisely, let I be a subset of $\{0, \dots, n - 1\}$ of cardinality ℓ , and $\varphi_I: \mathbb{Z}^n \rightarrow \mathbb{Z}^I$ be the projection $(u_i)_{0 \leq i < n} \mapsto (u_i)_{i \in I}$. Then we also have that $\varphi_I(\mathbf{z}_1)$ is a close vector to the sublattice L_I of \mathbb{Z}^I generated by $q\mathbb{Z}^I$ and the images under φ_I of the \mathbf{w}_i 's; and the difference should be $\varphi_I(\mathbf{s}_1)$.

Equivalently, using Babai's nearest plane approach to the closest vector problem, we hope to show that $(\varphi_I(\mathbf{s}_1), B)$, for a suitably chosen positive constant B , is the shortest vector in the sublattice L'_I of $\mathbb{Z}^I \times \mathbb{Z}$ generated by $(\varphi_I(\mathbf{v}), B)$ as well as the vectors $(\varphi_I(\mathbf{w}_i), 0)$ and $q\mathbb{Z}^I \times \{0\}$.

² In the other Fiat-Shamir schemes such as [20], the distribution of each coefficient is uniform in some interval rather than Gaussian, but this doesn't affect our attack strategy at all.

The volume of L'_I is given by:

$$\text{vol}(L'_I) = B \cdot \text{vol}(L_I) = B \cdot \frac{\text{vol}(q\mathbb{Z}^I)}{[L_I : q\mathbb{Z}^I]} = Bq^{\ell-r}$$

where r is the rank of the family $(\varphi_I(\mathbf{w}_0), \dots, \varphi_I(\mathbf{w}_{m-1}))$ in \mathbb{Z}_q^I , which is at most m . Hence $\text{vol}(L'_I) \geq Bq^{\ell-m}$, and the Gaussian heuristic predicts that the shortest vector should be of norm:

$$\lambda_I \approx \sqrt{\frac{\ell+1}{2\pi e}} \cdot \text{vol}(L'_I)^{1/(\ell+1)} \gtrsim \sqrt{\frac{\ell+1}{2\pi e}} \cdot B^{1/(\ell+1)} q^{1-(m+1)/(\ell+1)}.$$

Thus, we expect that $(\varphi_I(\mathbf{s}_1), B)$ will actually be the shortest vector of L'_I provided that its norm is significantly smaller than this bound λ_I . Now $\varphi_I(\mathbf{s}_1)$ has roughly $\delta_1\ell$ entries equal to ± 1 , $\delta_2\ell$ entries equal to ± 2 and the rest are zeroes; therefore, the norm of $(\varphi_I(\mathbf{s}_1), B)$ is around $\sqrt{(\delta_1 + 4\delta_2)\ell + B^2}$. Let us choose $B = \lceil \sqrt{\delta_1 + 4\delta_2} \rceil$. The condition for \mathbf{s}_1 to be the shortest vector L_I can thus be written as:

$$\sqrt{(\delta_1 + 4\delta_2) \cdot (\ell + 1)} \ll \sqrt{\frac{\ell + 1}{2\pi e}} \cdot B^{1/(\ell+1)} q^{1-(m+1)/(\ell+1)}$$

or equivalently:

$$\ell + 1 \gtrsim \frac{m + 1 + \frac{\log \sqrt{\delta_1 + 4\delta_2}}{\log q}}{1 - \frac{\log \sqrt{2\pi e(\delta_1 + 4\delta_2)}}{\log q}}. \tag{2}$$

The denominator of the right-hand side of (2) ranges from about 0.91 for the BLISS-I and BLISS-II parameter sets down to about 0.87 for BLISS-IV. In all cases, we thus expect to recover $\varphi_I(\mathbf{s}_1)$ if we can solve the shortest vector problem in a lattice of dimension slightly larger than m . This is quite feasible with the LLL algorithm for m up to about 50, and with BKZ for m up to 100 or so.

To complete the attack, it suffices to apply the above to a family of subsets I of $\{0, \dots, n-1\}$ covering the whole set of indices, which reveals the entire vector \mathbf{s}_1 . The second component of the secret key is then obtained as $\mathbf{s}_2 = \mathbf{a}_1\mathbf{s}_1/2 \bmod q$.

Simulations using our Sage implementation (see the full version of this paper [13]) confirm the theoretical estimates, and show that full key recovery can be achieved in practice in a time ranging from a few seconds to a few hours depending on m . Detailed experimental results are reported in Table 1.

Remark 1. A variant of that attack which is possibly slightly simpler consists in observing that $\varphi_I(\mathbf{s}_1)$ should be the shortest vector in the lattice generated by L_I and $\varphi_I(\mathbf{v})$. The bound on the lattice dimension becomes essentially the same as (2). The drawback of that approach, however, is that we obtain each $\varphi_I(\mathbf{s}_1)$ up to sign, and so one needs to use overlapping subsets I to ensure the consistency of those signs.

Table 1. Experimental success rate of the attack and average CPU time for key recovery for several values of m , the iteration after which the loop-abort fault is injected. We attack the BLISS-II parameter set $(n, q, \sigma, \delta_1, \delta_2, \kappa) = (512, 12289, 10, 0.3, 0, 23)$ from [9]. Since the choice of ℓ has no effect on the concrete fault injection (e.g. it does not affect the required number of faulty signatures, which is always 1), we did not attempt to optimize it very closely. The simulation was carried out using our Sage implementation (see the full version of this paper [13]) on a single core of an Intel Xeon E5-2697v3 workstation, using 100 trial runs for each value of m .

Fault after iteration number $m =$	2	5	10	20	40	60	80	100
Theoretical minimum dimension ℓ_{\min}	3	6	11	22	44	66	88	110
Dimension ℓ in our experiment	3	6	12	24	50	80	110	150
Lattice reduction algorithm	LLL	LLL	LLL	LLL	BKZ-20	BKZ-25	BKZ-25	BKZ-25
Success probability (%)	100	99	100	100	100	100	100	98
Avg. CPU time to recover ℓ coeffs. (s)	0.002	0.005	0.022	0.23	7.3	119	941	33655
Avg. CPU time for full key recovery	0.5 s	0.5 s	1 s	5 s	80 s	14 min	80 min	38 h

Remark 2. Note that a single *faulty* signature is enough to recover the entire secret key with this attack, a successful key recovery may require several *fault injections*. This is due to rejection sampling: after a faulty \mathbf{y}_1 is generated, the whole signature may be thrown away in the rejection step. On average, the fault attacker may thus need to inject the same number of faults as the repetition rate of the scheme, which is a small constant ranging from 1.6 to 7.4 depending on chosen parameters [9], and even smaller with the improved analysis of BLISS-B [8].

Remark 3. Finally, we note that in certain hardware settings, fault injection may yield a faulty value of \mathbf{y}_1 in which all coefficients upwards of a certain degree bound are non zero but equal to a common constant (see the discussion in Sect. 5.3). Our attack adapts to that setting in a straightforward way: that simply means that \mathbf{y}_1 is a linear combination of the \mathbf{x}^i for small i and of the all-one vector $(1, \dots, 1)$, so it suffices to add that vector to the set of lattice generators.

4 Attack on Hash-and-Sign Type Lattice-Based Signatures

Our second attack targets the practical hash-and-sign signature scheme of Ducas et al. [11], which is based on GPV-style lattice trapdoors. More precisely, the faults we consider are again early loop aborts, this time in the lattice-point Gaussian sampling routine used in signature generation.

4.1 Description of the Attack

The attack can be described as follows. A correctly generated signature element is of the form $\mathbf{z} = \mathbf{R} \cdot \mathbf{f} + \mathbf{r} \cdot \mathbf{F} \in \mathbb{Z}[\mathbf{x}]/(\mathbf{x}^n + 1)$, where the short polynomials \mathbf{f} and \mathbf{F} are components of the secret key, and \mathbf{r}, \mathbf{R} are short random polynomials sampled in such a way that \mathbf{z} follows a suitable Gaussian distribution. In fact, \mathbf{r}, \mathbf{R} are generated coefficient by coefficient, in a single loop with $2n$ iterations, going from the top-degree coefficient of \mathbf{r} down to the constant coefficient of \mathbf{R} .

Therefore, if we inject a fault aborting the loop after $m \leq n$ iterations (in the first half of the loop), the resulting signature simply has the form:

$$\mathbf{z} = r_0 \mathbf{x}^{n-1} \mathbf{F} + r_1 \mathbf{x}^{n-2} \mathbf{F} + \dots + r_{m-1} \mathbf{x}^{n-m} \mathbf{F}.$$

Any such faulty signature is, in particular, in the lattice L of rank m generated by the vectors $\mathbf{x}^{n-i} \mathbf{F}$, $i = 1, \dots, m$, in $\mathbb{Z}[\mathbf{x}]/(\mathbf{x}^n + 1)$.

Suppose then that we obtain several signatures $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(\ell)}$ of the previous form. If ℓ is large enough (slightly more than m is sufficient; see Sect. 4.2 below for an analysis of success probability depending on ℓ), the corresponding vectors will then generate the lattice L . Assuming the lattice dimension is not too large, we should then be able to use lattice reduction to recover a shortest vector in L , which is expected to be one of the signed shifts $\pm \mathbf{x}^{n-i} \mathbf{F}$, $i = 1, \dots, m$, since the polynomial \mathbf{F} is constructed in a such a way as to make it quite short relative to the Gram–Schmidt norm of the ideal lattice it generates. Hence, we can recover \mathbf{F} among a small set of at most $2m$ candidates.

And recovering \mathbf{F} is actually sufficient to reconstruct the entire secret key $(\mathbf{f}, \mathbf{g}, \mathbf{F}, \mathbf{G})$, and hence completely break the scheme. This is due to the particular structure of the NTRU lattice. On the one hand, \mathbf{G} is linked to \mathbf{F} via the public key polynomial \mathbf{h} : $\mathbf{G} = \mathbf{F} \cdot \mathbf{h} \bmod q$, so we obtain it directly. On the other hand, the basis completion algorithm of Hoffstein et al. [21] allows to recover the pair (\mathbf{f}, \mathbf{g}) from (\mathbf{F}, \mathbf{G}) via the defining relation $\mathbf{f} \cdot \mathbf{G} - \mathbf{g} \cdot \mathbf{F} = q$. This is actually used in the opposite direction in the key generation algorithm of the scheme of Ducas et al. (i.e. they construct (\mathbf{F}, \mathbf{G}) from (\mathbf{f}, \mathbf{g}) : see steps 5–12 of KEYGEN in Fig. 2), but applying [21, Theorem 1], the technique is easily seen to work in both ways.

Moreover, if we start from a polynomial of the form $\zeta \mathbf{F}$ where ζ is of the form $\pm \mathbf{x}^\alpha$, then applying the previous steps yields the quadruple $(\zeta \mathbf{f}, \zeta \mathbf{g}, \zeta \mathbf{F}, \zeta \mathbf{G})$, which is also a valid secret key equivalent to $(\mathbf{f}, \mathbf{g}, \mathbf{F}, \mathbf{G})$, in the sense that signing with either keys produces signatures with exactly the same distributions. Thus,

we don't even need to carry out an exhaustive search on several possible values of \mathbf{F} after the lattice reduction step: it suffices to use the first vector of the reduced basis directly.

4.2 How Many Faults Do We Need?

Let us analyze the probability of success of the attack depending on the iteration m at which the iteration is inserted and the number $\ell > m$ of faulty signatures $\mathbf{z}^{(i)}$ available. As we have seen, a sufficient condition for the attack to succeed (provided that our lattice reduction algorithm actually finds a shortest vector) is that the ℓ faulty signatures generate the rank- m lattice L defined above. This is not actually necessary (the attack works as soon as *one* of the shifts of \mathbf{F} is in sub-lattice generated by the signatures, rather than all of them), but we will be content with a lower bound on the probability of success.

Now, that condition is equivalent to saying that the vectors $(r_0^{(i)}, \dots, r_{m-1}^{(i)}) \in \mathbb{Z}^m$ (sampled according to the distribution given by the GPV algorithm) that define the faulty signatures:

$$\mathbf{z}^{(i)} = r_0^{(i)} \mathbf{x}^{n-1} \mathbf{F} + \dots + r_{m-1}^{(i)} \mathbf{x}^{n-m} \mathbf{F}$$

generate the whole integer lattice \mathbb{Z}^m . But the probability that $\ell > m$ random vectors generate \mathbb{Z}^m has been computed by Maze et al. [28] (see also [14]), and

Table 2. Experimental success probability of the attack and average CPU time for key recovery for several values of m , the iteration after which the loop-abort fault is injected. We consider the attack with $\ell = m + 1$ and $\ell = m + 2$ faulty signatures. The attacked parameters are $(n, q) = (256, 1021)$ as suggested in [11] for signatures. The simulation was carried out using our Sage implementation (see the full version of this paper [13]) on a single core of an Intel Xeon E5-2697v3 workstation, using 100 trial runs for each pair (ℓ, m) .

Fault after iteration number $m =$	2	5	10	20	40	60	80	100
Lattice reduction algorithm	LLL	LLL	LLL	LLL	LLL	LLL	BKZ-20	BKZ-20
Success probability for $\ell = m + 1$ (%)	75	77	90	93	94	94	95	95
Avg. CPU time for $\ell = m + 1$ (s)	0.001	0.003	0.016	0.19	2.1	8.1	21.7	104
Success probability for $\ell = m + 2$ (%)	89	95	100	100	99	99	100	100
Avg. CPU time for $\ell = m + 2$ (s)	0.001	0.003	0.017	0.19	2.1	8.2	21.6	146

is asymptotically equal to $\prod_{k=\ell-m+1}^{\ell} \zeta(k)^{-1}$. In particular, if $\ell = m + d$ for some integer d , it is bounded below by:

$$p_d = \prod_{k=d+1}^{+\infty} \frac{1}{\zeta(k)}.$$

Thus, if we take $\ell = m + 1$ (resp. $\ell = m + 2$, $\ell = m + 3$), we expect the attack to succeed with probability at least $p_1 \approx 43\%$ (resp. $p_2 \approx 71\%$, $p_3 \approx 86\%$).

As shown in Table 2, this is well verified in practice (and the lower bound is in fact quite pessimistic). Moreover, the attack is quite fast even for relatively large values of m : only a couple of minutes for full key recovery for $m = 100$.

5 Implementation of the Faults

Once again, due to the obvious similarities between the four instances of the Fiat-Shamir family that we choose to attack, we only give details of the attack on the BLISS scheme. We also give details for the GPV scheme but they are essentially the same as the one for BLISS since the underlying fault introduced is strictly identical.

In this section we investigate how an attacker may obtain helpful faulty signatures for the proposed attacks. We base our discussion on two available implementations of BLISS signature, namely the software implementation from Ducas and Lepoint [10] and the FPGA implementation by Pöppelmann *et al.* [36], and on Prest's software implementation of the GPV-based scheme of Ducas *et al.* [37]. Notice that the discussion on the hardware implementation is also valid for the implementation of [20] since both share some common components and architecture that we exploit (for instance BRAM storage).

We emphasize the fact that those three implementations were not supposed to have any resilience with respect to fault attacks and were only developed as proofs of concept to illustrate the efficiency properties of the schemes. The point here is to show that the fault attacks presented in this paper are relevant based on the analysis of freely available and published implementations to put forward the need of dedicated protections against faults attacks (when attackers have such abilities).

5.1 Classical Fault Models

Faults during a computation may be induced by different means as a laser beam shot, electromagnetic injection, under-powering, glitches, etc. These faults are mainly characterized by their

- range: impacting a single bit or many bits (e.g. register or memory word);
- effect: typically target chunk is set to a chosen value, random value or all-zero/all-one value;
- persistence: a fault may only modify the target for a short period or it may be definitive.

Obviously, some fault models are close from being purely theoretical: it is very unlikely to be able to set a 32-bit register to `0xbad00dad` during precisely 2 cycles. Nevertheless many recent works have been published showing that some faults models that seemed overdone are actually obtained during lab experiments. One example is the work of Ordas et al. at CARDIS 2014 [33] showing that with finely tuned EM probes it is possible to induce a single-bit fault (bit-set or bit-reset).

In the next subsections we discuss which fault models³ may lead to faulty signatures relevant with respect to the attacks presented in this paper. We did not investigate clock glitches or under-powering which induce violation of the setup time and which actual side-effects are implementation and compilation-dependent (with large ranges of possible parameters to test). Nevertheless, they may not be overseen in the evaluation of a chip since they may also lead to the generation of relevant faulty signatures.

5.2 Fault Attacks on Software Implementations

Polynomial \mathbf{y}_1 can be generated using a loop over the n coefficients. This is, again, how the implementation in [10] is made: a loop is constructing polynomials \mathbf{y}_1 and \mathbf{y}_2 one coefficient at a time using a Gaussian sampler (function `Sign::signMessage`). The condition to perform the attack is rather few restrictive since we only require \mathbf{y}_1 to have at most (roughly) a quarter of unknown coefficients. Such result can be obtain by going out the loop after a few iterations. A random fault on the loop counter or skipping the jump operation will lead to such result.

Notice here that it is less trivial here to decide whether a faulty signature will be helpful or not. Hopefully, the timing precision is much less important here since the attack will succeed even with 50 unknown coefficients out of 512. This means that the time-window for the fault to occur is composed of decades of loop iterations. Moreover, we may use side-channel analysis to detect the loop iteration pattern to trigger the fault injection. Such pattern is likely to be detected after much less than 50 iterations and thus it seems that the synchronization here will be relatively easy.

Similarly, the short random polynomials \mathbf{R} and \mathbf{r} used in the GPV scheme are generated in a single loop [37] ranging from leading coefficient of \mathbf{r} to the constant term in \mathbf{R} which allows to fault both polynomials using a single fault. Again, a random fault on the counter or skipping a jump makes it work and the time-window large according to the results shown in Table 2.

To conclude, these attacks seems to be a real threat since synchronization (which is a major difficulty when performing fault attacks) is eased by the loose condition on the number of known coefficients in faulted polynomials.

5.3 Fault Attacks on Hardware Implementations

Generation of polynomial \mathbf{y}_1 requires n random coefficients. It is very unlikely that all these coefficients are obtained at the same time (n is too large) thus \mathbf{y}_1

³ We only focus on single fault attacks here.

generation will be sequential. This is the case in the implementation we took as example where the super memory is linked to the sampler through a 14-bit port. We may fault a flag or a state register to fool the control logic (here the bliss processor) and keep part of the BRAM cells to their initial state. If this initial state is known then we know all the corresponding coefficients and hopefully the number of unknown ones will be small enough for the attack to work. The large number of unknown coefficients handled by the attack again helps the attacker by providing a large time window for the fault to occur. The feasibility of the attack will mostly depend on the precise flag/state implementation and the knowledge of memory cells previous/initial value.

There is a second way of performing the fault injection here. The value of \mathbf{y}_1 has to be stored somehow until the computation of \mathbf{z}_1 (close to the end of the signature generation). In the example implementation a BRAM is used. We may fault BRAM access to fix some coefficients to a known value. A possible fault would be to set the `rstram` or `rstreg` signal to one (Xilinx’s nomenclature). Indeed, when set to one, this will set the output latches (*resp.* register) of the RAM block to some fixed value `SRVAL` defined by the designer. We may notice two points to understand why this kind of fault enables the proposed attack.

- (i) The value \mathbf{y}_1 used to compute \mathbf{u} will not be the faulted one but this has no impact on the attack.
- (ii) If we do not know the default value for the output register, all coefficients are unknown but a big part of them are equal to the same unknown default value. In that case, the attack is still applicable by adding one generator to the constructed lattice: see Remark 3 in Sect. 3.

Again a large time window is given to the attacker due to sequential read induced by the size of \mathbf{y}_1 .

The BRAM storage of \mathbf{y}_1 helps here the attacker since a single bit-set fault may have effects on many coefficients. The only difficulty seems to be able to perform a single-bit fault—which seems to be possible according to [33]—and the `rstram` signal localization⁴.

6 Conclusion and Possible Countermeasures

We have shown that unprotected implementations of the lattice-based signature schemes that we considered are vulnerable to fault attacks, in fault models that our analysis suggests are quite realistic: the faulty signatures required by our attacks can be obtained on actual implementations. As a result, countermeasures should be added in applications where such a physical attacker is relevant to the threat model.

⁴ Since \mathbf{y}_1 is not directly outputted checking if the attack actually worked is a bit more tricky. Again side-channel collision analysis may help here. We may also notice that if the faulty \mathbf{y}_1 is sparse (that is known coefficients have been set to zero) then the number of non-zero coefficients in the corresponding \mathbf{z}_1 should be significantly smaller than for a \mathbf{z}_1 corresponding to a dense \mathbf{y}_1 .

Simple countermeasures exist to thwart the single fault attacks proposed. There are simple, non-cryptographic countermeasures that consist in validating that the full loop have been correctly performed. This can be achieved for instance by adding a second loop counter and doing a consistency check after exiting the loop. Such a countermeasure is very cheap and we therefore recommend introducing it in all deployed implementations.

Nevertheless, it will only detect early-abort faults while an attacker may succeed in getting the same kind of faulty signature using another technique. For instance, we mentioned the possibility of faulting BRAM blocks so that they output a fixed value. For software implementations, the compiler may decide to put the coefficient in some RAM location which address could be faulted to point to another part of the memory leading in many coefficients having the same value. A single fault may also alter instruction cache leading to a `nop` operation instead of a load from memory and thus not updating the coefficient. We propose now other countermeasures that may deal with this issue for both types of signature schemes we considered.

We have described our attack on the Fiat-Shamir schemes in a setting where the attacker can obtain a commitment polynomial \mathbf{y} of low degree, and it works more generally with a sparse \mathbf{y} , provided that the attackers *knows* where the non zero coefficients are located. If the locations are unknown, however, the attack does not work, so one possible countermeasure is to randomize the order of the loop generating \mathbf{y} . One should be careful that this may not protect against faults introduced after the very first few iterations, however: in the case of BLISS, for example, we have seen that we could easily attack polynomials \mathbf{y} in which the non zero coefficients are located in the 20% lower degree coefficients, say; then, if a fault attacker can collect a few hundred faulty signatures with \mathbf{y} of very low Hamming weight (say 3 or 4) at random positions, they have a good chance of finding one fault with all non zero coefficients in the lower 20%, and hence be able to attack.

Another possible approach for the Fiat-Shamir schemes is to check that the degree of the generated \mathbf{y} is not too low. One cannot demand that all its coefficients are non zero, as this would skew the distribution and invalidate the security argument, but verifying that the top $\varepsilon \cdot n$ coefficients of \mathbf{y} are not all zero for some small constant $\varepsilon > 0$, say $\varepsilon = 1/16$, would be a practical countermeasure that does not affect the security proof. Indeed, in the case of BLISS for example, the probability that all of these coefficients vanish is roughly $(1/\sigma\sqrt{2\pi})^{\varepsilon n}$, which is exponentially small. Thus, the resulting distribution of \mathbf{y} after this check is statistically indistinguishable from the original distribution, and security is therefore preserved. Moreover, the lattice dimension required to mount our fault attack is then greater than $(1 - \varepsilon)n$, so it will not work. An additional advantage of that countermeasure is that it also adapts easily to thwart faults that cause all the top coefficients of \mathbf{y} to be equal to some constant non-zero value.

Regarding the hash-and-sign signature of Ducas et al., one possible countermeasure is to simply check the validity of generated signatures. This will

usually work due to the fact that a faulty signature generated from an early loop abort from the GAUSSIAN SAMPLER algorithm is of significantly larger norm than a valid signature: a rough estimate of the norm after $m \leq n$ iterations is $\|\mathbf{F}\|_2 \sqrt{mq/12}$ (as $q/12$ is the variance of a uniform random variable in $\{-(q-1)/2, \dots, (q-1)/2\}$), which is too large for correct verification even for very small values of m . An added benefit of that countermeasure is that even the correct signature generation algorithm has a very small but non zero probability of generating an invalid signature, so this countermeasure doubles up as a safeguard against those rare accidental failures.

Acknowledgments. We would like to thank Keita Xagawa and anonymous reviewers for useful comments on earlier versions of this paper.

References

1. Commercial national security algorithm suite and quantum computing FAQ. Technical report, National Security Agency, January 2016. <https://www.iad.gov/iad/library/ia-guidance/ia-solutions-for-classified/algorithm-guidance/cnsa-suite-and-quantum-computing-faq.cfm>
2. Akleyek, S., Bindel, N., Buchmann, J., Krämer, J., Marson, G.A.: An efficient lattice-based signature scheme with provably secure instantiation. In: Pointcheval, D., Nitaj, A., Rachidi, T. (eds.) AFRICACRYPT 2016. LNCS, vol. 9646, pp. 44–60. Springer, Cham (2016). doi:10.1007/978-3-319-31517-1_3
3. Biehl, I., Meyer, B., Müller, V.: Differential fault attacks on elliptic curve cryptosystems. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 131–146. Springer, Heidelberg (2000). doi:10.1007/3-540-44598-6_8
4. Bindel, N., Buchmann, J.A., Krämer, J.: Lattice-based signature schemes and their sensitivity to fault attacks. IACR Cryptology ePrint Archive 2016:415 (2016)
5. Boneh, D., DeMillo, R.A., Lipton, R.J.: On the importance of eliminating errors in cryptographic computations. *J. Cryptol.* **14**(2), 101–119 (2001)
6. Chen, L., Jordan, S., Liu, Y.-K., Moody, D., Peralta, R., Perlmutter, R., Smith-Tone, D.: Report on post-quantum cryptography. Technical report, National Institute of Standards and Technology, February 2016. <http://csrc.nist.gov/publications/drafts/nistir-8105/nistir.8105-draft.pdf>
7. Denchev, V.S., Boixo, S., Isakov, S.V., Ding, N., Babbush, R., Smelyanskiy, V., Martinis, J., Neven, H.: What is the Computational Value of Finite Range Tunneling? ArXiv e-prints, December 2015
8. Ducas, L.: Accelerating BLISS: the geometry of ternary polynomials. Cryptology ePrint Archive, Report 2014/874 (2014). <http://eprint.iacr.org/>
9. Ducas, L., Durmus, A., Lepoint, T., Lyubashevsky, V.: Lattice signatures and bimodal Gaussians. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013. LNCS, vol. 8042, pp. 40–56. Springer, Heidelberg (2013). doi:10.1007/978-3-642-40041-4_3
10. Ducas, L., Lepoint, T.: A proof-of-concept implementation of BLISS. Available under the CeCILL License at <http://bliss.di.ens.fr>
11. Ducas, L., Lyubashevsky, V., Prest, T.: Efficient identity-based encryption over NTRU lattices. In: Sarkar, P., Iwata, T. (eds.) ASIACRYPT 2014. LNCS, vol. 8874, pp. 22–41. Springer, Heidelberg (2014). doi:10.1007/978-3-662-45608-8_2

12. Ducas, L., Nguyen, P.Q.: Learning a zonotope and more: cryptanalysis of NTRUSign countermeasures. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 433–450. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-34961-4_27](https://doi.org/10.1007/978-3-642-34961-4_27)
13. Espitau, T., Fouque, P., Gérard, B., Tibouchi, M.: Loop-abort faults on lattice-based Fiat-Shamir and hash-and-sign signatures. IACR Cryptology ePrint Archive (2016). Full version of this paper
14. Fontein, F., Wocjan, P.: On the probability of generating a lattice. *J. Symb. Comput.* **64**, 3–15 (2014)
15. Gentry, C., Jonsson, J., Stern, J., Szydlo, M.: Cryptanalysis of the NTRU signature scheme (NSS) from Eurocrypt 2001. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 1–20. Springer, Heidelberg (2001). doi:[10.1007/3-540-45682-1_1](https://doi.org/10.1007/3-540-45682-1_1)
16. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Dwork, C. (ed.), STOC, pp. 197–206. ACM (2008)
17. Gentry, C., Szydlo, M.: Cryptanalysis of the revised NTRU signature scheme. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 299–320. Springer, Heidelberg (2002). doi:[10.1007/3-540-46035-7_20](https://doi.org/10.1007/3-540-46035-7_20)
18. Goldreich, O., Goldwasser, S., Halevi, S.: Public-key cryptosystems from lattice reduction problems. In: Kaliski, B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 112–131. Springer, Heidelberg (1997). doi:[10.1007/BFb0052231](https://doi.org/10.1007/BFb0052231)
19. Bruinderink, L.G., Hülsing, A., Lange, T., Yarom, Y.: Flush, gauss, and reload - a cache attack on the BLISS lattice-based signature scheme. IACR Cryptology ePrint Archive 2016:300 (2016)
20. Güneysu, T., Lyubashevsky, V., Pöppelmann, T.: Practical lattice-based cryptography: a signature scheme for embedded systems. In: Prouff, E., Schaumont, P. (eds.) CHES 2012. LNCS, vol. 7428, pp. 530–547. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-33027-8_31](https://doi.org/10.1007/978-3-642-33027-8_31)
21. Hoffstein, J., Howgrave-Graham, N., Pipher, J., Silverman, J.H., Whyte, W.: NTRUSign: digital signatures using the NTRU lattice. In: Joye, M. (ed.) CT-RSA 2003. LNCS, vol. 2612, pp. 122–140. Springer, Heidelberg (2003). doi:[10.1007/3-540-36563-X_9](https://doi.org/10.1007/3-540-36563-X_9)
22. Hoffstein, J., Pipher, J., Schanck, J.M., Silverman, J.H., Whyte, W.: Practical signatures from the partial fourier recovery problem. In: Boureanu, I., Owesarski, P., Vaudenay, S. (eds.) ACNS 2014. LNCS, vol. 8479, pp. 476–493. Springer, Cham (2014). doi:[10.1007/978-3-319-07536-5_28](https://doi.org/10.1007/978-3-319-07536-5_28)
23. Howe, J., Pöppelmann, T., O’Neill, M., O’Sullivan, E., Güneysu, T.: Practical lattice-based digital signature schemes. *ACM Trans. Embed. Comput. Syst.* **14**(3), 41 (2015)
24. Howe, J., Pöppelmann, T., O’Neill, M., O’Sullivan, E., Güneysu, T., Lyubashevsky, V.: Practical lattice-based digital signature schemes. In: Slides of the Presentation at the NIST Workshop of Cybersecurity in a Post-Quantum World (2015). <http://csrc.nist.gov/groups/ST/post-quantum-2015/presentations/session9-oneill-maire.pdf>
25. Kamal, A.A., Youssef, A.M.: Fault analysis of the NTRUSign digital signature scheme. *Cryptogr. Commun.* **4**(2), 131–144 (2012)
26. Lyubashevsky, V.: Fiat-Shamir with aborts: applications to lattice and factoring-based signatures. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 598–616. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-10366-7_35](https://doi.org/10.1007/978-3-642-10366-7_35)
27. Lyubashevsky, V.: Lattice signatures without trapdoors. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 738–755. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-29011-4_43](https://doi.org/10.1007/978-3-642-29011-4_43)

28. Maze, G., Rosenthal, J., Wagner, U.: Natural density of rectangular unimodular integer matrices. *Linear Algebra Appl.* **434**(5), 1319–1324 (2011)
29. Melchor, C.A., Boyen, X., Deneuville, J.-C., Gaborit, P.: Sealing the leak on classical NTRU signatures. In: Mosca, M. (ed.) *PQCrypto 2014*. LNCS, vol. 8772, pp. 1–21. Springer, Cham (2014). doi:[10.1007/978-3-319-11659-4_1](https://doi.org/10.1007/978-3-319-11659-4_1)
30. Micciancio, D., Peikert, C.: Trapdoors for lattices: simpler, tighter, faster, smaller. In: Pointcheval, D., Johansson, T. (eds.) *EUROCRYPT 2012*. LNCS, vol. 7237, pp. 700–718. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-29011-4_41](https://doi.org/10.1007/978-3-642-29011-4_41)
31. Naccache, D., Nguyen, P.Q., Tunstall, M., Whelan, C.: Experimenting with faults, lattices and the DSA. In: Vaudenay, S. (ed.) *PKC 2005*. LNCS, vol. 3386, pp. 16–28. Springer, Heidelberg (2005). doi:[10.1007/978-3-540-30580-4_3](https://doi.org/10.1007/978-3-540-30580-4_3)
32. Nguyen, P.Q., Regev, O.: Learning a parallelepiped: cryptanalysis of GGH and NTRU signatures. *J. Cryptol.* **22**(2), 139–160 (2009)
33. Ordas, S., Guillaume-Sage, L., Tobich, K., Dutertre, J.-M., Maurine, P.: Evidence of a larger EM-induced fault model. In: Joye, M., Moradi, A. (eds.) *CARDIS 2014*. LNCS, vol. 8968, pp. 245–259. Springer, Cham (2015). doi:[10.1007/978-3-319-16763-3_15](https://doi.org/10.1007/978-3-319-16763-3_15)
34. Page, D., Vercauteren, F.: A fault attack on pairing-based cryptography. *IEEE Trans. Comput.* **55**(9), 1075–1080 (2006)
35. Peikert, C.: A decade of lattice cryptography. *Cryptology ePrint Archive*, Report 2015/939 (2015). <http://eprint.iacr.org/>
36. Pöppelmann, T., Ducas, L., Güneysu, T.: Enhanced lattice-based signatures on reconfigurable hardware. In: Batina, L., Robshaw, M. (eds.) *CHES 2014*. LNCS, vol. 8731, pp. 353–370. Springer, Heidelberg (2014). doi:[10.1007/978-3-662-44709-3_20](https://doi.org/10.1007/978-3-662-44709-3_20)
37. Prest, T.: Implementation of the GPV-based scheme of Ducas et al. <https://github.com/tprest/Lattice-IBE>
38. Stein, W., et al.: Sage Mathematics Software (Version 7.0) (2016). <http://www.sagemath.org>
39. Taha, M., Eisenbarth, T.: Implementation attacks on post-quantum cryptographic schemes. In: Aleisa, E.A. (ed.) *ICACC*. IEEE Social Implications of Technology Society (2015)