# Practical Signatures from the Partial Fourier Recovery Problem 

Jeff Hoffstein, Jill Pipher, John Schanck, Joseph H. Silverman, and William Whyte


#### Abstract

We present PASSSign, a variant of the prior PASS and PASS-2 proposals, as a candidate for a practical post-quantum signature scheme. Its hardness is based on the problem of recovering a ring element with small norm from an incomplete description of its Chinese remainder representation. For our particular instantiation, this corresponds to the recovery of a signal with small infinity norm from a limited set of its Fourier coefficients.

The key improvement over previous versions of PASS is the introduction of a rejection sampling technique from Lyubashevsky (2009) which assures that transcript distributions are completely decoupled from the keys that generate them.

Although the scheme is not supported by a formal security reduction, we present extensive arguments for its security and derive concrete parameters based on the performance of state of the art lattice reduction and enumeration techniques.


## 1. Introduction

In the late 1990s two authors of the present paper proposed authentication and signature schemes based on the problem of recovering a polynomial with tightly concentrated coefficients given a small number of evaluations of that polynomial. The heuristic justification for the security of the scheme was that the uncertainty principle severely restricts how concentrated a signal can be in two mutually incoherent bases.

An early incarnation of the scheme is found in [12], and a later version, called PASS-2 was published in [13]. A rough description goes as follows. Let $N$ be a positive integer, and choose a prime $q=r N+1$, with $r \geq 1$. We will denote by $R_{q}$ the ring $\mathbb{Z}_{q}[x] /\left(x^{N}-1\right)$, though we will often treat elements of $R_{q}$ as vectors in $\mathbb{Z}_{q}^{N}$ equipped with the $\star$-multiplication of $R_{q}$. To avoid confusion, we will denote component-wise multiplication of vectors by $\odot$. For any $\beta$, with $(\beta, q)=1$, it follows from Fermat's little theorem that $\beta^{r N} \equiv 1(\bmod q)$. Consequently, the mapping $\boldsymbol{f} \rightarrow \boldsymbol{f}\left(\beta^{r}\right)$ is well defined for any $\boldsymbol{f}$ in $R_{q}$. In addition to being well defined, it is also a ring homomorphism, for the simple reason that for any $\boldsymbol{f}_{1}, \boldsymbol{f}_{2} \in R_{q}$,

$$
\left(\boldsymbol{f}_{1}+\boldsymbol{f}_{2}\right)\left(\beta^{r}\right)=\boldsymbol{f}_{1}\left(\beta^{r}\right)+\boldsymbol{f}_{2}\left(\beta^{r}\right) \quad \text { and } \quad\left(\boldsymbol{f}_{1} \star \boldsymbol{f}_{2}\right)\left(\beta^{r}\right)=\boldsymbol{f}_{1}\left(\beta^{r}\right) \odot \boldsymbol{f}_{2}\left(\beta^{r}\right) .
$$

More generally, for any $\Omega=\left\{\beta_{1}^{r}, \beta_{2}^{r}, \ldots, \beta_{t}^{r}\right\}$, the mapping $\mathcal{F}: R_{q} \rightarrow \mathbb{Z}_{q}^{t}$ given by

$$
\mathcal{F}_{\Omega} \boldsymbol{f}=\left(\boldsymbol{f}\left(\beta_{1}^{r}\right), \boldsymbol{f}\left(\beta_{2}^{r}\right), \ldots, \boldsymbol{f}\left(\beta_{t}^{r}\right)\right)^{T}
$$

is a ring homomorphism, with addition and $\odot$-multiplication modulo $q$ done on the right hand side. This is an example of the more general phenomenon of the ring homomorphism mapping functions to their Fourier transforms.

In the above setting, the uncertainty principle implies that a ring element with a coefficient vector drawn from a small region of $\mathbb{Z}_{q}^{N}$ will have widely dispersed discrete Fourier coefficients. For instance a vector with small infinity norm, e.g. with coefficients in $\{-1,0,1\}$, will likely be supported on all powers of a primitive $N^{t h}$ root $\omega$ and will have Fourier coefficients which are essentially uniformly distributed in $\mathbb{Z}_{q}$.

The hard problem in PASS can be stated as the following underdetermined linear inversion problem, which we will refer to as the partial Fourier recovery problem. Let $\omega$ be a primitive $N^{t h}$ root of unity modulo $q$. We define the discrete Fourier transform over $\mathbb{Z}_{q}$ to be the linear transformation $\mathcal{F} \boldsymbol{f}=\widehat{\boldsymbol{f}}: \mathbb{Z}_{q}^{N} \rightarrow \mathbb{Z}_{q}^{N}$ given by

$$
(\mathcal{F})_{i, j}=\omega^{i j}
$$

Furthermore, let $\mathcal{F}_{\Omega}$ be the restriction of $\mathcal{F}$ to the set of $t$ rows specified by an index set $\Omega$,

$$
\left(\mathcal{F}_{\Omega}\right)_{i, j}=\omega^{\Omega_{i} j}
$$

The partial Fourier recovery problem is: given an evaluation $\left.\widehat{\boldsymbol{f}}\right|_{\Omega} \in \mathbb{Z}_{q}^{t}$, find $\boldsymbol{x}$ with small norm such that $\left.\widehat{\boldsymbol{x}}\right|_{\Omega}=\left.\widehat{\boldsymbol{f}}\right|_{\Omega}(\bmod q)$.

The problem of recovering a signal from a restricted number of its Fourier coefficients is well studied and known to be quite difficult in general. The restricted image $\left.\widehat{\boldsymbol{f}}\right|_{\Omega}$ is expected to contain very little information about the unobserved Fourier coefficients (the evaluations of $\boldsymbol{f}$ on $\omega^{i}$ for $i$ not in $\Omega$ ), and often the only way to recover $\boldsymbol{f}$ will be an expensive combinatorial optimization procedure. However, there are cases (some quite surprising) in which the problem is known to be easy.

Certainly, if $t \log q$ is small, brute force search over $\boldsymbol{f}^{\prime}$ with appropriate norm may be a viable solution - each randomly chosen candidate having essentially a $q^{-t}$ chance of evaluating to $\left.\widehat{\boldsymbol{f}}\right|_{\Omega}$.

The problem is trivial in the large $t$ regime, $t \geq N$, since any rank $N$ submatrix of the chosen Vandermonde matrix will be invertible. As $t$ decreases slightly below $N$, or we allow some portion of the coefficients to be corrupted, the problem essentially becomes that of decoding Reed-Solomon codes and we can expect to recover $\boldsymbol{f}$ by list-decoding or similar techniques. Efficient recovery of general signals when $t$ is much less than $N$ would have significant coding theoretic implications.

For $t$ in an intermediate range, say $t \approx N / 2$, the situation is more complicated. Were one to consider the complex Fourier transform rather than the number theoretic transform, one might be able to apply techniques from the field of compressed sensing. Recent work in this field has delineated cases in which a sparse signal can be recovered from a limited number of its (complex) Fourier coefficients by an $L^{1}$ optimization procedure. For this to be successful the signals must be very sparse, having a number of non-zero coefficients which is less than $|\Omega| / 2$. [2]. It is not clear how these results translate into the finite field setting.

As far as we are aware, the best technique for solving the partial Fourier recovery problem is by solving an associated closest vector problem. Specifically, let $\Lambda^{\perp}\left(\mathcal{F}_{\Omega}\right)$ be the lattice of vectors in the kernel of $\mathcal{F}_{\Omega}$. That is,

$$
\Lambda^{\perp}\left(\mathcal{F}_{\Omega}\right)=\left\{\boldsymbol{a} \in \mathbb{Z}_{q}^{N}: \mathcal{F}_{\Omega} \boldsymbol{a}=\mathbf{0} \quad(\bmod q)\right\}
$$

If, given $\boldsymbol{y} \in \mathbb{Z}_{q}^{N}$, a point $\boldsymbol{x} \in \Lambda^{\perp}\left(\mathcal{F}_{\Omega}\right)$ can be found such that $\|\boldsymbol{y}-\boldsymbol{x}\|_{\infty} \leq \beta$, then $\mathcal{F}_{\Omega}(\boldsymbol{y}-\boldsymbol{x})=\left.\widehat{\boldsymbol{y}}\right|_{\Omega}$ and $\|\boldsymbol{y}-\boldsymbol{x}\|_{\infty} \leq \beta$. Since one can easily find (large) $\boldsymbol{y}$ such that $\left.\widehat{\boldsymbol{y}}\right|_{\Omega}=\left.\widehat{\boldsymbol{f}}\right|_{\Omega}$ for any evaluation set $\left.\widehat{\boldsymbol{f}}\right|_{\Omega}$, the ability to solve CVP in $\Lambda^{\perp}\left(\mathcal{F}_{\Omega}\right)$ implies the ability to solve arbitrary partial Fourier recovery instances

While there is no known reduction from standard lattice problems to the partial Fourier recovery problem, there is at very least a superficial relationship between finding short preimages of $\mathcal{F}_{\Omega}$ and another well studied hard problem. A great deal of the research in lattice based cryptography throughout the last decade has focused on a type of underdetermined linear inverse problem referred to as the small integer solution (SIS) problem.

SIS is the problem of of finding a vector $\boldsymbol{y}$ in the kernel of a specified linear transformation $\boldsymbol{A}: \mathbb{Z}_{q}^{n} \rightarrow \mathbb{Z}_{q}^{m}$ such that $\boldsymbol{y}$ is small with respect to a given norm. That is, the goal is to solve

$$
\boldsymbol{A} \boldsymbol{y}=\mathbf{0} \quad(\bmod q) \quad \text { and } \quad\|\boldsymbol{y}\| \leq \beta
$$

Ajtai showed in [1] that, for certain parameters and uniform random $\boldsymbol{A}$, SIS enjoys a remarkable average-case correspondence with worst-case lattice problems. That is to say that the ability to solve random SIS instances with non-negligible probability implies an ability to find short vectors in any lattice. This correspondence between worst and average cases is attractive from a provable security point of view, offering strong assurance that easy to generate instances of the SIS problem will be hard to solve, but it does not yield particularly efficient cryptosystems without additional assumptions.

The most efficient and compact SIS schemes in the literature are based on the Ideal-SIS problem, wherein the matrix $\boldsymbol{A}$ is replaced by several uniform random elements, $\boldsymbol{a}_{1}, \boldsymbol{a}_{2}, \ldots \boldsymbol{a}_{k}$ of a quotient ring $R_{q}^{\varphi}=\mathbb{Z}_{q}[x] /(\varphi)$. The polynomial $\varphi$ is typically, but not necessarily, cyclotomic. A solution to Ideal-SIS is $\boldsymbol{y}_{1}, \boldsymbol{y}_{2}, \ldots \boldsymbol{y}_{k}$ in the ring such that:

$$
\sum_{i=1}^{k} \boldsymbol{a}_{i} \star \boldsymbol{y}_{i}=\mathbf{0} \quad \text { and } \quad \sum_{i=1}^{k}\left\|\boldsymbol{y}_{i}\right\|^{2} \leq \beta^{2}
$$

These schemes derive their security from the presumed hardness of Ideal-SVP - the shortest vector problem in the restricted class of lattices generated by matrix representations of elements of $R_{q}^{\varphi}$. Reductions from worst-case Ideal-SVP to average-case Ideal-SIS were presented in [17] [20]. Unfortunately, even with the reduced storage requirements and fast multiplication algorithms available in some rings, provably secure Ideal-SIS based constructions are still too inefficient to be competitive with existing (non-quantum resistant) schemes.

The security of PASS can be said to rest on the assumed average-case hardness of Vandermonde-SIS. We are not aware of any technique for reducing a worst-case lattice problem to Vandermonde-SIS, nor will we postulate the existence of such a
reduction. We do however raise the question of whether there might be a characterization of hard instances of SIS which does not rely on structural properties of the matrix $\boldsymbol{A}$. Or more generally, when is a constrained linear inverse problem hard?

We believe an answer to this problem would likely simultaneously explain the hardness of Uniform-, Ideal- and Vandermonde-SIS, as well as delineate new classes.

## 2. Related Work

2.1. The original PASS protocols. Given a (padded) message $\mu$, a secret key $\boldsymbol{f}$ with small norm, and a public key $\left.\widehat{\boldsymbol{f}}\right|_{\Omega}=\mathcal{F}_{\Omega} \boldsymbol{f}$, the objective is to construct a signature that mixes $\boldsymbol{f}$ and $\mu$ and can be verified by means of $\left.\widehat{\boldsymbol{f}}\right|_{\Omega}$. A prototype of this was presented in [12].

To sign, Alice

- Computes and keeps secret a short polynomial $\boldsymbol{g} \in R_{q}$ and reveals the commitment $\left.\widehat{\boldsymbol{g}}\right|_{\Omega}=\mathcal{F}_{\Omega} \boldsymbol{g}$.
- Computes and reveals a short challenge polynomial $\boldsymbol{c} \in R_{q}$ from $\operatorname{Hash}\left(\left.\widehat{\boldsymbol{g}}\right|_{\Omega}, \mu\right)$.
- Computes and reveals $\boldsymbol{h}=\boldsymbol{g} \star(\boldsymbol{f}+\boldsymbol{c})$.

To verify, Bob

- Verifies that $\boldsymbol{h}$ has norm less than a specific upper bound.
- Verifies that $\boldsymbol{c}=\operatorname{Hash}\left(\left.\widehat{\boldsymbol{h}}\right|_{\Omega} /\left(\left.\widehat{\boldsymbol{f}}\right|_{\Omega}+\left.\widehat{\boldsymbol{c}}\right|_{\Omega}\right), \mu\right)$

The first condition for verification is met because

$$
\|\boldsymbol{g} \star(\boldsymbol{f}+\boldsymbol{c})\| \approx\|\boldsymbol{g}\|\|\boldsymbol{f}+\boldsymbol{c}\|
$$

The fact that $\|\boldsymbol{f}\|,\|\boldsymbol{g}\|,\|\boldsymbol{c}\|$ are small thus implies that $\|\boldsymbol{h}\|$ is small ${ }^{1}$. The second condition is true because $\mathcal{F}_{\Omega}$ is a ring homomorphism.

To forge a signature, a third party would need to produce an $\boldsymbol{h}$ which is short, and which satisfies the required evaluations at points in $\Omega$. It is conjectured that finding such an $\boldsymbol{h}$ is no easier than solving the associated CVP.
2.2. Transcript weaknesses in previous PASS protocols. The difficulty with this PASS prototype is that a transcript of signatures produced by a single signer on any set of messages leaks information about that signer's secret key. One way to see this is via a ring homomorphism $\rho: R_{q} \rightarrow R_{q}$ given by

$$
\rho\left(a_{0}+a_{1} \boldsymbol{x}+a_{2} \boldsymbol{x}^{2}+\cdots+a_{N-1} \boldsymbol{x}^{N-1}\right)=a_{0}+a_{N-1} \boldsymbol{x}+a_{N-2} \boldsymbol{x}^{2}+\cdots+a_{1} \boldsymbol{x}^{N-1}
$$

The homomorphism $\rho$ plays the same role that conjugation would play if $x$ were replaced by a primitive $N^{t h}$ root of unity. If a polynomial $\boldsymbol{p} \in R_{q}$ is drawn randomly from a distribution, let $\mathbb{E}[\boldsymbol{p}]$ denote the expectation of $\boldsymbol{p}$, that is, the average of $\boldsymbol{p}$ over many samples. A third party observing many examples of $\boldsymbol{g} \star(\boldsymbol{f}+\boldsymbol{c})$ could compute

$$
\mathbb{E}[\boldsymbol{g} \star(\boldsymbol{f}+\boldsymbol{c}) \star \rho(\boldsymbol{g} \star(\boldsymbol{f}+\boldsymbol{c}))]=\mathbb{E}[\boldsymbol{g} \star \rho(\boldsymbol{g})] \mathbb{E}[(\boldsymbol{f}+\boldsymbol{c}) \star \rho(\boldsymbol{f}+\boldsymbol{c})]
$$

For simplicity assume that $\mathbb{E}[\boldsymbol{c}]=0$, then, since $\boldsymbol{f}$ is constant, the above becomes

$$
\mathbb{E}[\boldsymbol{g} \star \rho(\boldsymbol{g})](\mathbb{E}[\boldsymbol{c} \star \rho(\boldsymbol{c})]+\boldsymbol{f} \star \rho(\boldsymbol{f})) .
$$

[^0]The distributions from which $\boldsymbol{c}$ and $\boldsymbol{g}$ are drawn are known, and thus a sufficiently long transcript will reveal $\boldsymbol{f} \star \rho(\boldsymbol{f})$ from which $\boldsymbol{f}$ may be computed by a technique from Gentry and Szydlo [8].
2.3. Recent developments and countermeasures. The problem with PASS was not that individual signatures leaked information about the secret key, but rather that an average over a collection of signatures would converge to a secret key dependent value. This is not a concern for signature schemes based on number theoretic trapdoor permutations, as such schemes enjoy relatively simple proofs that their signatures are uniformly distributed over the full range of possibilities. However, the requirement that PASS signatures have small norm, i.e. that they occupy a small region of the full domain, necessitates throwing out much of the algebraic structure that makes such uniformity guarantees possible. Full decoupling of secret keys from transcripts was a difficult barrier for the construction of secure lattice based signature schemes, and more so for the construction of efficient schemes.

The first successful decoupling, the signature scheme of Gentry, Peikert, and Vaikuntanathan [7], involved computing a candidate signature point $\boldsymbol{x}$ and then adding noise sampled from a discrete Gaussian distribution centered at $-\boldsymbol{x}$. The resulting signatures have a distribution which is computationally indistinguishable from a spherical discrete Gaussian centered at the origin.

Lyubashevsky, in [14], constructed a lattice based identification scheme which avoids transcript analysis attacks with a technique he called "aborting." In this scheme, provers are capable of determining when their response to a challenge will leak information about their secret key. Whenever this occurs they abort the protocol rather than supply a response.

In [15], Lyubashevsky improved his aborting technique and constructed a signature scheme through the Fiat-Shamir transform with hardness based on the RingSIS problem. Improvements and variants of this scheme with different hardness assumptions were presented in [16].

The first truly practical lattice signature scheme to avoid transcript attacks was developed by Güneysu, Lyubashevsky, and Pöppelmann [9]. Their scheme is a highly optimized variant of [16] and relies on a stronger hardness assumption.

The current state of the art would appear to be the new scheme, called BLISS, by Ducas, Durmus, Lepoint, and Lyubashevsky [4]. This scheme makes use of an NTRU-like key generation procedure and a bimodal discrete Gaussian noise distribution to produce very compact signatures. The efficiency of the scheme is also very impressive, especially considering the complexity of sampling discrete Gaussians.

## 3. PASS with Rejection Sampling

We now present a new variant of PASS which completely decouples the transcript distribution from the secret key. Table 1 lists the public parameters of the system and gives a brief description of each.

Some notes on notation: $R_{q}$ is the ring $\mathbb{Z}_{q}[x] /\left(x^{N}-1\right)$; elements $\boldsymbol{a} \in R_{q}$ are represented as polynomials $\boldsymbol{a}=a_{0}+a_{1} \boldsymbol{x}+a_{2} \boldsymbol{x}^{2}+\cdots+a_{N-1} \boldsymbol{x}^{N-1}$, with coefficients in $a_{i} \in \mathbb{Z}_{q}$. We freely transition between this polynomial representation and a coefficient vector representation, $\boldsymbol{a}=\left[a_{0}, a_{1}, a_{2}, \ldots, a_{N-1}\right]^{T}$, wherever convenient.
$N$ Dimension
$q \quad$ Prime $\equiv 1(\bmod N)$
$g \quad$ a primitive $N^{t h}$ root of unity in $\mathbb{Z}_{q}$
$\Omega \quad$ A subset of $\left\{g^{j}: 1 \leq j \leq N-1\right\}$
$t \quad|\Omega|$
$k$ Infinity norm of noise polynomials
$b \quad$ 1-norm of challenge polynomials
TABLE 1. List of public parameters.

Norms, such as $\|\boldsymbol{a}\|_{\infty}$ and $\|\boldsymbol{a}\|_{1}$, are the standard $L^{p}$ norms on coefficient vectors; for numerical calculations we consistently identify $a_{i}$ with an integer such that $\left|a_{i}\right| \leq q / 2$.

We write $\mathcal{B}^{1}(b)$ to denote the elements of $R_{q}$ with 1-norm $\leq b$, and $\mathcal{B}^{\infty}(k)$ to denote the elements of $R_{q}$ with $\infty$-norm $\leq k$.

Lastly, The indicator function $\mathbf{1}_{\mathcal{S}}(x)$ yields 1 if $x \in \mathcal{S}$ and 0 otherwise.
3.1. Key Generation. A secret key is a polynomial with $L^{\infty}$ norm equal to 1 . We recommend the simple strategy of choosing each coefficient independently and uniformly from $\{-1,0,1\}$. Binary coefficients, though attractive for several reasons, would open the system up to a UniqueSVP gap amplification attack similar to that used by Nguyen in his cryptanalysis of GGH [19].

The public key corresponding to the secret key $\boldsymbol{f}$ is $\left.\widehat{\boldsymbol{f}}\right|_{\Omega}=\mathcal{F}_{\Omega} \boldsymbol{f}$.
3.2. Signing. Signing is an iterated process consisting of the generation of a candidate signature followed by a rejection sampling step to prevent the publication of candidates that could leak secret key information.

A party with secret key $\boldsymbol{f}$, who wishes to sign a message $\mu$, first selects a commitment polynomial $\boldsymbol{y}$ uniformly at random from $\mathcal{B}^{\infty}(k)$. The commitment $\boldsymbol{y}$ serves to mask the private key and must be treated with the same care as the private key itself. The signer then computes and stores $\left.\widehat{\boldsymbol{y}}\right|_{\Omega}=\mathcal{F}_{\Omega} \boldsymbol{y}$, which will ultimately be made public if the candidate passes rejection sampling. ${ }^{2}$

Next, the signer computes a challenge, $\boldsymbol{c}$, which binds $\left.\widehat{\boldsymbol{y}}\right|_{\Omega}$ to $\mu$. To do so she makes use of the public algorithms:

$$
\begin{aligned}
& \text { Hash : } \mathbb{Z}_{q}^{t} \times\{0,1\}^{*} \rightarrow\{0,1\}^{\ell}, \text { and } \\
& \quad \text { FormatC }:\{0,1\}^{\ell} \hookrightarrow \mathcal{B}^{1}(b) .
\end{aligned}
$$

Hash concatenates its inputs and passes the result through a cryptographic hash function such as sha512. FormatC maps the set of bitstrings output by Hash into a set of sparse polynomials. We avoid further description of the algorithms for now and simply say that

$$
\boldsymbol{c}=\operatorname{FormatC}\left(\operatorname{Hash}\left(\left.\widehat{\boldsymbol{y}}\right|_{\Omega}, \mu\right)\right)
$$

Finally, the signer computes a candidate signature point

$$
\boldsymbol{z}=\boldsymbol{f} \star \boldsymbol{c}+\boldsymbol{y} \in R_{q},
$$

[^1]if any of the coefficients of $\boldsymbol{z}$ fall outside the interval $[-k+b, k-b]$, then $\boldsymbol{y}, \boldsymbol{c}$, and $\boldsymbol{z}$ are discarded and the signing process is repeated. Otherwise, the signer outputs the signature $(\boldsymbol{c}, \boldsymbol{z}, \mu)$.

In section 4 we will prove that signatures that pass the rejection sampling procedure have $\boldsymbol{z}$ values that are uniformly distributed over $\mathcal{B}^{\infty}(k-b)$.
3.3. Verification. The signature $(\boldsymbol{c}, \boldsymbol{z}, \mu)$ is valid if $\boldsymbol{z}$ is in $\mathcal{B}^{\infty}(k-b)$ and if

$$
\boldsymbol{c}=\operatorname{FormatC}\left(\operatorname{Hash}\left(\left.\widehat{\boldsymbol{z}}\right|_{\Omega}-\left.\left.\widehat{\boldsymbol{f}}\right|_{\Omega} \odot \widehat{\boldsymbol{c}}\right|_{\Omega}, \mu\right)\right)
$$

Since $\mathcal{F}_{\Omega}$ is a ring homomorphism, it is the case that $\left.\widehat{\boldsymbol{z}}\right|_{\Omega}=\left.\left.\widehat{\boldsymbol{f}}\right|_{\Omega} \odot \widehat{\boldsymbol{c}}\right|_{\Omega}+\left.\widehat{\boldsymbol{y}}\right|_{\Omega}$. Therefore, on receipt of $(\boldsymbol{c}, \boldsymbol{z}, \mu)$, any verifier in possession of the appropriate public key $\left.\widehat{\boldsymbol{f}}\right|_{\Omega}$ can evaluate $\boldsymbol{z}$ and $\boldsymbol{c}$ and compute $\left.\widehat{\boldsymbol{y}}\right|_{\Omega}=\left.\widehat{\boldsymbol{z}}\right|_{\Omega}-\left.\left.\widehat{\boldsymbol{f}}\right|_{\Omega} \odot \widehat{\boldsymbol{c}}\right|_{\Omega}$. The correctness of the scheme is immediate.

```
Algorithm 1 Sign
Input: \((\mu, \boldsymbol{f})\)
    repeat
        \(\boldsymbol{y} \stackrel{\$}{\leftarrow} \mathcal{B}^{\infty}(k)\)
        \(h \leftarrow \operatorname{Hash}\left(\left.\widehat{\boldsymbol{y}}\right|_{\Omega}, \mu\right)\)
        \(\boldsymbol{c} \leftarrow\) FormatC \((h)\)
        \(\boldsymbol{z} \leftarrow \boldsymbol{f} \star \boldsymbol{c}+\boldsymbol{y}\)
    until \(\boldsymbol{z} \in \mathcal{B}^{\infty}(k-b)\)
Output: \((\boldsymbol{c}, \boldsymbol{z}, \mu)\)
```

```
Algorithm 2 Verify
Input: \(\left(\boldsymbol{c}, \boldsymbol{z}, \mu,\left.\widehat{\boldsymbol{f}}\right|_{\Omega}\right)\)
    result \(\leftarrow\) invalid
    if \(\boldsymbol{z} \in \mathcal{B}^{\infty}(k-b)\) then
        \(h^{\prime} \leftarrow \operatorname{Hash}\left(\left.\widehat{\boldsymbol{z}}\right|_{\Omega}-\left.\left.\widehat{\boldsymbol{f}}\right|_{\Omega} \odot \widehat{\boldsymbol{c}}\right|_{\Omega}, \mu\right)\)
        \(c^{\prime} \leftarrow\) Format \(C\left(h^{\prime}\right)\)
        if \(c=c^{\prime}\) then
            result \(\leftarrow\) valid
        end if
    end if
Output: result
```


## 4. Rejection Sampling

Each iteration of the signature generation routine produces a candidate signature which is accepted or rejected based on its infinity norm alone. In this section we will argue that this rejection sampling procedure completely decouples the distribution of signature points from the private key.

We will make use of the following fact:
FACT 1. Each candidate signature $\boldsymbol{z}$ is a member of $\mathcal{B}^{\infty}(k+b)$.
Proof. By definition we have $\|\boldsymbol{z}\|_{\infty}=\|\boldsymbol{f} \star \boldsymbol{c}+\boldsymbol{y}\|_{\infty}$ and by the triangle inequality $\|\boldsymbol{f} \star \boldsymbol{c}+\boldsymbol{y}\|_{\infty} \leq\|\boldsymbol{f} \star \boldsymbol{c}\|_{\infty}+\|\boldsymbol{y}\|_{\infty}$. Again by the triangle inequality $\|\boldsymbol{f} \star \boldsymbol{c}\|_{\infty} \leq\|\boldsymbol{f}\|_{\infty}\|\boldsymbol{c}\|_{1}$, thus

$$
\|\boldsymbol{z}\|_{\infty} \leq\|\boldsymbol{f}\|_{\infty}\|\boldsymbol{c}\|_{1}+\|\boldsymbol{y}\|_{\infty} \leq b+k
$$

We will also make use of the following assumption on instantiations of Hash and FormatC.

Assumption 1. Let the public parameters $(N, q, k, b, \Omega)$ be fixed and let $\boldsymbol{c} \in$ $\mathcal{B}^{1}(b), \boldsymbol{y} \in \mathcal{B}^{\infty}(k), \mu \in\{0,1\}^{*}$ be random variables related by

$$
\boldsymbol{c}=\operatorname{FormatC}\left(\operatorname{Hash}\left(\left.\widehat{\boldsymbol{y}}\right|_{\Omega}, \mu\right)\right)
$$

We assume that Hash is a collision resistant hash function, that $\boldsymbol{c}$ and $\boldsymbol{y}$ are independent, and that $\boldsymbol{c}$ is uniform over the range of FormatC. More explicitly, for any fixed $\boldsymbol{c}_{0} \in \mathcal{B}^{1}(b)$ and fixed $\boldsymbol{y}_{0} \in \mathcal{B}^{\infty}(k)$,

$$
\operatorname{Pr}\left[\boldsymbol{c}=\boldsymbol{c}_{0} \mid \boldsymbol{y}=\boldsymbol{y}_{0}\right]=\frac{\operatorname{Pr}\left[\boldsymbol{c}=\boldsymbol{c}_{0}\right] \operatorname{Pr}\left[\boldsymbol{y}=\boldsymbol{y}_{0}\right]}{\operatorname{Pr}\left[\boldsymbol{y}=\boldsymbol{y}_{0}\right]}=\left|\mathcal{B}^{1}(b)\right|^{-1}
$$

Note that assumption 1 is no stronger than the standard random oracle assumption, so the reader may assume we are working in the random oracle model. We state the assumption in the above form to aid in the analysis of concrete instantiations. Clearly the assumption that the joint distribution of $\boldsymbol{y}$ and $\boldsymbol{c}$ factors is untenable - no deterministic instantiation of Hash can satisfy it while maintaining collision resistance. Yet by choosing an appropriate padding scheme for $\mu$ one should be able to approximately satisfy the assumption. We leave the exploration of padding schemes and analysis of the practical impact of assumption 1 to future work.

The following proposition describes the distribution of candidate signatures.
Proposition 4.1. Fix vectors $\boldsymbol{f}_{0} \in \mathcal{B}^{\infty}(1)$ and $\boldsymbol{z}_{0} \in \mathcal{B}^{\infty}(k+b)$. Then as the pair $(\boldsymbol{c}, \boldsymbol{y})$ is chosen uniformly from the space $\mathcal{B}^{1}(1) \times \mathcal{B}^{\infty}(k)$, we have

$$
\operatorname{Pr}\left[\boldsymbol{f}_{0} \star \boldsymbol{c}+\boldsymbol{y}=\boldsymbol{z}_{0}\right]=\left|\mathcal{B}^{\infty}(k)\right|^{-1} \sum_{\boldsymbol{c}_{0} \in \mathcal{B}^{1}(b)} \operatorname{Pr}\left[\boldsymbol{c}=\boldsymbol{c}_{0}\right] \mathbf{1}_{\mathcal{B}^{\infty}(k)}\left(\boldsymbol{z}_{0}-\boldsymbol{f}_{0} \star \boldsymbol{c}_{0}\right)
$$

Proof. For any fixed $\boldsymbol{c}_{0} \in \mathcal{B}^{1}(b)$ we have

$$
\begin{aligned}
\operatorname{Pr}\left[\boldsymbol{f}_{0} \star \boldsymbol{c}_{0}+\boldsymbol{y}=\boldsymbol{z}_{0}\right] & =\operatorname{Pr}\left[\boldsymbol{y}=\boldsymbol{z}_{0}-\boldsymbol{f}_{0} \star \boldsymbol{c}_{0}\right] \\
& = \begin{cases}\left|\mathcal{B}^{\infty}(k)\right|^{-1} & \text { if }\left(\boldsymbol{z}_{0}-\boldsymbol{f}_{0} \star \boldsymbol{c}_{0}\right) \in \mathcal{B}^{\infty}(k) \\
0 & \text { otherwise } .\end{cases}
\end{aligned}
$$

By application of the law of total probability and the assumption that the $\boldsymbol{c}$ and $\boldsymbol{y}$ are independent:

$$
\begin{aligned}
\operatorname{Pr}\left[\boldsymbol{f}_{0} \star \boldsymbol{c}+\boldsymbol{y}=\boldsymbol{z}_{0}\right] & =\sum_{\boldsymbol{c}_{0} \in \mathcal{B}^{1}(b)} \operatorname{Pr}\left[\boldsymbol{c}=\boldsymbol{c}_{0}\right] \operatorname{Pr}\left[\boldsymbol{f}_{0} \star \boldsymbol{c}+\boldsymbol{y}=\boldsymbol{z}_{0} \mid \boldsymbol{c}=\boldsymbol{c}_{0}\right] \\
& =\sum_{\boldsymbol{c}_{0} \in \mathcal{B}^{1}(b)} \operatorname{Pr}\left[\boldsymbol{c}=\boldsymbol{c}_{0}\right] \operatorname{Pr}\left[\boldsymbol{y}=\boldsymbol{z}_{0}-\boldsymbol{f}_{0} \star \boldsymbol{c}_{0}\right] \\
& =\left|\mathcal{B}^{\infty}(k)\right|^{-1} \sum_{\boldsymbol{c}_{0} \in \mathcal{B}^{1}(b)} \operatorname{Pr}\left[\boldsymbol{c}=\boldsymbol{c}_{0}\right] \mathbf{1}_{\mathcal{B}^{\infty}(k)}\left(\boldsymbol{z}_{0}-\boldsymbol{f}_{0} \star \boldsymbol{c}_{0}\right) .
\end{aligned}
$$

Recall from section 3.2 that a candidate signature is rejected unless its $\boldsymbol{z}$ component is contained in $\mathcal{B}^{\infty}(k-b)$. The following proposition shows that each point in $\mathcal{B}^{\infty}(k-b)$ is selected as a candidate signature with equal probability.

Proposition 4.2. Fix vectors $\boldsymbol{f}_{0}$ in $\mathcal{B}^{\infty}(1)$ and $\boldsymbol{z}_{0}$ in $\mathcal{B}^{\infty}(k-b)$. Then as the pair $(\boldsymbol{c}, \boldsymbol{y})$ is chosen uniformly from the space $\mathcal{B}^{1}(b) \times \mathcal{B}^{\infty}(k)$, we have

$$
\operatorname{Pr}\left[\boldsymbol{f}_{0} \star \boldsymbol{c}+\boldsymbol{y}=\boldsymbol{z}_{0}\right]=\left|\mathcal{B}^{\infty}(k)\right|^{-1} .
$$

Proof. We first note that $\mathcal{B}^{\infty}(k-b)$ is contained within $\mathcal{B}^{\infty}(k+b)$, so proposition 4.1 applies. Additionally, it is the case that $\left\|\boldsymbol{z}_{0}\right\|_{\infty} \leq k-b$ and consequently, for any fixed $\boldsymbol{c}_{0} \in \mathcal{B}^{1}(b)$, we have $\left\|\boldsymbol{z}_{0}-\boldsymbol{f}_{0} \star \boldsymbol{c}_{0}\right\|_{\infty} \leq k$. Thus $\boldsymbol{z}_{0}-\boldsymbol{f}_{0} \star \boldsymbol{c}_{0}$ is contained in $\mathcal{B}^{\infty}(k)$ and the indicator function in proposition 4.1 is unconditionally satisfied. Therefore,

$$
\operatorname{Pr}\left[\boldsymbol{f}_{0} \star \boldsymbol{c}+\boldsymbol{y}=\boldsymbol{z}_{0}\right]=\left|\mathcal{B}^{\infty}(k)\right|^{-1} \sum_{\boldsymbol{c}_{0} \in \mathcal{B}^{1}(b)} \operatorname{Pr}\left[\boldsymbol{c}=\boldsymbol{c}_{0}\right]=\left|\mathcal{B}^{\infty}(k)\right|^{-1}
$$

Proposition 4.2 informs us that each of the $\left|\mathcal{B}^{\infty}(k-b)\right|$ acceptable signature points is chosen with probability $\left|\mathcal{B}^{\infty}(k)\right|^{-1}$. We infer that each pass through the signature generation routine has probability

$$
\operatorname{Pr}[\text { accept }]=\frac{\left|\mathcal{B}^{\infty}(k-b)\right|}{\left|\mathcal{B}^{\infty}(k)\right|}=\left(1-\frac{2 b}{2 k+1}\right)^{N} \approx e^{-\frac{N b}{k}}
$$

of generating a valid signature point, where the approximation is valid provided that both $N$ and $k / b$ are large.

A transcript is a set of signatures published by an honest signer. For instance, a signer who uses private key $\boldsymbol{f}$ to sign messages $\mu_{1}, \mu_{2}, \ldots, \mu_{k}$ produces a transcript

$$
T=\left\{\left(\boldsymbol{c}_{i}, \boldsymbol{z}_{i}\right):\left(\boldsymbol{c}_{i}, \boldsymbol{z}_{i}, \mu_{i}\right)=\operatorname{Sign}\left(\mu_{i}, \boldsymbol{f}\right)\right\} .
$$

Proposition 4.3. A transcript $T$ generated by an honest signer with private key $\boldsymbol{f}$ is indistinguishable from a set of points drawn uniformly from $\mathcal{B}^{1}(b) \times \mathcal{B}^{\infty}(k-$ b).

Furthermore, for any fixed $\boldsymbol{c}_{0} \in \mathcal{B}^{1}(b), \boldsymbol{z}_{0} \in \mathcal{B}^{\infty}(k-b)$ and $\boldsymbol{f}_{0} \in \mathcal{B}^{1}(1)$, the events $\left(\boldsymbol{c}_{0}, \boldsymbol{z}_{0}\right) \in T$ and $\boldsymbol{f}=\boldsymbol{f}_{0}$ are independent.

Proof. The $\boldsymbol{c}$ components of $T$ are uniformly distributed over $\mathcal{B}^{1}(b)$ by assumption 1. Proposition 4.2 establishes not only that the $\boldsymbol{z}$ components of $T$ are uniformly distributed over $\mathcal{B}^{\infty}(k-b)$, but also that the distribution of $\boldsymbol{z}$ depends only on the distribution of $\boldsymbol{y}$. Again by assumption 1, $\boldsymbol{c}$ and $\boldsymbol{y}$ are independent and therefore $\boldsymbol{c}$ and $\boldsymbol{z}$ are independent. The distribution of transcript points is consequently the product distribution of $\boldsymbol{c}$ and $\boldsymbol{z}$, i.e. uniform over $\mathcal{B}^{1}(b) \times \mathcal{B}^{\infty}(k-b)$.

Independence of transcript points from the secret key follows from the fact that proposition 4.2 holds for all choices of $\boldsymbol{f}_{0}$ in $\mathcal{B}^{\infty}(1)$.

## 5. Security Analysis

Our security analysis will focus on two types of attacks, those that target the hash function (or the combination FormatCoHash), and those that target the partial Fourier transform $\mathcal{F}_{\Omega}$. Other attacks may be possible, and investigating them is an area for future work.

As our aim is to develop a practical quantum-resistant signature scheme, we will assume that the adversary has access to a quantum computer. Relatively little
is known about the existence or non-existence of quantum algorithms for lattice problems, so our assumptions related to quantum computers will only address their ability to solve $k$-element black-box search problems in $\Theta(\sqrt{k})$ time.
5.1. Attacks on the hash function. The most obvious constraint on the security of the system comes from the entropy of $\boldsymbol{c}$. An adversary who can find a Hash preimage of a particular $\boldsymbol{c}$ can produce forgeries on structured messages from any user's public key. To do so, the adversary:
(1) Chooses arbitrary $\boldsymbol{z}$ and $\boldsymbol{c}$ from the appropriate domains.
(2) Computes $\left.\widehat{\boldsymbol{g}}\right|_{\Omega}=\left.\widehat{\boldsymbol{z}}\right|_{\Omega}-\left.\left.\widehat{\boldsymbol{f}}\right|_{\Omega} \widehat{\boldsymbol{c}}\right|_{\Omega}$, where $\left.\widehat{\boldsymbol{f}}\right|_{\Omega}$ is the victim's public key.
(3) Finds a preimage of $\boldsymbol{c}$ in $\operatorname{Hash}\left(\left.\widehat{\boldsymbol{g}}\right|_{\Omega}, \cdot\right)$.

While attacks against specific hash functions can have arbitrarily low complexity, we will assume that a strong hash function is chosen, and only consider generic attacks. If the output of Hash is $r$ bits, a quantum adversary can find preimages in time $\Theta\left(2^{r / 2}\right)$. For $\kappa$-bit security, the range of FormatC $\circ$ Hash should produce an essentially uniform distribution on a set of cardinality $2^{2 \kappa}$.
5.2. Attacks on the partial Fourier transform. An adversary who can find $\mathcal{F}_{\Omega}$ preimages which are in $\mathcal{B}^{\infty}(k-b)$ can forge signatures on arbitrary messages from any user's public key.
(1) Adversary chooses random point $\boldsymbol{g}_{F}$ in $\mathcal{B}^{\infty}(k)$
(2) $\boldsymbol{c}_{F}=$ FormatC $\left(\operatorname{Hash}\left(\mathcal{F}_{\Omega} \boldsymbol{g}_{F}, \mu\right)\right)$
(3) $\left.\widehat{\boldsymbol{z}_{F}}\right|_{\Omega}=\left.\widehat{\boldsymbol{g}_{F}}\right|_{\Omega}+\left.\left.\widehat{\boldsymbol{f}}\right|_{\Omega} \widehat{\boldsymbol{c}_{F}}\right|_{\Omega}$
(4) Adversary uses preimage attack on $\left.\widehat{\boldsymbol{z}_{F}}\right|_{\Omega}$ to find appropriate $\boldsymbol{z}_{F}$.

Adversaries could also try to recover the secret key directly with their preimage algorithm, but in order for this to be effective they must be able to find exceptionally short preimages. The problem of secret key recovery seems, at least intuitively then, to be harder than forgery. Yet, surprisingly, given the particular parameters of the scheme, lattice attacks may be better suited for solving the secret key recovery problem than they are for forging messages. Some care must be taken when choosing parameters to balance the difficulty of the two problems.
5.2.1. Lattice attacks on $\mathcal{F}_{\Omega}$. As mentioned briefly in the introduction, the partial Fourier recovery problem PFRP can easily be seen to be no harder than a specific class of closest vector problem CVP. Presented with the evaluation set, $\Omega$, for a PFRP, and a partial Fourier representation $\left.\widehat{\boldsymbol{z}}\right|_{\Omega}$, an adversary can construct a lattice in which solving the CVP associated to any arbitrary preimage of $\left.\widehat{\boldsymbol{z}}\right|_{\Omega}$ allows them to construct a short preimage of $\left.\widehat{\boldsymbol{z}}\right|_{\Omega}$.

That lattice, which we denote $\Lambda^{\perp}\left(\mathcal{F}_{\Omega}\right)$, is equivalent to the kernel of $\mathcal{F}_{\Omega}$,

$$
\Lambda^{\perp}\left(\mathcal{F}_{\Omega}\right)=\left\{\boldsymbol{a} \in \mathbb{Z}_{q}^{N}: \mathcal{F}_{\Omega} \boldsymbol{a}=\mathbf{0} \quad(\bmod q)\right\}
$$

In practice, CVP instances are almost always solved by transforming the problem into an SVP in dimension $N+1$. If $\boldsymbol{z}^{\prime}$ is an arbitrary preimage of the target $\left.\widehat{\boldsymbol{z}}\right|_{\Omega}$, i.e. $\mathcal{F}_{\Omega} \boldsymbol{z}^{\prime}=\left.\widehat{\boldsymbol{z}}\right|_{\Omega}$ but $\left\|\boldsymbol{z}^{\prime}\right\|$ is large, and $\left\{\boldsymbol{b}_{1}, \boldsymbol{b}_{2}, \ldots, \boldsymbol{b}_{m}\right\}$ form a Hermite Normal Form basis for $\Lambda^{\perp}\left(\mathcal{F}_{\Omega}\right)$, then solving SVP in the lattice generated by the columns of

$$
\mathcal{L}_{\boldsymbol{z}^{\prime}}^{\mathrm{SVP}}=\left(\begin{array}{ccccccc}
q & & 0 & \boldsymbol{b}_{1,0} & \ldots & \boldsymbol{b}_{m, 0} & \boldsymbol{z}_{0}^{\prime} \\
& \ddots & & \vdots & \ddots & \vdots & \vdots \\
0 & & q & \boldsymbol{b}_{1, t-1} & \ldots & \boldsymbol{b}_{m, t-1} & \boldsymbol{z}_{t-1}^{\prime} \\
0 & \ldots & 0 & \vdots & \ddots & \vdots & \vdots \\
\vdots & \ddots & \vdots & \boldsymbol{b}_{1, N-1} & \ldots & \boldsymbol{b}_{m, N-1} & \boldsymbol{z}_{N-1}^{\prime} \\
0 & \ldots & 0 & 0 & \ldots & 0 & 1
\end{array}\right)
$$

is likely to yield a short $\boldsymbol{z}$ such that $\mathcal{F}_{\Omega} \boldsymbol{z}=\left.\widehat{\boldsymbol{z}}\right|_{\Omega}$.
Experiments by Micciancio and Regev [18] have demonstrated that lattice reduction algorithms perform best against the kernel lattices, $\Lambda^{\perp}(\boldsymbol{A})$, of $t \times N$ matrices $\boldsymbol{A}$ when $N \approx \sqrt{t \log (q) / \log (\gamma)}$ for some $\gamma \approx 1.01$ determined precisely by the reduction algorithm. In the PASS setting this places restrictions on $t$ and $q$ that we have obeyed in all of our proposed parameter sets. As such there should be no benefit to attacking a sublattice of $\mathcal{L}^{\mathrm{SVP}}$, and we proceed under this assumption.

The performance of lattice reduction algorithms, particularly LLL and BKZ, on lattices such as $\mathcal{L}_{z^{\prime}}^{\text {SVP }}$ is difficult to analyze in practice. Perhaps the most surprising complicating factor is that the performance depends crucially on the coset of $\mathbb{Z}_{q}^{N} / \Lambda^{\perp}\left(\mathcal{F}_{\Omega}\right)$ to which $\boldsymbol{z}^{\prime}$ belongs, and not strongly on $\boldsymbol{z}^{\prime}$ itself. This dependence gives rise to two regimes that we will analyze separately. The extreme case, when $\boldsymbol{z}^{\prime}$ is very close to the kernel lattice, produces instances of the UniqueSVP problem and determines the difficulty of the secret key recovery problem in PASS. The average case produces instances of ApproxSVP which will inform our discussion of the signature forgery problem.

UniqueSVP is the problem of finding a shortest vector in a lattice that is known to have a significant gap between the lengths of its first and second successive minima. Such is the case ${ }^{3}$ in the lattices $\mathcal{L}_{\boldsymbol{f}^{\prime}}^{\text {SVP }}$, as the the secret key, $\boldsymbol{f}$, has an expected norm of $\sqrt{2 N / 3}$ and $[\boldsymbol{f}, 1]^{T} \in \mathcal{L}_{\boldsymbol{f}^{\prime}}^{\text {SVP. }}$.

Lattice reduction algorithms can be ranked according to the so-called Hermite factor that they achieve. Algorithms that achieve Hermite factor $\gamma$ can be expected to find the shortest vector in a lattice when the UniqueSVP-gap, $\lambda_{2}(\mathcal{L}) / \lambda_{1}(\mathcal{L})$, is greater than a constant fraction of $\gamma$. This behavior was first examined by Gama and Nguyen, whose experiments determined that for a certain class of random lattices the constant is approximately 0.48 [5]. They exhibited classes of lattices for which the constant was smaller, but these appear to be somewhat exceptional. Ducas et al. [4] performed similar experiments on the lattices that occur in BLISS, and found the constant again to be 0.48 , and we have found the same to be true of the lattices related to PASS.

Table 2 contains estimates on the Hermite factor needed to recover PASS secret keys at several concrete parameter levels. We estimate $\lambda_{2}\left(\mathcal{L}_{\boldsymbol{f}^{\prime}}^{\text {SVP }}\right)$ by the Gaussian heuristic in the $L^{2}$ norm. This predicts that $N$ successive minima of a lattice will be tightly clustered around the radius of the smallest $N$-ball that has volume equal to the determinant of the lattice. The $q$-ary lattices, $\Lambda^{\perp}\left(\mathcal{F}_{\Omega}\right)$, have determinant $q^{t}$,

[^2]and the Gaussian heuristic therefore predicts
$$
\lambda_{2}\left(\mathcal{L}_{f^{\prime}}^{\mathrm{SVP}}\right)=\lambda_{1}\left(\Lambda^{\perp}\left(\mathcal{F}_{\Omega}\right)\right) \approx \operatorname{det}\left(\Lambda^{\perp}\left(\mathcal{F}_{\Omega}\right)\right)^{1 / N} \sqrt{\frac{N}{2 \pi \mathrm{e}}}=q^{t / N} \sqrt{\frac{N}{2 \pi \mathrm{e}}}
$$

As mentioned above, we estimate $\lambda_{1}$ as $\sqrt{2 N / 3}$, the length of the secret key. This gives us a UniqueSVP-gap, $\lambda_{2} / \lambda_{1} \approx q^{t / N} \sqrt{3 /(4 \pi \mathbf{e})}$. Incorporating the constant 0.48 adjustment, we find that lattice reduction algorithms must achieve Hermite factor

$$
\begin{equation*}
\gamma=0.62 \cdot q^{t / N} \tag{1}
\end{equation*}
$$

in order to recover PASS secret keys.

The analysis for forgery attacks is very similar, only now the target $\left.\widehat{\boldsymbol{z}}\right|_{\Omega}$ will lie in an essentially random coset of $\mathbb{Z}_{q}^{N} / \Lambda^{\perp}\left(\mathcal{F}_{\Omega}\right)$. The relevant problem is now ApproxSVP ${ }_{\alpha}$ the problem of finding a short vector that is more than $\alpha$ factor of being optimal, in other words a vector that is no longer than $\alpha \lambda_{1}\left(\mathcal{L}_{\boldsymbol{z}^{\prime}}^{\mathrm{SVP}}\right)$. Lattice reduction algorithms that achieve Hermite factor $\gamma$ can solve ApproxSVP with factor $\alpha=\gamma^{2}$ in the worst case. That said, $\alpha=\gamma$ seems achievable on average [5], so we use this estimate in our analysis.

PASS signatures are validated by the $L^{\infty}$ norm, but lattice reduction algorithms typically only guarantee the $L^{2}$ norm of their results. A vector of $L^{2}$ norm $\sqrt{N}$. $(k-b)$ could potentially serve as a forgery, but this is highly unlikely. We estimate the approximation factor to be the ratio of the expected length of a forgery to the volume of the lattice, which is

$$
\begin{equation*}
\alpha=\sqrt{N} \cdot V / q^{t / N} \tag{2}
\end{equation*}
$$

where $V$ is the variance of the discrete uniform distribution on $[-k+b, k-b]$.
5.2.2. Concrete performance of lattice reduction algorithms. Current folklore is that lattice reduction algorithms can achieve Hermite factor $\approx 1.01^{N}$ in reasonable time but that Hermite factor $1.005^{N}$ is completely out of reach. These are useful heuristics, but they reflect more our ignorance about the concrete performance of lattice reduction and enumeration algorithms than they do our knowledge. Unfortunately, it seems that we know far too little about how these algorithms perform in high dimension to give precise "bit-security" estimates. We can, however, roughly determine which of the currently available lattice reduction algorithms might be useful for attacking PASS.

Experiments by Schneider and Buchmann [21] indicate that the Hermite factor reachable by BKZ with blocksize $\beta$ is approximately:

$$
1.01655-0.000196185 \cdot \beta,
$$

which for Hermite factors relevant to our parameter sets yields:

| Blocksize $(\beta)$ | 15 | 30 | 40 | 55 |
| :--- | :---: | :---: | :---: | :---: |
| Root Hermite factor | 1.0136 | 1.0107 | 1.0087 | 1.0058 |

Table 2 lists several PASS parameter sets, the line labeled "Lattice security factor" represents our best guess as to the Hermite factor needed to launch either a key recovery or forgery attack (whichever is easier). We expect that our toy parameter set, $N=433$, could be defeated by running BKZ-15 to completion.

Although we do not have a good estimate on how long this would take, it should be possible with current technology.

Our other parameter sets should be significantly more difficult to attack. While Hermite factor $1.01^{N}$ is nominally within reach of today's technology, this has only been verified in relatively small dimensions. We know very little about how the algorithms will perform in dimension 577. Key recovery attacks on this parameter set should be possible with BKZ-30, but other approaches are likely needed to make the attack practical.

Chen and Nguyen have had impressive success with their BKZ-2.0 algorithm [3], which combines extreme pruning, developed in [6], with an early termination procedure, theoretically justified by [11]. BKZ-2.0 runs BKZ at phenomenally high blocksizes for a small number of rounds under the, experimentally justified, belief that most of the progress of BKZ is made in the early rounds. It is difficult to extrapolate security estimates from the results published thus far on BKZ-2.0's performance, but it would appear that our 577,769 , and 1153 parameter sets could be within reach of terminated BKZ-75, 122, and 229 respectively.

For $N=577$, our experiments with a BKZ-2.0 simulator similar to that presented in [3] indicate that 56 rounds of BKZ-75 would be sufficient to reach root hermite factor 1.0106 ; for $N=769,47$ rounds of BKZ-122 would suffice to reach 1.0084; and for $N=1153$, 42 rounds of BKZ-229 would reach 1.0058.

Following the analysis of [3], we expecte enumeration to be the most expensive subrouting of BKZ-2.0. Each round consists of approximately $N$ enumerations, and the cost of each enumeration depends on the the number of nodes visited in the enumeration tree. The estimated bit security is

$$
\log _{2}(N \cdot \text { rounds })+\log _{2}(\text { nodes per enumeration })+\log _{2}(\text { cost per node })
$$

Using number-of-node and cost-per-node estimates from [3], we have that the estimated security of our $N=769$ parameter is $\log _{2}(769 \cdot 47)+53+\log _{2}(200) \approx 76$ bits.

For $N=1153$, a single enumeration in BKZ-229 is expected to take over $2^{130}$ time, which is greater than the expected time for a quantum attack on the hash function.

## 6. Reference Implementation

We have created a reference implementation of PASS in C and made it available ${ }^{4}$ under the GNU General Public License. Table 3 gives some idea of the performance of PASS relative to the recent proposal of Ducas et al. (BLISS [4]) and to RSA and ECDSA. BLISS was tested using the June 13, 2013 version ${ }^{5}$. The implementations of RSA and ECDSA are from OpenSSL 1.0.1e. All benchmarks were run on a single 2.8 GHz core of an Intel Core i7-2640M with hyper threading and turbo boost disabled. We make no claims as to the accuracy of these benchmarks - the timing methods used internally by the three libraries tested are incommensurate and many variables have been left uncontrolled. However, we do feel that these preliminary performance estimates are worth reporting, as they indicate that the schemes are competitive with each other and that further comparisons would be interesting.

[^3]6.1. Performance considerations. The two most computationally intensive parts of PASS are the number theoretic transforms (NTT) used to compute $\mathcal{F}_{\Omega}$, and the sparse cyclic convolution used in computing $\boldsymbol{z}=\boldsymbol{f} \star \boldsymbol{c}+\boldsymbol{y}$. To compute $\mathcal{F}_{\Omega}$ we use Rader's algorithm to decompose the prime length NTT into cyclic convolution of length $N-1$. We compute the resulting convolution as a pair of Fourier transforms over $\mathbb{C}$ using version 3.3.3 of FFTW. For all of the parameter sets presented above we have chosen chosen $N$ to be a Pierpont prime (a prime of the form $2^{u} \cdot 3^{v}+1$ ) as these yield very fast Fourier transform algorithms. Fermat primes $\left(2^{u}+1\right)$ would yield a faster transforms, but there are no Fermat primes in our preferred parameter range.

We have made little effort to optimize the computation of sparse convolutions, and these often dominate the running time of the signing process.
6.2. Concrete instantiations of public functions. Our reference implementation uses sha512 to instantiate Hash for all parameter sets. The input passed to sha512 is the concatenation of the low order byte of each coefficient of $\left.\widehat{\boldsymbol{y}}\right|_{\Omega}$ followed by the sha512 digest of $\mu$.

$$
\operatorname{Hash}(\widehat{\boldsymbol{y}}, \mu)=\operatorname{sha512}\left(\text { lowbyte }\left(\widehat{\boldsymbol{y}}_{0}\right)|\ldots| \operatorname{lowbyte}\left(\widehat{\boldsymbol{y}}_{t-1}\right) \mid \operatorname{sha512}(\mu)\right)
$$

We have not implemented any message padding.
Our instantiation of FormatC sets aside the first 64 bits $^{6}$ of $h_{0}=\operatorname{Hash}\left(\left.\widehat{\boldsymbol{y}}\right|_{\Omega}, \mu\right)$ to use as signs of the nonzero coefficients of $\boldsymbol{c}$. The remaining bits of $h_{0}$ are used, 16 at a time, in a rejection sampling procedure to generate uniform random values in the interval $[0, N-1]$. Each such value becomes the index of a non-zero coefficient of $\boldsymbol{c}$. If the pool of bits is ever exhausted, the process continues on $h_{i}=\operatorname{sha512}\left(h_{i-1}\right)$.

The random coefficients of $\boldsymbol{y}$ are generated by a rejection sampling procedure on the output of a stream cipher. Specifically we use the procedure from [10] of keying the Salsa20 stream cipher with a short seed from the Linux kernel random number generator.

[^4]| $N$ | 433 | 577 | 769 | 1153 |
| :---: | ---: | ---: | ---: | ---: |
| $q$ | 775937 | 743177 | 1047379 | 968521 |
| $g$ | 268673 | 296108 | 421722 | 56574 |
| $k$ | $2^{12}-1$ | $2^{14}-1$ | $2^{15}-1$ | $2^{15}-1$ |
| $b$ | 19 | 24 | 29 | 36 |
| $t$ | 200 | 280 | 386 | 600 |
| $\operatorname{Pr}[$ Accept $]$ | 0.78 | 0.57 | 0.49 | 0.72 |
| UniqueSVP gap | 1.0117 | 1.0093 | 1.0075 | 1.0052 |
| ApproxSVP factor | 1.0105 | 1.0101 | 1.0081 | 1.0054 |
| Lattice security factor | 1.0134 | 1.0106 | 1.0084 | 1.0058 |
| Entropy of $\boldsymbol{c}$ | 124 | 160 | 200 | 260 |
| Bit-security bound | $<62$ | $<80$ | $<100$ | $\leq 130$ |

TABLE 2. Parameter sets and security indicators. UniqueSVP gap refers to $\lambda_{2} / \lambda_{1}$ without any correction for the performance of specific lattice reduction algorithms.

| Algorithm | Parameter Set | Sign $(\mu \mathrm{s})$ | Verify $(\mu \mathrm{s})$ | Sig. (bytes) | Pub. key (bytes) |
| :---: | ---: | ---: | ---: | ---: | ---: |
| PASS | 577 | 62 | 31 | 1115 | 700 |
|  | 769 | 73 | 40 | 1578 | 965 |
|  | 1153 | 203 | 69 | 2360 | 1500 |
| BLISS | 0 | 321 | 25 | 413 | 413 |
|  | I | 164 | 44 | 700 | 875 |
|  | II | 642 | 43 | 625 | 875 |
|  | III | 270 | 45 | 750 | 875 |
|  | IV | 496 | 47 | 813 | 875 |
| RSA | 1024 | 225 | 15 | 128 | 128 |
|  | 2048 | 1591 | 50 | 256 | 256 |
|  | 4096 | 1532 | 185 | 512 | 512 |
| ECDSA | secp160r1 | 80 | 270 | 40 | 20 |
|  | nistp256 | 146 | 348 | 64 | 32 |
|  | nistp384 | 268 | 1151 | 96 | 48 |

TABLE 3. Benchmarks. Times are averages over many operations.

| Parameter Set | Sign |  | Verify |  |
| :---: | ---: | ---: | ---: | ---: |
|  | Median | Average | Median | Average |
| 577 | 121996 | 171753 | 86828 | 87031 |
| 769 | 174900 | 205456 | 120204 | 120374 |
| 1153 | 421904 | 584230 | 172428 | 172641 |

TABLE 4. Sandy Bridge cycle counts for PASS. 100k samples

## References

1. M. Ajtai, Generating hard instances of lattice problems (extended abstract), Proceedings of the twenty-eighth annual ACM symposium on Theory of computing (New York, NY, USA), STOC '96, ACM, 1996, pp. 99-108.
2. E.J. Candes, J. Romberg, and T. Tao, Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information, IEEE Transactions on Information Theory 52 (2006), no. 2, 489-509.
3. Yuanmi Chen and Phong Q Nguyen, BKZ 2.0: better lattice security estimates, Advances in Cryptology-ASIACRYPT 2011, Springer, 2011, pp. 1-20.
4. Léo Ducas, Alain Durmus, Tancrède Lepoint, and Vadim Lyubashevsky, Lattice signatures and bimodal gaussians, Advances in Cryptology - CRYPTO 2013 (Ran Canetti and Juan A. Garay, eds.), Lecture Notes in Computer Science, no. 8042, Springer Berlin Heidelberg, January 2013, pp. 40-56.
5. Nicolas Gama and Phong Q. Nguyen, Predicting lattice reduction, Advances in Cryptology EUROCRYPT 2008 (Nigel Smart, ed.), Lecture Notes in Computer Science, no. 4965, Springer Berlin Heidelberg, January 2008, pp. 31-51.
6. Nicolas Gama, Phong Q. Nguyen, and Oded Regev, Lattice enumeration using extreme pruning, Proceedings of the 29th Annual international conference on Theory and Applications of Cryptographic Techniques (Berlin, Heidelberg), EUROCRYPT'10, Springer-Verlag, 2010, pp. 257-278.
7. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan, Trapdoors for hard lattices and new cryptographic constructions, Proceedings of the 40 th annual ACM symposium on Theory of computing (New York, NY, USA), STOC '08, ACM, 2008, pp. 197-206.
8. Craig Gentry and Mike Szydlo, Cryptanalysis of the revised NTRU signature scheme, Advances in Cryptology EUROCRYPT 2002 (Lars R. Knudsen, ed.), Lecture Notes in Computer Science, no. 2332, Springer Berlin Heidelberg, January 2002, pp. 299-320 (en).
9. Tim Güneysu, Vadim Lyubashevsky, and Thomas Pöppelmann, Practical lattice-based cryptography: A signature scheme for embedded systems, Cryptographic Hardware and Embedded Systems - CHES 2012 (Emmanuel Prouff and Patrick Schaumont, eds.), Lecture Notes in Computer Science, no. 7428, Springer Berlin Heidelberg, January 2012, pp. 530-547.
10. Tim Güneysu, Tobias Oder, Thomas Pöppelmann, and Peter Schwabe, Software speed records for lattice-based signatures, Post-Quantum Cryptography (Philippe Gaborit, ed.), Lecture Notes in Computer Science, vol. 7932, Springer-Verlag Berlin Heidelberg, 2013, Document ID: d67aa537a6de60813845a45505c313, http://cryptojedi.org/papers/\#lattisigns, pp. 67-82.
11. Guillaume Hanrot, Xavier Pujol, and Damien Stehlé, Analyzing blockwise lattice algorithms using dynamical systems, Advances in Cryptology - CRYPTO 2011 (Phillip Rogaway, ed.), Lecture Notes in Computer Science, no. 6841, Springer Berlin Heidelberg, January 2011, pp. 447-464.
12. Jeffrey Hoffstein, Burton S. Jr. Kaliski, Daniel Bennett Lieman, Matthew John Barton Robshaw, and Yiqun Lisa Yin, Secure user identification based on constrained polynomials, June 2000, U.S. Classification: 713/168; 380/28; 380/30; 713/170; 713/176 International Classification: H04L 932; H04L 928; H04L 930.
13. Jeffrey Hoffstein and Joseph H. Silverman, Polynomial rings and efficient public key authentication II, Cryptography and Computational Number Theory (Kwok-Yan Lam, Igor Shparlinski, Huaxiong Wang, and Chaoping Xing, eds.), Progress in Computer Science and Applied Logic, no. 20, Birkhäuser Basel, January 2001, pp. 269-286 (en).
14. Vadim Lyubashevsky, Lattice-based identification schemes secure under active attacks, Public Key Cryptography - PKC 2008 (Ronald Cramer, ed.), Lecture Notes in Computer Science, no. 4939, Springer Berlin Heidelberg, January 2008, pp. 162-179.
15. , Fiat-shamir with aborts: Applications to lattice and factoring-based signatures, Advances in Cryptology-ASIACRYPT 2009, Springer, 2009, pp. 598-616.
16. $\qquad$ , Lattice signatures without trapdoors, Advances in Cryptology - EUROCRYPT 2012 (David Pointcheval and Thomas Johansson, eds.), Lecture Notes in Computer Science, no. 7237, Springer Berlin Heidelberg, January 2012, pp. 738-755.
17. Vadim Lyubashevsky and Daniele Micciancio, Generalized compact knapsacks are collision resistant, Automata, Languages and Programming (Michele Bugliesi, Bart Preneel, Vladimiro

Sassone, and Ingo Wegener, eds.), Lecture Notes in Computer Science, no. 4052, Springer Berlin Heidelberg, January 2006, pp. 144-155.
18. Daniele Micciancio and Oded Regev, Lattice-based cryptography, Post-Quantum Cryptography (Daniel J. Bernstein, Johannes Buchmann, and Erik Dahmen, eds.), Springer Berlin Heidelberg, January 2009, pp. 147-191.
19. Phong Nguyen, Cryptanalysis of the goldreich-goldwasser-halevi cryptosystem from crypto 97, Advances in Cryptology CRYPTO 99 (Michael Wiener, ed.), Lecture Notes in Computer Science, no. 1666, Springer Berlin Heidelberg, January 1999, pp. 288-304 (en).
20. Chris Peikert and Alon Rosen, Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices, Theory of Cryptography (Shai Halevi and Tal Rabin, eds.), Lecture Notes in Computer Science, no. 3876, Springer Berlin Heidelberg, January 2006, pp. 145-166.
21. Michael Schneider and Johannes Buchmann, Extended lattice reduction experiments using the BKZ algorithm, Sicherheit 2010170 (2010), 241-252.
E-mail address: jhoff@math.brown.edu
Department of Mathematics, Brown University, Providence, RI, 02912
E-mail address: jpipher@math.brown.edu
Department of Mathematics, Brown University, Providence, RI, 02912
E-mail address: jschanck@securityinnovation.com
Security Innovation, Wilmington, MA 01887
E-mail address: jhs@math.brown.edu
Department of Mathematics, Brown University, Providence, RI, 02912
E-mail address: wwhyte@securityinnovation.com
Security Innovation, Wilmington, MA 01887


[^0]:    ${ }^{1}$ The original PASS protocol used the centered $L^{2}$ norm - the $L^{2}$ norm about the mean of the vector. This norm can be seen to enjoy the above quasi-multiplicative property for independent random polynomials by considering the product in the complex Fourier domain, noting that the centering operation has the effect of zeroing the constant terms, and by applying Parseval's theorem.

[^1]:    ${ }^{2}$ Note that the generation of $\boldsymbol{y}$ and the computation of $\left.\widehat{\boldsymbol{y}}\right|_{\Omega}$ can both be done offline, oblivious to the message to be signed.

[^2]:    ${ }^{3}$ Curiously, the fact that the kernel lattice always contains the exceptionally short vector $[1,1, \ldots, 1]$ seems to have no impact here.

[^3]:    ${ }^{4}$ The code is available at https://github.com/NTRUOpenSourceProject/ntru-crypto
    ${ }^{5}$ http://bliss.di.ens.fr/

[^4]:    ${ }^{6}$ The number 64 was chosen, arbitrarily, as a reasonable upper bound on $b$

