

MORE EFFICIENT POST-QUANTUM SIGNATURES

5 CROSS REFERENCE(S) TO RELATED APPLICATIONS

[0001] This application is an international application of and claims the benefit of the filing date of U.S. Provisional Application No. 62/803,325, filed on February 8, 2019, which is herein incorporated by reference in its entirety.

BACKGROUND

10 [0002] Many cryptosystems today, e.g. RSA encryption, are based on problems such as prime factorization or modular exponentiation. These problems are hard for conventional computers, but may not be for quantum computers. Thus new advances in quantum computing may render cryptosystems based on those problems insecure. There is a need for new cryptosystems based on problems that are hard for both conventional computers and
15 quantum computers.

[0003] Lattice-based problems can be used as the basis for cryptosystems. One example of a lattice problem may be the short integer solution problem, which relates to finding the shortest vector in an n -dimensional lattice of points. Lattice problems can be shown to be hard for both conventional and quantum computers. However, computing the cryptographic
20 keys for lattice-based cryptosystems can be computationally expensive, and can result in large keys that also make it difficult to sign and verify messages. Therefore, there is a need to make lattice-based cryptography more efficient.

[0004] Embodiments of the present disclosure address these and other problems, individually and collectively.

25 BRIEF SUMMARY

[0005] Some embodiments of the present disclosure are directed to methods of generating a lattice-based cryptographic verification key and an associated trapdoor matrix. A generating device can generate a gadget matrix $\mathbf{G}$. The generating device can then drop the first l columns of the gadget matrix $\mathbf{G}$ to generate a reduced gadget matrix $\mathbf{F}$. Then, the generating
30 device can sample a distribution to form a trapdoor matrix $\mathbf{R}$. The trapdoor matrix $\mathbf{R}$ can function as a secret key for generating a digital signature, and may be an approximate trapdoor. The generating device can then use the reduced gadget matrix $\mathbf{F}$ and the trapdoor

matrix **R** to generate a verification matrix **A**. The generating device can transmit the verification matrix **A** to one or more verification devices. The one or more verification devices can use the verification matrix **A** to verify a digital signature **x** generated by a sending device to sign a message *m*, if the digital signature **x** was generated using the trapdoor matrix **R**.

[0006] Other embodiments of the present disclosure are directed to methods of generating a lattice-based signature. A sending device can receive a message *m*. The message may be a hash $h(m)$, or the sending device may hash the message *m* to generate the hash $h(m)$. The sending device may also store a verification matrix **A** and a trapdoor matrix **R**. The sending device may generate a signature vector **x** using the verification matrix **A** and the trapdoor matrix **R**. The signature vector **x** may be generated such that, for some modulus *q* and an error vector **e**, the relationship $\mathbf{Ax} = h(m) + \mathbf{e} \bmod q$ is satisfied. The modulus *q* can be selected to give a desired level of data security (e.g., selecting $q = 2^{24}$ may give 100-bit security). The signing device may transmit the message *m* and the signature vector **x** to a verification device. The verification device may have previously received the verification matrix **A**, or may receive the verification matrix **A** from the sending device, and may use the verification matrix **A** to verify that the signature vector **x** was generated by the sending device. This may allow the verification device to verify that the message *m* is from the sending device as well.

[0007] Other embodiments of the present disclosure are directed to methods of verifying a lattice-based signature. A verification device can store a verification matrix **A**. The verification device can also receive a message *m* and a signature vector **x** from a sending device. The message may be a hash $h(m)$, or the verification device may hash the message *m* to generate the hash $h(m)$. The verification device may compute an error vector **e** and by solving $\mathbf{Ax} = h(m) + \mathbf{e} \bmod q$ for some modulus *q* and quantify the size of the error vector **e**. The verification device can compare the error vector **e** to a threshold and verify the signature vector **x** if the error vector **e** is less than the threshold. For example, if the length of the error vector **e** is less than the dimension of the hash $h(m)$.

[0008] Other embodiments of the present disclosure are directed to systems and computer readable media associated with the above-described methods.

[0009] These and other embodiments of the invention are described in further detail below with reference to the Figures and the Detailed Description.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0010] FIG. 1 shows a block diagram of a system according to some embodiments.
- [0011] FIG. 2 shows a general flow diagram for lattice-based verification matrix and trapdoor generation according to some embodiments.
- 5 [0012] FIG. 3 shows a general flow diagram for lattice-based signature generation according to some embodiments.
- [0013] FIG. 4 shows a general flow diagram for verification of lattice-based signatures according to some embodiments.
- [0014] FIG. 5A-B show a simplified diagram of signature verification with exact trapdoors
10 and approximate trapdoors in lattice-based signing systems.
- [0015] FIG. 6 shows a summary table of concrete parameters according to some embodiments.
- [0016] FIG. 7 shows algorithms for verification matrix generation and signature vector generation according to some embodiments.
- 15 [0017] FIG. 8 shows another table of concrete parameters according to some embodiments.
- [0018] FIG. 9 shows another table of concrete parameters according to some embodiments.
- [0019] FIG. 10 shows a block diagram of an example computer system usable with systems and methods according to some embodiments.

20

DETAILED DESCRIPTION

- [0020] Lattice-based cryptography, cryptography based on lattice problems, can provide cryptographic solutions that are robust especially with the development of quantum computing that can reduce the effectiveness of more traditional cryptosystems. Lattice problems include, for example, the inhomogeneous short integer solution (ISIS) problem,
25 which relates to finding the shortest vector in an n -dimensional lattice of points. The inhomogeneous short integer solution (ISIS) problem asks to find a vector $\mathbf{x}$ that satisfies the relationship $\mathbf{A} \cdot \mathbf{x} = \mathbf{y} \pmod{q}$. In this example, the vector $\mathbf{y}$ may be a message, or a hash of the message, and the vector $\mathbf{x}$ may be a signature associated with the message. Lattice problems can be computationally difficult, even with quantum computing, making them well
30 suited for cryptographic applications. Solving lattice problems can be made easier with a

trapdoor matrix, which can be a matrix that makes a hard problem (such as the ISIS problem) easy, or computable within a reasonable time. In the example of the ISIS problem, a trapdoor matrix can be used to compute the vector $\mathbf{x}$ without needing to solve the equation. Thus in a cryptographic application, a trapdoor can function as a secret key.

- 5 **[0021]** Different constructions of a trapdoor matrix can have different impacts on the size and usefulness of the lattice problem as a cryptosystem. Trapdoor matrices that solve exactly solve the ISIS problem may result in the public key (the verification matrix $\mathbf{A}$) and the private key (the trapdoor matrix) being very large, and thus not very practical for signing messages and transmitting signatures.
- 10 **[0022]** Instead, embodiments can use an approximate trapdoor which relaxes the solution to the ISIS problems to find an approximate solution instead of an exact solution, i.e., $\mathbf{A} \cdot \mathbf{x} = \mathbf{y} + \mathbf{e} \pmod{q}$, with an error vector $\mathbf{e}$. The signature can still be verified if the error vector is small. In particular, a gadget matrix (a matrix where columns comprise ascending powers of a small base) can be used to find an exact trapdoor matrix, and by removing a few columns of
- 15 the gadget matrix, a reduced gadget matrix can be used to find an approximate trapdoor matrix. With the approximate solution, the signature $\mathbf{x}$, the trapdoor matrix and the verification matrix $\mathbf{A}$ can be smaller and thus more practical.

- [0023]** In embodiments, a generating device can compute a verification matrix $\mathbf{A}$ and a trapdoor matrix $\mathbf{R}$. The verification matrix $\mathbf{A}$, functioning as a public key, can be distributed
- 20 or even made public. In particular, the verification matrix $\mathbf{A}$ can be sent to a sending device, that will send signed messages, and a verification device, that will verify the signature on the messages to authenticate the sending device. The trapdoor matrix $\mathbf{R}$, functioning as a private key, can be sent securely to the sending device, enabling the sending device to generate a signature.

25 **I. INTRODUCTION**

- [0024]** In the past two decades, lattice-based cryptography has emerged as an active area of research. It can enable both advanced cryptographic capabilities, such as fully homomorphic encryption [Gen09]; and practical post-quantum secure public-key encryptions and signatures, as observed in the ongoing NIST post-quantum cryptography (PQC)
- 30 standardization procedure [AAAS+19]. A large fraction of the lattice-based cryptosystems uses lattice trapdoors. Those cryptosystems include basic primitives like public-key

encryption and signature schemes [GGH97, HPS98, HHP+03, GPV08], as well as advanced primitives such as identity-based encryption [GPV08, ABB10, CHKP12], attribute-based encryption [GVW13], and graded encodings [GGH15].

[0025] In this work, we focus on the trapdoor for the lattice-based one-way function defined by Ajtai [Ajt96], and its application in digital signatures [GPV08]. Given a wide, random matrix $\mathbf{A}$, and a target vector $\mathbf{y}$. The inhomogeneous short integer solution (ISIS) problem asks to find a short vector $\mathbf{x}$ as a preimage of $\mathbf{y}$, i.e. $\mathbf{A} \cdot \mathbf{x} = \mathbf{y} \pmod{q}$:

[0026] Without the trapdoor, finding a short preimage is proven to be as hard as solving certain lattice problems in the worst case [Ajt96]. A trapdoor for the matrix $\mathbf{A}$, on the other hand, allows its owner to efficiently produce a short preimage. An explicit construction of the trapdoor for Ajtai's function was first given in [Ajt99] and later simplified by [AP11, MP12].

[0027] Towards the proper use of lattice trapdoors in cryptography, what really gives the trapdoor a punch is the work of Gentry, Peikert and Vaikuntanathan [GPV08]. They show how to sample a short preimage from a *trapdoor-independent distribution*, instead of a distribution which may leak information about the trapdoor (as observed by the attacks [GS02, NR06] on the initial attempts of building lattice-based signatures [GGH97, HHP+03]). The trapdoor-independent preimage sampling algorithm allows [GPV08] to securely build a hash-and-sign signature as follows: Let the matrix $\mathbf{A}$ be the public verification key, the trapdoor of $\mathbf{A}$ be the secret signing key. To sign a message m , first hash it to a vector $\mathbf{y}$, then use the trapdoor to sample a short preimage $\mathbf{x}$ as the signature. The secret signing key is guaranteed to be hidden from the signatures, since the signatures are produced from a trapdoor-independent distribution.

[0028] Despite its elegant design, the hash-and-sign signature based on Ajtai's function suffers from practical inefficiency due to its large key size and signature size. Indeed, all the three lattice-based signature candidates that enter the second round of NIST PQC standardization [AAAS+19] are built from two alternative approaches—Falcon [FHK+18] is based on the hash-and-sign paradigm over NTRU lattices; Dilithium [DKL+18] and qTESLA [ABB+19] are based on the rejection sampling approach [Lyu12, BG14]. The suggested parameters for the three candidates lead to competitive performance measures. For example, for 128-bit security, the sizes of the public keys & signatures for all the three candidates are below 5 kB & 4 kB (respectively). By contrast, for the hash-and-sign signature based on

Ajtai's function, the sizes of the public keys & signatures are more than 35 kB & 25 kB according to the implementation benchmarks of [BB13, BFRS18, GPR+18].

[0029] We can define a relaxed notion of lattice trapdoors called approximate trapdoors, which can be used to solve the ISIS problem approximately instead of exactly. The primary motivation is to improve the efficiency of hash-and-sign signatures based on Ajtai's one-way function. With a relaxation of the correctness requirement, it is possible to generate smaller public matrices, trapdoors, and preimages for Ajtai's function, which translate to smaller public-keys, secret-keys, and signatures for the hash-and-sign signature scheme.

[0030] Embodiments of the disclosure show that the gadget trapdoor proposed by Micciancio and Peikert [MP12] can be modified to an approximate trapdoor. In particular, we show how to use the approximate gadget trapdoor to sample preimages from a trapdoor-independent distribution. The analysis of the distribution can use linear transformations of discrete Gaussians on lattices.

[0031] Our approximate gadget trapdoor can be used together with existing optimization techniques, such as using the Hermite normal form and using a bigger base in the gadget, to improve the concrete performance of the hash-and-sign signature in the random oracle model under RingLWE and RingSIS assumptions. We show that the sizes of the public key & signature can be reduced to 5 kB & 4.45 kB for an estimation of 100-bit security, and 11.25 kB & 9.38 kB for an estimation of 192-bit security. Those are much closer to the sizes of the signatures based on the rejection sampling approach [Lyu12, BG14, DKL+18, ABB+19].

A. Example cryptographic context

[0032] FIG. 1 shows a sequence diagram of a system 100 according to embodiments. System 100 may comprise a generating device 105, a sending device 115, and a verification device 125. The system 100 may also comprise an attacker 135. System 100 may be a system that uses signature generation and verification according to embodiments of the disclosure. For example, in one embodiment, sending device 115 may be a device of a user sending a message to transfer funds from one account to another account. Verification device 125 may be an issuer verifying the message to prevent a fraudulent transfer. In another embodiment, sending device 115 may be a device sending a message to request access to a secure database. Verification device 125 may be an access gateway that controls access to the secure database and may verify the message prior to granting access. In some embodiments, generating

device 105 and sending device 115 may be the same device. In other embodiments, generating device 105 and verification device 125 may be the same device.

[0033] In step 102, the generating device 105 can generate a verification key vk and a secret key sk . For example, the verification key vk may be a verification matrix $\mathbf{A}$ and the
 5 secret key sk may be a trapdoor matrix $\mathbf{R}$. The generating device 105 can then publicly transmit the verification key vk to the verification device 125. An attacker 135 may also receive the verification key vk . The generating device 105 can also privately send the secret key sk to the sending device 115.

[0034] In step 104, the sending device 115 can receive a message m . The sending device
 10 115 can then hash the message with a hashing function to generate a hashed message $h(m)$. Alternatively, the sending device 115 may generate the message m or may receive the hash $h(m)$.

[0035] In step 106, the sending device 115 can generate a signature sig based on the hash $h(m)$ and the secret key sk . For example, the sending device 115 can generate the signature
 15 sig compute a function of the hash $h(m)$ and the secret key sk . The sending device 115 may generate the signature sig such that the product of the verification key vk and the signature sig is the hashed message $h(m)$.

[0036] In step 108, the sending device 115 can send the message m and the signature sig to a verification device 125. The attacker 135 may be able to intercept the message m and the
 20 signature sig as they are sent.

[0037] In step 110, the attacker 135 may attempt to modify the message m to generate a fraudulent message m' . The attacker 135 may also attempt to modify the signature sig to generate a fraudulent signature sig' . However, because the attacker 135 does not have access to the secret key sk , they may not be able to generate a valid signature. Therefore, the
 25 fraudulent signature sig' may not be a valid signature. The attacker 135 may then send the fraudulent message m' and the fraudulent signature sig' to the verification device 125.

[0038] In step 112, the verification device 125 can attempt to verify the message m from the sending device 115 and the fraudulent message m' from the attacker 135. For example, the verification device 125 may multiply the signature sig and the verification key vk , and
 30 compare the result to the hashed message $h(m)$. If the result is equal to the hash $h(m)$, or within an accepted threshold, then the message m can be verified. If the result is not equal to

the hashed message $h(m)$, or within an accepted threshold, then the message m may not be verified. More generally, verifying may involve a function of the message, the hashed message, the signature, and/or the verification key. Signatures that were not generated with the secret key sk may not give valid result and may not be verified. The verification device
 5 125 can thus differentiate between authentic messages sent by the sending device 115 and fraudulent messages sent by the attacker 135.

B. Generation of verification matrix

[0039] A verification matrix can be generated that is used as a public key to verify a signature. A trapdoor matrix (as a private key) can also be generated, which can be used in
 10 generating the verification matrix and in generating the signature.

[0040] FIG. 2 shows a flow diagram of a key generation process according to embodiments of the disclosure. In some embodiments, key generation may be done by a generating device (e.g., 105 in FIG. 1), which may also be a sending device. The key may be a lattice-based cryptographic verification key.

15 [0041] In step 202, the generating device may generate a gadget matrix $\mathbf{G}$. The gadget matrix $\mathbf{G}$ may be a matrix where all entries are powers of a small integer b , e.g., 2, 3, 4, or 5. For example, the entries of the gadget matrix $\mathbf{G}$ may be ascending powers of 2. The gadget matrix G may have dimensions $n \times \log q$. n may be a security parameter, such as 64, 128, 256, and 512. q may be set to $n\sqrt{n}$. In some embodiments, the modulus q may be between
 20 2^{16} to 2^{24} .

[0042] In step 204, the generating device may reduce the dimensions of the gadget matrix $\mathbf{G}$ by dropping selected columns to generate a reduced gadget matrix $\mathbf{F}$. For example, the generating device may drop the first l columns of the gadget matrix $\mathbf{G}$. The parameter l may be chosen, for example, to be between 0 and $\frac{\log_b q}{2}$, or between 0 and $\log q$. l may be
 25 experimentally determined to provide a desired security level and key size. A larger value of l (resulting a smaller reduced gadget matrix $\mathbf{F}$) may result in smaller keys, but may also be less secure. A larger value of l (resulting in a larger reduced gadget matrix $\mathbf{F}$) may result in greater security but with larger keys. For example, for a security parameter $n = 512$ and base $b = 2$, the modulus may be set to $q = 2^{24}$. Then l may be chosen from the range between 0
 30 and 12.

[0043] In step 206, the generating device may sample a distribution to form a trapdoor matrix $\mathbf{R}$. For example, each element of the trapdoor matrix $\mathbf{R}$ may be sampled independently from a Gaussian distribution with a small standard deviation. The trapdoor matrix $\mathbf{R}$ can also be used to calculate a small matrix $\mathbf{D}$, where $\mathbf{D} = \begin{bmatrix} \mathbf{R} \\ \mathbf{I} \end{bmatrix}$. $\mathbf{D}$ may thus be the trapdoor matrix $\mathbf{R}$ vertically augmented with an identity matrix of the same dimensions. A small matrix may be a matrix where each entry is less than n . The trapdoor matrix $\mathbf{R}$ can function as a secret key for generating a digital signature.

[0044] In step 208, the generating device may generate a verification matrix $\mathbf{A}$. The generating device may first generate a matrix $\bar{\mathbf{A}}$, based on a uniformly random matrix $\hat{\mathbf{A}}$. Each element in the uniformly random matrix $\hat{\mathbf{A}}$ may be chosen with equal probability. $\mathbf{A}$ can then be calculated from $\mathbf{F}$, $\mathbf{R}$, and $\bar{\mathbf{A}}$ as $\mathbf{A} := [\bar{\mathbf{A}} \mid \mathbf{F} - \bar{\mathbf{A}}\mathbf{R}]$. Thus, the verification matrix $\mathbf{A}$ is the matrix $\bar{\mathbf{A}}$ augmented with the matrix resulting from the expression $\mathbf{F} - \bar{\mathbf{A}}\mathbf{R}$.

[0045] In step 210, the generating device may transmit the verification matrix $\mathbf{A}$ to a sending device (e.g., 115 in FIG. 1). The verification matrix $\mathbf{A}$ may serve as a publicly known verification key. If the generating device is not the sending device, then the generating device may also transmit the trapdoor matrix $\mathbf{R}$ and/or the small matrix $\mathbf{D}$ to the sending device. The trapdoor matrix $\mathbf{R}$ may be sent, for example, by a secure channel or using a secret sharing algorithm. A secure channel may include, for example, an Internet communication channel secured with SSH encryption or manually writing the trapdoor matrix $\mathbf{R}$ and sending to the sending device.

C. Signature generation

[0046] FIG. 3 shows a flow diagram of signature generation process according to embodiments of the disclosure. The signature may be a lattice-based signature vector. In some embodiments, signature generation may be done by a sending device (e.g., sending device 115 of FIG. 1).

[0047] In step 302, the sending device can receive a message m . In some embodiments, the sending device may have generated the message m . The message m may be, for example, an authorization request message or a payment transaction message. The message m may be hashed. The sending device can hash the message m to generate a hashed message $h(m)$. For

example, the sending device may take the modulus of the message m . The message m and the hash $h(m)$ may be vectors.

[0048] In step 304, the sending device can store a verification matrix $\mathbf{A}$ and a trapdoor matrix $\mathbf{R}$. The verification matrix $\mathbf{A}$ and the trapdoor matrix $\mathbf{R}$ may be received from a generating device. In some embodiments, the sending device may generate the verification matrix $\mathbf{A}$ and the trapdoor matrix $\mathbf{R}$. The sending device may receive a small matrix $\mathbf{D}$ instead of, or in addition to, the trapdoor matrix $\mathbf{R}$. The trapdoor matrix $\mathbf{R}$ and the verification matrix $\mathbf{A}$ may be received, for example, through a secure channel or using a secret sharing algorithm. A secure channel may include, for example, an Internet communication channel secured with SSH encryption.

[0049] In step 306, the sending device may generate a signature vector $\mathbf{x}$ using the hash $h(m)$, the verification matrix $\mathbf{A}$ and the trapdoor matrix $\mathbf{R}$. The sending device may first sample a probability distribution (e.g., a Gaussian distribution) to form a perturbation vector $\mathbf{p}$. Each element in the perturbation vector $\mathbf{p}$ can be independently sampled from the probability distribution. The sending device may then compute a lattice vector $\mathbf{v} = h(m) - \mathbf{A}\mathbf{p}$. The lattice vector $\mathbf{v}$ may then be used to sample an intermediate vector $\mathbf{z}$ from an lattice Gaussian distribution based on a lattice formed from the lattice vector $\mathbf{v}$. A lattice Gaussian distribution may be a discrete Gaussian distribution over a lattice, in particular the lattice defined by lattice vector $\mathbf{v}$. Then, the sending device can then generate the signature vector $\mathbf{x}$ from a product of the intermediate vector $\mathbf{z}$ and the trapdoor matrix $\mathbf{R}$ as $\mathbf{x} = \mathbf{p} + \begin{bmatrix} \mathbf{R} \\ \mathbf{I} \end{bmatrix} \mathbf{z}$. The perturbation $\mathbf{p}$ may prevent an attacker from determining the trapdoor matrix $\mathbf{R}$ based on a message-signature pair. As the perturbation $\mathbf{p}$ can be sampled from a distribution that does not depend on the message m or the trapdoor matrix $\mathbf{R}$, it can mask the trapdoor matrix $\mathbf{R}$ in the signature vector $\mathbf{x}$.

[0050] In step 308, the sending device may transmit the message m and the signature vector $\mathbf{x}$. These can be transmitted to a verification device. The sending device may also transmit the verification matrix $\mathbf{A}$ to the verification device. The verification device can then verify the message m . The verification device can multiply the verification matrix $\mathbf{A}$ by the signature vector $\mathbf{x}$, and if the result matches the hashed message with a small error then the message m can be verified. The verification device can also verify that the signature vector $\mathbf{x}$ was generated by the sending device.

D. Signature verification

[0051] FIG. 4 shows a flow diagram of verifying a signature vector $\mathbf{x}$ with a verification matrix $\mathbf{A}$. In some embodiments, verification may be done by a verification device (e.g., verification device 125 in FIG. 1).

5 [0052] In step 402, the verification device can store a verification matrix $\mathbf{A}$. The verification matrix $\mathbf{A}$ may be received from a generating device. Alternatively, the verification device may receive the verification matrix $\mathbf{A}$ from a sending device, or may be publically available. The verification matrix $\mathbf{A}$ may alternatively be generated by the verification device.

10 [0053] In step 404, the verification device can receive a message m and a signature vector $\mathbf{x}$ from a sending device. The verification device may receive a hashed message $h(m)$ in addition to, or instead of the message m . The verification device may hash the message m to generate the hashed message $h(m)$. The message m and the hash $h(m)$ may be vectors.

[0054] In step 406, the verification device can use the verification matrix $\mathbf{A}$ to verify the
 15 signature vector $\mathbf{x}$. In order to verify the signature vector $\mathbf{x}$, the verification device can compute an error vector $\mathbf{e}$ by solving the equation $\mathbf{Ax} = h(m) + \mathbf{e} \bmod q$, for some modulus q . In particular, because $\mathbf{Ax} - h(m) = \mathbf{e} \bmod q$, the verification device can compute $\mathbf{Ax} - h(m)$ with the verification matrix $\mathbf{A}$, the signature vector $\mathbf{x}$, and the hash $h(m)$ to determine the error vector $\mathbf{e}$. The modulus q may be determined as $n\sqrt{n}$ for a
 20 security parameter n . For example, for a parameter $n = 128$, q may be 2^{23} . In some embodiments, the modulus q may be selected from a range between 2^{16} and 2^{24} . Because the signature vector $\mathbf{x}$ and verification matrix $\mathbf{A}$ are generated with an approximate trapdoor, the product may not be exactly the hash $h(m)$, but instead may differ by the error vector $\mathbf{e}$. If the signature vector $\mathbf{x}$ had been generated with an exact trapdoor, the error vector $\mathbf{e}$ would be 0.
 25 The verification device can then quantify the size of the error vector $\mathbf{e}$, such as by finding the Euclidean length of the error vector $\mathbf{e}$, and compare the size to a threshold. In other embodiments, the size of the error vector $\mathbf{e}$ can be quantified using a different l_p -norm, or an infinity norm.

[0055] If the error vector $\mathbf{e}$ is less than the threshold then the signature vector $\mathbf{x}$ can be
 30 verified. In some embodiments, the threshold may be the dimension of the hash $h(m)$ (e.g., the dimension of a vector of the hash $h(m)$). In other embodiments, the threshold may be

chosen between 0 and $\frac{q}{4}$. For example, if $q = 2^{23}$, the threshold may be 2^{12} . If the error vector $\mathbf{e}$ is large, this can indicate that the message m was not signed with signature vector $\mathbf{x}$ generated from the same trapdoor as the verification matrix $\mathbf{A}$, and thus the message m is likely fraudulent or otherwise compromised.

- 5 [0056] After the signature is verified, the verification device can act upon information in the message. As one example, the message may include a request to access a secure area, such as a building or venue, or to access secure data or other resource that may reside on a networked computer. After verifying the signature and the message, the verification device can then send an indication to an access device to allow access. As another example, the
- 10 message may include a request to establish a secure connection between the sending device and the verification device. After verifying the message, the verification device can establish the secure connection, e.g., by performing a key exchange. Additionally, or alternatively, the verification device may respond to the sending device. The response may include information such as a cryptogram, an encryption key, or secure data, such as an access token. Some
- 15 embodiments may be applied to a payment processing system. For example, the message may be a payment request message, and the verification device may use the information in the payment request message as part of generating an authorization response, e.g., by transmitting the authentication (verification) to a processing network or to an authorization server that uses the authentication to prepare the authorization response.

20 E. Approximate trapdoors

[0057] FIG. 5 depicts signature verification with an exact trapdoor scheme and with an approximate trapdoor scheme. FIG. 5A shows an example calculation for verification of a signature generated with an exact trapdoor. FIG. 5B shows an example calculation for verification of a signature generated with an approximate trapdoor.

- 25 [0058] In FIG. 5A, a verification matrix $\mathbf{A}$, that has dimensions $n \times m$, can be multiplied by a signature vector $\text{Sign}(h(m), \text{sk})$, that is generated with a secret key sk and a hashed message $h(m)$. Each entry of the verification matrix may be in $\mathbb{Z}_q$. In some embodiments, n may be a security parameter, such as 128. q may be set to $n\sqrt{n}$. The signature vector $\text{Sign}(h(m), \text{sk})$ has dimension m , and is chosen such that when it is multiplied by the
- 30 verification matrix $\mathbf{A}$, it results in the hashed message. The hashed message is a vector with n dimensions. By hashing the message, the size of n and m can be fixed instead of depending

on the size of the message. Note that if m is relatively large, then both the verification matrix $\mathbf{A}$ and the signature vector $\text{Sign}(h(m), sk)$ will be large.

[0059] FIG. 5B depicts the effects of a verification matrix $\mathbf{A}$ with smaller dimension. The verification matrix $\mathbf{A}$ may now have dimensions $n \times m'$, where $m' < n + m$. Using an
 5 approximate trapdoor can allow the dimensions to be reduced. Because m' is smaller, the associated signature vector for a hashed message (still of length n) is smaller as well. However, by using an approximate trapdoor, the product $\mathbf{d}$ of the matrix multiplication may be the hashed message $h(m)$, plus an error vector $\mathbf{e}$. The verification matrix $\mathbf{A}$ can be constructed such that the error vector $\mathbf{e}$ is small, that is, the length of the error vector $\mathbf{e}$ is less
 10 than n . Therefore the condition for verification is that the product $\mathbf{d}$ is within known error bounds of the hashed message. A fraudulent signature may result in a product $\mathbf{d}$ that differs from the hashed message by a large amount, and thus the verification process can still be secure.

[0060] We can now discuss in more detail the difference between approximate trapdoors
 15 and exact trapdoors.

[0061] Given a public matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ where $m = O(n \log q)$, and a target $\mathbf{y}$. We call a vector $\mathbf{x} \in \mathbb{Z}^m$ an approximate short preimage of $\mathbf{y}$ if

$$\mathbf{A} \cdot \mathbf{x} = \mathbf{y} + \mathbf{z} \pmod{q}$$

for some $\mathbf{z} \in \mathbb{Z}^n$, and both $\mathbf{x}$ and $\mathbf{z}$ are short. An approximate trapdoor for $\mathbf{A}$ is defined to be
 20 a string that allows its owner to efficiently find an approximate short preimage given a target $\mathbf{y}$.

[0062] To make sense of the word “trapdoor”, we can argue that solving the approximate version of ISIS is hard without the trapdoor. Under proper settings of parameters, we show the approximate ISIS problem is as hard as the standard ISIS problem, or no easier than
 25 LWE. The reductions extensively use the Hermite normal form (HNF) and are pretty straightforward.

[0063] The approximate ISIS problem and the approximate trapdoor are natural generalizations of their exact variants. Indeed, both notions have been used in the literature, at least on an informal level. For example, the approximate ISIS problem was used in the

work of Bai et al. [BGLS19] to improve the combinatorial algorithms of the exact ISIS problem.

[0064] For the approximate trapdoor, an exact trapdoor of a public matrix in the HNF, say a trapdoor for $\mathbf{A} = [\mathbf{I}_n \mid \mathbf{A}']$, can be used as an approximate trapdoor for $\mathbf{A}'$. Such a method
 5 was often used in the implementation of signatures to decrease the sizes of the public key and the signature by a dimension of n . Our goal is thus to further reduce the sizes compared to the HNF approach, while preserving the quality of the trapdoor, i.e. at least not increasing the norm of the preimage.

[0065] One contribution of embodiments is to show that the gadget trapdoor (G-trapdoor)
 10 proposed by Micciancio and Peikert [MP12] can be modified to an approximate trapdoor, in a way that further reduces the sizes of the public matrix, the trapdoor, and the preimage. An aspect of the G-trapdoor is a specific “gadget” matrix of base b ,

$$\mathbf{G} := \mathbf{I}_n \otimes g^t := \mathbf{I}_n \otimes (1, b, \dots, b^{k-1}) \in \mathbb{Z}^{n \times (nk)}$$

where $k := \lceil \log_b q \rceil$. The base b can be chosen as 2 for simplicity, or a larger value in
 15 practical implementations.

[0066] Micciancio and Peikert [MP12] show how to generate a random matrix $\mathbf{A}$ together with a matrix $\mathbf{D}$ of small norm such that $\mathbf{A} \cdot \mathbf{D} = \mathbf{G} \pmod{q}$. In particular, $\mathbf{A}$ can be designed as

$$\mathbf{A} = [\bar{\mathbf{A}} \mid \mathbf{G} - \bar{\mathbf{A}}\mathbf{R}]$$

20 where $\mathbf{R}$ is a matrix with small entries and is the actual trapdoor. The matrix $\mathbf{D}$ is then equal to $\begin{bmatrix} \mathbf{R} \\ \mathbf{I}_{nk} \end{bmatrix}$. Since the kernel of the $\mathbf{G}$ matrix has a public short basis, one can first solve the ISIS problem under the public matrix $\mathbf{G}$, then use $\mathbf{D}$ to solve the ISIS problem under the public matrix $\mathbf{A}$.

[0067] We observe that if we drop a few (say l) entries corresponding to the small powers
 25 of b from the gadget matrix $\mathbf{G}$, i.e. let the following $\mathbf{F}$ matrix be a reduced gadget matrix

$$\mathbf{F} := \mathbf{I}_n \otimes f^t := \mathbf{I}_n \otimes (b^l, \dots, b^{k-1}) \in \mathbb{Z}^{n \times n(k-l)};$$

then we are still able to solve the ISIS problem with respect to the public matrix $\mathbf{F}$ up to a b^l -approximation of the solution (i.e., the norm of the error vector is proportional to b^l).

Replacing $\mathbf{G}$ by $\mathbf{F}$ in $\mathbf{A}$ gives

$$\mathbf{A} = [\bar{\mathbf{A}}|\mathbf{F} - \bar{\mathbf{A}}\mathbf{R}]$$

Then the dimensions of the trapdoor $\mathbf{R}$ and the public matrix $\mathbf{A}$ can be reduced.

[0068] Given a public matrix $\mathbf{A}$ together with its approximate G-trapdoor $\mathbf{R}$, finding an arbitrary approximate short preimage of a given target $\mathbf{u}$ is quite straightforward, but
 5 sampling the preimage from a trapdoor-independent distribution can be non-trivial. As mentioned, the ability to sample from a trapdoor-independent distribution is involved in many of the trapdoor applications including digital signatures.

[0069] We provide an algorithm that samples an approximate short preimage from a trapdoor-independent distribution. The algorithm itself can build on the perturbation-based
 10 discrete Gaussian sampler from [MP12], but the analyses of the preimage distribution from [MP12] are not easy to generalize. Our analyses of the preimage distribution and the approximation error distribution extensively use a linear transformation theorem on lattice distributions (cf. Theorem 2.7, implicitly used in [MP12, MP13, BPMW16]).

[0070] While our algorithm works for target images $\mathbf{u} \in \mathbb{Z}_q^n$ and does not cause any blow
 15 up in the standard deviation of the distribution, the analysis of trapdoor-independence only applies to a target image $\mathbf{u}$ sampled uniformly from $\mathbb{Z}_q^n$, as opposed to the analysis for the exact trapdoor in [GPV08, MP12] which is able to spell out the distribution of the preimages of all the $\mathbf{u} \in \mathbb{Z}_q^n$. To briefly explain the reason for this gap, we observe that the methods we have tried to handle all the target images require significant increases in the smoothing
 20 parameters of the lattice intersections required in the linear transformation theorem (Theorem 2.7). In other words, the norm of the resulting preimage increases significantly rendering the result meaningless.

[0071] Still, sampling the approximate preimages for uniform targets from a trapdoor-independent distribution suffices for replacing the use of exact lattice trapdoors in many
 25 cryptosystems, including the hash-and-sign signature [GPV08], IBE [GPV08, ABB10, CHKP12], ABE [GVW13], and the special purpose obfuscation schemes [GKW17, WZ17], the private-constrained PRF and the reusable garbled circuits [CC17, CVW18] built on the GGH15 graded encoding scheme [GGH15].

[0072] We now explain the efficiency gain of using our approximate trapdoor compared to
 30 the exact trapdoor and the other existing optimization techniques, with a focus on the signature application. Our goal is to set the parameters to achieve the following “win-win-

win” scenario: 1) save on the preimage size (bandwidth); 2) save on the size for the public matrix $\mathbf{A}$; and 3) retain, or even gain, concrete security, which is related to the discrete Gaussian width of the preimage and the norm of the error term.

[0073] Let us start with an understanding of the dependency of the savings on the variable l , i.e, the number of entries dropped from the gadget $\mathbf{g}$. In Table 1 we provide a comparison of the parameters between the exact G-trapdoor of [MP12] and the approximate G-trapdoor samplers in this paper. In both cases the public matrices are instantiated in the pseudorandom mode. For the approximate trapdoor, the dimension of the trapdoor decreases from nk to $n(k - l)$. The dimension m of the public matrix and the preimage decreases. The width s of the preimage distribution also decreases slightly following the decreasing of m . However, the norm of the error factor in the image grows with l . So in the concrete instantiation of the hash-and-sign signature discussed later, we need to coordinate the value of l with the norms of the preimage and the error, which can determine the security estimation together.

[0074] Our algorithm inherits the $O(\log q)$ -space, $O(n \log q)$ -time G-preimage sample subroutine from [MP12, GM18]. So the saving of space and time in the sampling of the perturbation is proportional to the saving in the dimension m .

[0075] Let us make a quick remark for the applications beyond signatures. The saving of the dimension m is of significant importance to the applications built on the GGH15 graded encoding scheme (implemented in [HHSS17, CGM+18]). In those applications, the modulus q is proportional to m^d (where $d \in \mathbb{N}$ is the number of “levels” of the graded encodings; larger d supports richer functionalities). So reducing the dimension m would dramatically reduce the overall parameters.

Parameters	Exact G-trapdoor	Approximate G-trapdoor
m	$n(2 + k)$	$n(2 + (k - l))$
σ	$\sqrt{b^2 + 1} \cdot \omega(\sqrt{\log n})$	$\sqrt{b^2 + 1} \cdot \omega(\sqrt{\log n})$
s	$C \cdot \tau \cdot (\sqrt{m} + 2\sqrt{n}) \cdot \sigma$	$C \cdot \tau \cdot (\sqrt{m} + 2\sqrt{n}) \cdot \sigma$
ν	0	$b^l \cdot \sigma$

Table 1: A brief comparison of the parameters.

[0076] The parameters in table 1 are derived under a fixed lattice dimension n , a fixed modulus $q \geq \sqrt{n}$, and a fixed base b . Let $k = \lceil \log_b q \rceil$. Let l denote the number of entries removed from $\mathbf{g}$ ($1 \leq l < k$). Then we list m as the dimension of the public matrix and the preimage; σ as the width of the gadget preimage distribution; s as the width of the final

preimage distribution (where $C > 0$ is a universal constant; τ as the width, or subgaussian parameter, of the distribution of the entries in the trapdoor); v as the length bound of the error for each entry in the image. Note that some embodiments may use δ in place of τ .

F. Example parameters for the signatures

5 **[0077]** We give a proof-of-concept implementation of the hash-and-sign signature based on our approximate trapdoor. The security is analyzed in the random oracle model, assuming the hardness of RingLWE for the pseudorandomness of the public key and RingSIS for the unforgeability of the signature. Here we provide a short summary and leave more details in Section 5.2.

10 **[0078]** FIG. 6 shows a summary table of concrete parameters according to some embodiments. Let us first remark that different implementation results of the hash-and-sign signatures [BB13, BFRS18, GPR+18] possibly use different ways of measuring sizes and security, and not all the details are recoverable from the papers. So we also include our reference implementation of the exact trapdoor as a fair comparison. For an estimation of
15 100-bit security, our reference implementation for the exact trapdoor under the modulus $q \approx 2^{24}$ and base $b = 2$ matches the parameters reported in [BB13].

[0079] We also use smaller moduli and bigger bases to reduce the size and increase the security level. The parameters in FIG. 6 suggest that for the 3 choices of q and b , using the approximate gadget trapdoor by setting $l = \lceil (\log_b q)/2 \rceil$ saves about half of the sizes in the
20 public key and signatures comparing to using the exact trapdoor, with even a slight increase in the security estimation.

[0080] Our implementation shows that the sizes of the public-key & signature can be reduced to 5 kB & 4.45 kB for an estimation of 100-bit security, and 11.25 kB & 9.38 kB for an estimation of 192-bit security. Those are much closer to the sizes of the signatures based
25 on the rejection sampling approach [Lyu12, BG14, DKL+18, ABB+19]. As a reference, the sizes of the public-key & signature for qTESLA [ABB+19] are 4.03 kB & 3.05 kB for an estimation of 128-bit security, and 8.03 kB & 6.03 kB for an estimation of 192-bit security. The sizes for Dilithium [DKL+18] are even smaller. Let us remark that our implementation has not used many low-level optimizations like Dilithium [DKL+18] and qTESLA [ABB+19].
30 So it is reasonable to expect we have more room to improve after adding lower-level

optimizations. The parameters for Falcon [FHK+18] are the smallest due to the use of NTRU lattices, so they are rather incomparable with the ones based on RingLWE.

[0081] There are many folklore optimizations regarding trapdoors for Ajtai's one-way function. We discuss the comparison and compatibility of a few of them with our
5 construction. Throughout these comparisons we are concerned with the "win-win-win" scenario mentioned in the beginning.

[0082] First is the approximate trapdoor from the HNF optimization: $\mathbf{A} = [\mathbf{I}|\mathbf{A}']$. This barely achieves the "win-win-win" scenario with a slight savings on the public key and the signature. Our construction can be used in the pseudorandom-mode of the gadget trapdoor
10 which has automatically included the HNF optimization, and saves around 50% in addition.

[0083] Our method can also be used together with any base in the gadget, including a large base of size $b \approx \sqrt{q}$ (the resulting gadget is $\mathbf{g} = [1, \sqrt{q}]$), as was used in [dPLS18] when the modulus is large. This construction suffers from a large Gaussian width ($\sqrt{q}$), which can hurts concrete security and may be infeasible in the smaller modulus regime we implement in
15 Section 5. Specifically for the smaller moduli, the signature's Gaussian width is larger than the modulus, as was confirmed both on paper and in our experiments. So we use a moderately large base b .

[0084] One may also try to construct a short integer matrix $\mathbf{S}$ for $\mathbf{A} := [(\mathbf{I}, \bar{\mathbf{A}}')|\mathbf{F} - (\mathbf{I}, \bar{\mathbf{A}}')\mathbf{R}]$ (corresponds to the pseudorandom public key in Eqn. (1)) such that $\mathbf{AS} = \mathbf{G}$, and hope this
20 view provides a better approximate trapdoor. From here, the hash-and-sign signature scheme is to return $\mathbf{Sz} + \mathbf{p}$ where $\mathbf{p}$ is a perturbation and $\mathbf{z}$ is a $\mathbf{G}$ -lattice sample. However, such a matrix $\mathbf{S}$ requires a b^l term. So this method does save on the public key, but does not improve the signature size and, most importantly, increases the Gaussian width by a factor of b^l . The increase of the width decreases the concrete security of the underlying SIS problem. In fact,
25 to achieve the same saving in the public key, one can instead directly increase the base from b to b^l in the gadget.

II. PRELIMINARIES

[0085] Before describing specific examples in detail, we can describe aspects of lattice-based cryptosystems that may be utilized in embodiments of the present disclosure.

A. Notations and terminology.

[0086] In cryptography, the security parameter (denoted as λ) is a variable that is used to parameterize the computational complexity of the cryptographic algorithm or protocol, and the adversary's probability of breaking security. An algorithm is "efficient" if it runs in
 5 (probabilistic) polynomial time over λ .

[0087] When a variable v is drawn uniformly random from the set S we denote as $v \leftarrow U(S)$. We use $\approx_s$ and $\approx_c$ as the abbreviations for statistically close and computationally indistinguishable. For two distributions D_1, D_2 over the same support $\mathcal{X}$, we denote $D_1 \stackrel{\varepsilon}{\approx} D_2$ to denote that each $x \in \mathcal{X}$ has $D_1(x) \in [1 \pm \varepsilon]D_2(x)$ and $D_2(x) \in [1 \pm \varepsilon]D_1(x)$.

10 [0088] Let $\mathbb{R}, \mathbb{Z}, \mathbb{N}$ be the set of real numbers, integers and positive integers. Denote $\mathbb{Z}/(q\mathbb{Z})$ by $\mathbb{Z}_q$. For $n \in \mathbb{N}$, $[n] := 1, \dots, n$. A vector in $\mathbb{R}$ (represented in column form by default) is written as a bold lower-case letter, e.g. $\mathbf{v}$. For a vector $\mathbf{v}$, the i^{th} component of $\mathbf{v}$ will be denoted by v_i . For an integer base $b > 1$, we call a positive integer's "b-ary" decomposition the vector $(q_0, q_1, \dots, q_{k-1}) \in \{0, \dots, b-1\}^k$ where $k := \lceil \log_b q \rceil$, and $q =$
 15 $\sum q_i b^i$.

[0089] A matrix is written as a bold capital letter, e.g. $\mathbf{A}$. The i^{th} column vector of $\mathbf{A}$ is denoted $\mathbf{a}_i$. The length of a vector is the ℓ_p -norm $\|\mathbf{v}\|_p := (\sum v_i^p)^{1/p}$, or the infinity norm given by its largest entry $\|\mathbf{v}\|_\infty := \max_i \{v_i\}$. The length of a matrix is the norm of its longest column: $\|\mathbf{A}\|_p := \max_i \|\mathbf{a}_i\|_p$. By default we use ℓ_2 -norm unless explicitly
 20 mentioned. When a vector or matrix is called "small" or "short", we refer to its norm but not its dimension, unless explicitly mentioned. The thresholds of "small" or "short" can be precisely parameterized when necessary. The notations $[\mathbf{A} \mathbf{B}]$, $[\mathbf{A}, \mathbf{B}]$, and $[\mathbf{A}|\mathbf{B}]$ can all denote that a matrix $\mathbf{A}$ is augmented or concatenated horizontally with a second matrix $\mathbf{B}$. The notation $\begin{bmatrix} \mathbf{A} \\ \mathbf{B} \end{bmatrix}$ can denote that matrices $\mathbf{A}$ and $\mathbf{B}$ are vertically concatenated, and may also
 25 be called a stack.

B. Linear Algebra

[0090] Let $\{e_i\}_{i=1}^n$ be the canonical basis $\mathbb{R}^n$, with entries $\delta(j, k)$ where $\delta(j, k) = 1$ when $j = k$ and 0 otherwise. For any set $S \subseteq \mathbb{R}^n$, its span (denoted as $\text{span}(S)$) is the smallest subspace of $\mathbb{R}^n$ containing S . For a matrix, $\mathbf{M} \in \mathbb{R}^{n \times m}$, its span is the span of its columns

vectors, written as $\text{span}(\mathbf{M})$. We write matrix transpose as $\mathbf{M}^t$. Let $\tilde{\mathbf{B}}$ denote the Gram-Schmidt orthogonalization of $\mathbf{B}$. The GSO of an ordered basis $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_k]$ is assumed to be from left to right, $\widetilde{\mathbf{b}}_1 = \mathbf{b}_1$, unless stated otherwise.

[0091] Recall $\mathbf{M}$'s singular value decomposition (SVD), i.e. $\mathbf{M} = \mathbf{VDW} \in \mathbb{R}^{n \times m}$ where
 5 $\mathbf{V} \in \mathbb{R}^{n \times n}$ along with $\mathbf{W} \in \mathbb{R}^{m \times m}$ are unitary, and $\mathbf{D} \in \mathbb{R}^{n \times m}$ is a triangular matrix containing $\mathbf{M}$'s singular values. Further, let $q = \min\{n, m\}$ and $\mathbf{D}_q = \text{diag}(s_1, \dots, s_q)$ be the diagonal matrix containing $\mathbf{M}$'s singular values $\{0 \leq s_i\}$. Then, $\mathbf{D} = \mathbf{D}_q$ when $n = m$, $\mathbf{D} = [\mathbf{D}_q \ 0]$ when $m > n$, and $\mathbf{D} = \begin{bmatrix} \mathbf{D}_q \\ \mathbf{0} \end{bmatrix}$ in the case $m < n$.

[0092] A symmetric matrix $\Sigma \in \mathbb{R}^{n \times n}$ is *positive semi-definite* if for all $\vec{x} \in \mathbb{R}^n$, we have
 10 $\vec{x}^t \Sigma \vec{x} \geq 0$. It is *positive definite*, $\Sigma > 0$, if it is positive semi-definite and $\vec{x}^t \Sigma \vec{x} = 0$ implies $\vec{x} = \mathbf{0}$. We say $\Sigma_1 > \Sigma_2$ ($\geq$) if $\Sigma_1 - \Sigma_2$ is positive-(semi)definite. This forms a partial ordering on the set of positive semi-definite matrices, and we denote $\Sigma \geq \alpha I$ often as $\Sigma \geq \alpha$ for constants $\alpha \in \mathbb{R}^+$. For any positive semi-definite matrix Σ , we write $\sqrt{\Sigma}$ to be any full rank matrix T such that $\Sigma = TT^t$. We say T is a *square root* of Σ . For two positive semi-
 15 definite matrices, Σ_1 and Σ_2 , we denote the positive semi-definite matrix formed by their block diagonal concatenation as $\Sigma_1 \oplus \Sigma_2$. Let M^* denote Hermitian transpose. The (*Moore-Penrose*) *pseudoinverse* for matrix M with SVD $M = \mathbf{VDW}$ is $M^+ = \mathbf{W}D^+V^*$ where D^+ is given by transposing D and inverting M 's nonzero singular values. For example, $T = sI$ and $T^+ = s^{-1}I$ for a covariance $\Sigma = s^2I$. (An analogous T^+ is given for the non-spherical, full-
 20 rank case $\Sigma > 0$ using Σ 's diagonalization.)

C. Lattices background

[0093] An n -dimensional lattice of rank $k \leq n$ is a discrete additive subgroup of $\mathbb{R}^n$. Given k linearly independent basis vectors $B = b_1, \dots, b_k \in \mathbb{Z}^n$, the lattice generated by B is

$$(B) = (b_1, \dots, b_k) = \left\{ \sum_{i=1}^k x_i \cdot b_i, x_i \in \mathbb{Z} \right\}.$$

25 [0094] Given $n, m \in \mathbb{Z}$ and a modulus $q \geq 2$, we often use q -ary lattices and their cosets, denoted as

$$\text{for } A \in \mathbb{Z}_q^{n \times m}, \text{ denote } \Lambda^\perp[A] \text{ or } \Lambda_q^\perp[A] \text{ as } \{x \in \mathbb{Z}^m : A \cdot x = 0 \pmod{q}\};$$

for $A \in \mathbb{Z}_q^{n \times m}$, $w \in \mathbb{Z}_q^n$, denote $\Lambda_w^\perp[A]$ as $\{x \in \mathbb{Z}^m: A \cdot x = w(\text{mod } q)\}$.

[0095] We can define Gaussians on lattices. For any $s > 0$ define the Gaussian function on $\mathbb{R}^n$ with parameter s :

$$\forall \mathbf{x} \in \mathbb{R}^n, \rho_s(x) = e^{-\pi \|\mathbf{x}\|^2 / s^2}.$$

5 [0096] For any $c \in \mathbb{R}^n$, real $s > 0$, and n -dimensional lattice Λ , define the discrete Gaussian distribution $D_{\Lambda+c,s}$ as:

$$\forall x \in \Lambda + c, D_{\Lambda+c,s}(x) = \frac{\rho_s(x)}{\rho_s(\Lambda + c)}.$$

The subscripts s and c are taken to be 1 and 0 (respectively) when omitted.

10 [0097] For any positive semidefinite $\Sigma = T \cdot T^t$, define the non-spherical Gaussian function as

$$\forall x \in \text{span}(T) = \text{span}(\Sigma), \rho_T(x) = e^{-\pi x^t \Sigma^+ x},$$

and $\rho_T(x) = 0$ for all $\mathbf{x} \notin \text{span}(\Sigma)$. Note that $\rho_T(\cdot)$ only depends on Σ but not the specific choice of the T , so we may write $\rho_T(\cdot)$ as $\rho_{\sqrt{\Sigma}}(\cdot)$.

15 [0098] For any $c \in \mathbb{R}^n$, any positive semidefinite Σ , and n -dimensional lattice Λ such that $(\Lambda + c) \cap \text{span}(\Sigma)$ is non-empty, define the discrete Gaussian distribution $D_{\Lambda+c,\sqrt{\Sigma}}$ as:

$$\forall \mathbf{x} \in \Lambda + c, D_{\Lambda+c,\sqrt{\Sigma}}(x) = \frac{\rho_{\sqrt{\Sigma}}(x)}{\rho_{\sqrt{\Sigma}}(\Lambda + c)}.$$

Such a discrete Gaussian distribution may be referred to as a lattice Gaussian distribution.

[0099] We recall the definition of smoothing parameter and some useful facts.

20 [0100] **Definition 2.1** [Smoothing parameter [MR07]] For any lattice Λ and positive real $\varepsilon > 0$, the smoothing parameter $\eta_\varepsilon(\Lambda)$ is the smallest real $s > 0$ such that $\rho_{1/s}(\Lambda^* \setminus \{0\}) \leq \varepsilon$.

[0101] Notice that for two lattices of the same rank $\Lambda_1 \subseteq \Lambda_2$, the denser lattice always has the smaller smoothing parameter, i.e. $\eta_\varepsilon(\Lambda_2) \leq \eta_\varepsilon(\Lambda_1)$. We can also use a generalization of the smoothing parameter to the non-spherical Gaussian.

25 [0102] **Definition 2.2** For a positive semi-definite $\Sigma = \mathbf{T}\mathbf{T}^t$, an $\varepsilon > 0$, and a lattice Λ with $\text{span}(\Lambda) \subseteq \text{span}(\Sigma)$, we say $\eta_\varepsilon(\Lambda) \leq \sqrt{\Sigma}$ if $\eta_\varepsilon(\mathbf{T}^+ \Lambda) \leq 1$.

[0103] When the covariance matrix $\Sigma > 0$ and the lattice Λ are full-rank, $\sqrt{\Sigma} \geq \eta_\varepsilon(\Lambda)$ is equivalent to the minimum eigenvalue of Σ , $\lambda_{\min}(\Sigma)$, being at least $\eta_\varepsilon^2(\Lambda)$.

[0104] **Lemma 2.3** [Smoothing parameter bound from [GPV08]] For any n -dimensional lattice (B) and for any $\omega(\sqrt{\log n})$ function, there is a negligible $\varepsilon(n)$ for which

5 [0105] The following is a generalization of [GPV08, Corollary 2.8] for non-spherical Gaussian.

[0106] **Corollary 2.5** [Smooth over the cosets] Let Λ, Λ' be n -dimensional lattices s.t. $\Lambda' \subseteq \Lambda$. Then for any $\varepsilon > 0$, $\sqrt{\Sigma} \geq \eta_\varepsilon(\Lambda')$, and $c \in \text{span}(\Lambda)$, we have

$$\Delta(D_{\Lambda+c, \sqrt{\Sigma}} \text{mod}', U(\Lambda \text{ mod } \Lambda')) < 2\varepsilon$$

10 [0107] **Lemma 2.6.** [[PR06, MR07]] Let B be a basis of an n -dimensional lattice, and let $s \geq \|B\| \cdot \omega(\log n)$, then $\Pr_{x \in D_{B,s}}[\|x\| \geq s \cdot \sqrt{n} \vee x = 0] \leq (n)$.

[0108] We can use the following general theorem regarding the linear transformation, $\mathbf{T}$, of a discrete gaussian. It states that as long as the original discrete gaussian is smooth enough it the kernel of $\mathbf{T}$, then the distribution transformed by $\mathbf{T}$ is statistically close to another discrete
15 Gaussian.

[0109] **Theorem 2.7** For any positive definite Σ , vector $\mathbf{c}$, lattice coset $A := \Lambda + a \subset c + \text{span}(\Sigma)$, and linear transformation $\mathbf{T}$, if the lattice $\Lambda_T = \Lambda \cap \ker(\mathbf{T})$ satisfies $\text{span}(\Lambda_T) = \ker(\mathbf{T})$ and $\eta_\varepsilon(\Lambda_T) \leq \sqrt{\Sigma}$, then

$$T(D_{A,c,\sqrt{\Sigma}}) \stackrel{\bar{\varepsilon}}{\approx} D_{TA,Tc,T\sqrt{\Sigma}}$$

20 where $\bar{\varepsilon} = 2\varepsilon/(1 - \varepsilon)$. Let us remark that if T is injective (i.e. $\ker(T)$ is trivial), then

$$T(D_{A,c,\sqrt{\Sigma}}) = D_{TA,Tc,T\sqrt{\Sigma}}.$$

D. Gadgets, or G-Lattices

[0110] Let $\mathbf{G} = \mathbf{I}_n \otimes g^t \in \mathbb{Z}_q^{n \times nk}$ with $g^t = (1, b, \dots, b^{k-1})$, $k = \lceil \log_b q \rceil$. $\mathbf{G}$ is commonly referred to the gadget matrix. The gadget matrix's q -ary lattice, $\Lambda_q^\perp(\mathbf{G})$, is the direct sum of n
25 copies of the lattice $\Lambda_q^\perp(g^t)$. Further, $\Lambda_q^\perp(g^t)$ has a simple basis,

$$B_q = \begin{bmatrix} b & & & q_0 \\ -1 & \ddots & & \vdots \\ & \ddots & b & q_{k-2} \\ & & -1 & q_{k-1} \end{bmatrix}$$

where $(q_0, \dots, q_{k-1}) \in \{0, 1, \dots, b-1\}^k$ is the b-ary decomposition of the modulus, q . When $q = b^k$, we can set $q_0 = q_1 = \dots = q_{k-2} = 0$ and $q_{k-1} = b$. Either way, the integer cosets of $\Lambda_q^\perp(g^t)$ can be viewed as the syndromes of g^t as a check matrix, in the terminology of coding theory. These cosets are expressed as $\Lambda_u^\perp(g^t) = \{x \in \mathbb{Z}^k : g^t x = u \bmod q\} = \Lambda_q^\perp(g^t) + u$ where u can be any coset representative. A simple coset representative of $\Lambda_u^\perp(g^t)$ is the b-ary decomposition of u . The integer cosets of $\Lambda_q^\perp(\mathbf{G})$ are expressed through the direct-sum construction, $\Lambda_u^\perp(\mathbf{G}) = \Lambda_{u_0}^\perp(g^t) \oplus \dots \oplus \Lambda_{u_n}^\perp(g^t)$ where $u = (u_1, \dots, u_n) \in \mathbb{Z}_q^n$. We call $\mathbf{G}$ a gadget matrix since the following problems, SIS and LWE, are easily solved on the matrix $\mathbf{G}$ [MP12].

[0111] As an example, consider a gadget matrix $\mathbf{G}$ with base $b = 2$, $q = 8$ (and thus $k = 3$), and $n = 3$. The gadget matrix is then $\mathbf{G} = \mathbf{I}_3 \otimes g^t \in \mathbb{Z}_8^{3 \times 9}$ with $g^t = (1, 2, 2^2)$, or

$$\mathbf{G} = \begin{bmatrix} 1 & 2 & 4 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 2 & 4 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 2 & 4 \end{bmatrix}$$

[0112] We can then form a reduced gadget matrix $\mathbf{F}$ by dropping the l columns with the smallest powers b from the gadget matrix $\mathbf{G}$. For example, if $l = 1$, then this corresponds to dropping any column with a 1. The reduced gadget matrix is thus

$$\mathbf{F} = \begin{bmatrix} 2 & 4 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 4 & 0 & 0 \\ 0 & 0 & 0 & 0 & 2 & 4 \end{bmatrix}$$

E. SIS and LWE

[0113] We first recall the short integer solution (SIS) problem.

[0114] **Definition 2.8** [SIS [Ajt96]] For any $n, m, q \in \mathbb{Z}$ and $\beta \in \mathbb{R}$, define the short integer solution problem $SIS_{n,m,q}$, as follows: Given $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, find a non-zero vector $\mathbf{x} \in \mathbb{Z}^m$ such that $\|\mathbf{x}\| \leq \beta$, and

$$\mathbf{Ax} = 0 \bmod q.$$

[0115] **Definition 2.9** [ISIS] For any $n, m, q \in \mathbb{N}$, define the inhomogeneous short integer solution problem $ISIS_{n,m,q}$, as follows: Given $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{y} \in \mathbb{Z}_q^n$, find $\mathbf{x} \in \mathbb{Z}^m$ such that $\|\mathbf{x}\| \leq \beta$, and

$$\mathbf{Ax} = \mathbf{y} \bmod q.$$

5 [0116] **Lemma 2.10** [Hardness of (I)SIS based on the lattice problems in the worst case [Ajt96, MR07, GPV08]] For any $m = \text{poly}(n)$, any $\beta > 0$, and any sufficiently large $q \geq \beta \cdot \text{poly}(n)$, solving $ISIS_{n,m,q,\beta}$ or $ISIS_{n,m,q,\beta}$ (where $\mathbf{y}$ is sampled uniformly from $\mathbb{Z}_q^n$) with non-negligible probability is as hard as solving GapSVP_γ and SIVP_γ on arbitrary n -dimensional lattices with overwhelming probability, for some approximation factor $\gamma = \beta \cdot$
10 $\text{poly}(n)$.

[0117] All the (I)SIS problems and their variants admit the Hermite normal form (HNF), where the public matrix A is of the form $[I_n | A']$ where $A' \in \mathbb{Z}_q^{m-n}$. The HNF variant of (I)SIS is as hard as the standard (I)SIS. This can be seen by rewriting $A \in \mathbb{Z}_q^{n \times m}$ as $A = [A_1 | A_2] = A_1 \cdot [I_n | A_1^{-1} \cdot A_2]$ (we may work with n, q such that $A_1 \leftarrow U(\mathbb{Z}_q^{n \times n})$ is invertible with non-
15 negligible probability).

[0118] We can also recall the decisional learning with errors (LWE) problem.

[0119] **Definition 2.11** [Decisional learning with errors [Reg09]] For $n, m \in \mathbb{N}$ and modulus $q \geq 2$, distributions for secret vectors, public matrices, and error vectors $\theta, \pi, \chi \subseteq \mathbb{Z}_q$. An LWE sample is obtained from sampling $s \leftarrow \theta^n$, $A \leftarrow \pi^{n \times m}$, $e \leftarrow \chi^m$, and outputting
20 $(A, y^t := s^t A + e^t \bmod q)$.

[0120] We say that an algorithm solves $\text{LWE}_{n,m,q,\theta,\pi,\chi}$ if it distinguishes the LWE sample from a random sample distributed as $\pi^{n \times m} \times U(\mathbb{Z}_q^m)$ with probability greater than $1/2$ plus non-negligible.

[0121] **Lemma 2.12** [Hardness of LWE based on the lattice problems in the worst case [Reg09, Pei09, BLP+13, PRS17]] Given $n \in \mathbb{N}$, for any $m = \text{poly}(n)$, $q \leq 2^{\text{poly}(n)}$. Let $\theta = \pi = U(\mathbb{Z}_q)$, $\chi = D_{\mathbb{Z},s}$ where $s \geq 2\sqrt{n}$. If there exists an efficient (possibly quantum) algorithm that breaks $\text{LWE}_{n,m,q,\theta,\pi,\chi}$, then there exists an efficient (possibly quantum) algorithm for solving GapSVP_γ and SIVP_γ on arbitrary n -dimensional lattices with overwhelming probability, for some approximation factor $\gamma = \tilde{O}(nq/s)$.
25

[0122] The next lemma shows that LWE with the secret sampled from the error distribution is as hard as the standard LWE.

[0123] **Lemma 2.13** [[ACPS09, BLP+13]] For n, m, q, s chosen as was in Lemma 2.12, $LWE_{n,m',q,D_{\mathbb{Z},s},U(\mathbb{Z}_q),D_{\mathbb{Z},s}}$ is as hard as $LWE_{n,m,q,U(\mathbb{Z}_q),U(\mathbb{Z}_q),D_{\mathbb{Z},s}}$ for $m' \leq m - (16n + 4\log\log q)$.

[0124] The (I)SIS and LWE problems bear similarities. It is sometimes convenient to talk about one of the two problems, and an analogous result immediately applies for the other. On a high level they can be considered equally hard, since breaking one of them would morally break the other. But a careful examine of the current status of the reductions suggests that LWE is a stronger assumption than (I)SIS. Since if there is a polynomial time algorithm that solves the SIS problem with respect to a public matrix $\mathbf{A} \in \mathbb{Z}^{n \times m}$, we can simply use the SIS solution $\mathbf{x}$ to break the decisional LWE problem by computing the inner-product of $\mathbf{x}$ and the LWE challenge vector $\mathbf{y}$; on the other hand, given a polynomial time algorithm that solves the LWE problem, we know of a more involved polynomial time quantum algorithm that solves SIS [SSTX09].

[0125] Nevertheless, a *trapdoor* for a public matrix $A \in \mathbb{Z}_q^{n \times m}$ is defined in [Ajt99, GPV08] as anything that allows us to efficiently solve both the (I)SIS and LWE problems w.r.t. A .

III. APPROXIMATE TRAPDOOR FOR AJTAI'S FUNCTION

[0126] Given a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, define an *approximate trapdoor* of $\mathbf{A}$ as anything that allows us to efficiently solve the approximate version of the ISIS problem w.r.t. $\mathbf{A}$. We first define the approximate ISIS problem.

[0127] **Definition 3.1** (Approximate ISIS) For any $n, m, q \in \mathbb{N}$ and $\alpha, \beta \in \mathbb{R}$, define the *approximate inhomogeneous short integer solution problem* $\text{Approx. ISIS}_{n,m,q,\alpha,\beta}$ as follows:
Given $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{y} \in \mathbb{Z}_q^n$, find a vector $\mathbf{x} \in \mathbb{Z}^m$ such that $\|\mathbf{x}\| \leq \beta$, and there is a vector $\mathbf{z} \in \mathbb{Z}^n$ satisfying

$$\|\mathbf{z}\| \leq \alpha \text{ and } \mathbf{Ax} = \mathbf{y} + \mathbf{z} \pmod{q}.$$

[0128] That is to say, given a verification matrix $\mathbf{A}$ and a message (or hash) $\mathbf{y}$, the approximate ISIS relates to the problem of finding a signature vector $\mathbf{x}$ such that $\mathbf{Ax} = \mathbf{y}$ with

a small error $\mathbf{z}$. Let us remark that the approximate ISIS is only non-trivial when the bounds α, β are relatively small compared to the modulus q . Also, our definition chooses to allow the zero vector to be a valid solution, which means when $\|\mathbf{y}\| \leq \alpha$, the zero vector is trivially a solution. Such a choice does not cause a problem since the interesting case in the application
 5 is to handle all the $\mathbf{y} \in \mathbb{Z}_q^n$ or $\mathbf{y}$ sampled uniformly from $\mathbb{Z}_q^n$.

[0129] Definition 3.2 (Approximate trapdoor) *A string τ is called an (α, β) -approximate trapdoor for a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ if there is a polynomial time algorithm (in $n, m, \log q$) that given $\tau, \mathbf{A}$ and any $\mathbf{y} \in \mathbb{Z}_q^n$, outputs a non-zero vector $\mathbf{x} \in \mathbb{Z}^m$ such that $\|\mathbf{x}\| \leq \beta$, and there is a vector $\mathbf{z} \in \mathbb{Z}^n$ satisfying*

$$10 \quad \|\mathbf{z}\| \leq \alpha \text{ and } \mathbf{Ax} = \mathbf{y} + \mathbf{z}(\text{mod } q).$$

[0130] To make sense of the approximate trapdoor, we argue that for those who do not have the trapdoor, the approximate ISIS problem is a candidate one-way function under proper settings of parameters.

[0131] First, we observe a rather obvious reduction that bases the hardness of approximate
 15 ISIS on the hardness of decisional LWE with low-norm secret (e.g. when the secret is sampled from the error distribution). In the theorem statement below, when the norm symbol is applied on a distribution D , i.e. $\|D\|$, it denotes the lowest value $v \in \mathbb{R}^+$ such that $\Pr_{d \leftarrow D}[\|d\| < v] > 1 - \text{negl}(\lambda)$.

[0132] Theorem 3.3 *For $n, m, q \in \mathbb{Z}$, $\alpha, \beta \in \mathbb{R}^+$, θ, π, χ be distributions over $\mathbb{Z}$ such that
 20 $q > 4(\|\theta\| \cdot (\alpha + 1) + \|\theta^n\| \cdot \alpha + \|\chi^m\| \cdot \beta)$. Then $\text{LWE}_{n,m,q,\theta,\pi,\chi} \leq_p \text{Approx.ISIS}_{n,m,q,\alpha,\beta}$.*

[0133] Proof. Suppose there is a polynomial time adversary A that breaks $\text{Approx.ISIS}_{n,m,q,\alpha,\beta}$, we build a polynomial time adversary B that breaks decisional LWE.

[0134] Let $r = \lceil \alpha \rceil + 1$. Given an LWE challenge $(\mathbf{A}, \mathbf{w}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m$, where $\mathbf{w}$ is either
 25 an LWE sample or sampled uniformly from $\mathbb{Z}_q^m$. B picks a vector $\mathbf{y} := (r, 0, \dots, 0)^t \in \mathbb{Z}_q^n$, sends $\mathbf{A}$ and $\mathbf{y}$ to the adversary A as an approximate ISIS challenge. A replies with $\mathbf{x} \in \mathbb{Z}^m$ such that $\|\mathbf{x}\| \leq \beta$, and there is a vector $\mathbf{z} \in \mathbb{Z}^n$ satisfying

$$\|\mathbf{z}\| \leq \alpha \text{ and } \mathbf{Ax} = \mathbf{y} + \mathbf{z}(\text{mod } q).$$

[0135] Note that $\mathbf{x} \neq \mathbf{0}$ since $\|\mathbf{y}\| > \alpha$.

[0136] B then computes $v := \langle \mathbf{w}, \mathbf{x} \rangle$. If $\mathbf{w}^t = \mathbf{s}^t \mathbf{A} + \mathbf{e}^t$ for $\mathbf{s} \leftarrow \theta^n$, $\mathbf{e} \leftarrow \chi^m$, then

$$v = (\mathbf{s}^t \mathbf{A} + \mathbf{e}^t) \mathbf{x} = \mathbf{s}^t (\mathbf{y} + \mathbf{z}) + \mathbf{e}^t \mathbf{x} \Rightarrow \|v\| \leq \|\theta\| \cdot r + \|\theta^n\| \cdot \alpha + \|\chi^m\| \cdot \beta < q/4.$$

[0137] Otherwise v distributes uniformly random over $\mathbb{Z}_q$. So B can compare v with the threshold value and wins the decisional LWE challenge with probability $1/2$ plus non-negligible.

[0138] Alternatively, we can also prove that the approximate ISIS problem is as hard as the standard ISIS.

[0139] **Theorem 3.4** $ISIS_{n,n+m,q,\beta} \geq_p \text{Approx. ISIS}_{n,m,q,\alpha+\beta}$,
 $ISIS_{n,n+m,q,\alpha+\beta} \leq_p \text{Approx. ISIS}_{n,m,q,\alpha,\beta}$.

[0140] *Proof.* The reductions go through the HNFs of the ISIS and the approximate ISIS problems. We will show $ISIS = \text{HNF. ISIS} = \text{HNF. Approx. ISIS} = \text{Approx. ISIS}$ under proper settings of parameters.

[0141] Recall that $ISIS_{n,m,q,\beta} = \text{HNF. ISIS}_{n,m,q,\beta}$ as explained in the preliminary. Also, $\text{HNF. ISIS}_{n,m,q,\beta} \geq_p \text{HNF. Approx. ISIS}_{n,m,q,\alpha,\beta}$ for any $\alpha \geq 0$ by definition. It remains to show the rest of the connections.

[0142] **Lemma 3.5** $\text{HNF. ISIS}_{n,m,q,\alpha+\beta} \leq_p \text{HNF. Approx. ISIS}_{n,m,q,\alpha,\beta}$.

[0143] *Proof.* Suppose there is a polynomial time algorithm A that solves $\text{HNF. Approx. ISIS}_{n,m,q,\alpha,\beta}$, we build a polynomial time algorithm B that solves $\text{HNF. ISIS}_{n,m,q,\alpha+\beta}$. Given an HNF. ISIS instance $[\mathbf{I}_n | \mathbf{A}] \in \mathbb{Z}_q^{n \times m}$, $\mathbf{y}$, B passes the same instance to A , gets back a vector $\mathbf{x}$ such that

$$[\mathbf{I}_n | \mathbf{A}] \cdot \mathbf{x} = \mathbf{y} + \mathbf{z} \pmod{q}.$$

where $\|\mathbf{x}\| \leq \beta$, $\|\mathbf{z}\| \leq \alpha$. Now write $\mathbf{x} =: [\mathbf{x}_1^t | \mathbf{x}_2^t]^t$ where $\mathbf{x}_1 \in \mathbb{Z}^n$, $\mathbf{x}_2 \in \mathbb{Z}^m$. Then $\mathbf{x}' := [(\mathbf{x}_1 - \mathbf{z})^t | \mathbf{x}_2^t]^t$ satisfies

$$[\mathbf{I}_n | \mathbf{A}] \cdot \mathbf{x}' = \mathbf{y} \pmod{q},$$

and $\|\mathbf{x}'\| \leq \alpha + \beta$. So $\mathbf{x}'$ is a valid solution to HNF. ISIS .

[0144] **Lemma 3.6** $\text{HNF. Approx. ISIS}_{n,n+m,q,\alpha,\beta} \leq_p \text{Approx. ISIS}_{n,m,q,\alpha,\beta}$.

[0145] *Proof.* Suppose there is a polynomial time algorithm A that solves Approx. ISIS $_{n,m,q,\alpha,\beta}$, we build a polynomial time algorithm B that solves HNF. Approx. ISIS $_{n,n+m,q,\alpha,\beta}$. Given $[\mathbf{I}_n|\mathbf{A}] \in \mathbb{Z}_q^{n \times (n+m)}$, $\mathbf{y} \in \mathbb{Z}_q^n$ as an HNF. Approx. ISIS instance, B passes $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{y}$ to A , gets back a short vector $\mathbf{x} \in \mathbb{Z}^m$. Then $[\mathbf{0}_n^t | \mathbf{x}^t]^t$ is a valid solution to the HNF. Approx. ISIS instance.

[0146] **Lemma 3.7** *HNF. Approx. ISIS $_{n,n+m,q,\alpha,\beta} \geq_p$ Approx. ISIS $_{n,m,q,\alpha+\beta,\beta}$.*

[0147] *Proof.* Suppose there is a polynomial time algorithm A that solves HNF. Approx. ISIS $_{n,n+m,q,\alpha,\beta}$, we build a polynomial time algorithm B that solves Approx. ISIS $_{n,m,q,\alpha+\beta,\beta}$. Given an Approx. ISIS instance $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{y} \in \mathbb{Z}^n$, B passes $[\mathbf{I}_n|\mathbf{A}] \in \mathbb{Z}_q^{n \times (n+m)}$, $\mathbf{y}$ as an HNF. Approx. ISIS instance to A , gets back an answer $\mathbf{x} \in \mathbb{Z}^{m+n}$ such that

$$[\mathbf{I}_n|\mathbf{A}] \cdot \mathbf{x} = \mathbf{y} + \mathbf{z} \pmod{q}, \quad (1)$$

where $\|\mathbf{x}\| \leq \beta$, $\|\mathbf{z}\| \leq \alpha$.

[0148] Now write $\mathbf{x} =: [\mathbf{x}_1^t | \mathbf{x}_2^t]^t$ where $\mathbf{x}_1 \in \mathbb{Z}^n$, $\mathbf{x}_2 \in \mathbb{Z}^m$. Rewriting Eqn. (1) gives

15
$$\mathbf{A} \cdot \mathbf{x}_2 = \mathbf{y} + \mathbf{z} - \mathbf{x}_1 \pmod{q},$$

so $\mathbf{x}_2$ is a valid solution to Approx. ISIS $_{n,m,q,\alpha+\beta,\beta}$.

[0149] Theorem 3.4 then follows the lemmas above.

[0150] The following statement immediately follows the proof of Lemma 3.7.

[0151] **Corollary 3.8** *An (α, β) -approximate trapdoor for $[\mathbf{I}|\mathbf{A}]$ is an $(\alpha + \beta, \beta)$ -approximate trapdoor for $\mathbf{A}$.*

[0152] **Lemma 3.9** *HNF. Approx. ISIS $_{n,n+m,q,\alpha,\beta} \geq_p$ Approx. ISIS $_{n,m,q,2(\alpha+\beta),\beta}$.*

[0153] *Proof.* Suppose there is a polynomial time algorithm A that solves HNF. Approx. ISIS $_{n,n+m,q,\alpha,\beta}$, we build a polynomial time algorithm B that solves Approx. ISIS $_{n,m,q,\alpha+\beta,\beta}$. Given an Approx. ISIS instance $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{y} \in \mathbb{Z}^n$, B first check if $\|\mathbf{y}\| \leq \alpha + \beta$. If so, then find $\mathbf{y}' =: \mathbf{y} + \Delta\mathbf{y}$ such that $\|\mathbf{y}'\| > \alpha + \beta$ and $\|\mathbf{y}' - \mathbf{y}\| \leq \alpha + \beta$; if not then simply set $\mathbf{y}' =: \mathbf{y}$ and $\Delta\mathbf{y} =: \mathbf{0}$. B then passes $[\mathbf{I}_n|\mathbf{A}] \in \mathbb{Z}_q^{n \times (n+m)}$, $\mathbf{y}'$ as an HNF. Approx. ISIS instance to A , gets back an answer $\mathbf{x} \in \mathbb{Z}^{m+n}$ such that

$$[\mathbf{I}_n | \mathbf{A}] \cdot \mathbf{x} = \mathbf{y}' + \mathbf{z}(\text{mod } q), \quad (2)$$

where $\|\mathbf{x}\| \leq \beta$, $\|\mathbf{z}\| \leq \alpha$.

[0154] Now write $\mathbf{x} =: [\mathbf{x}_1^t | \mathbf{x}_2^t]^t$ where $\mathbf{x}_1 \in \mathbb{Z}^n$, $\mathbf{x}_2 \in \mathbb{Z}^m$. Since $\|\mathbf{y}'\| > \alpha + \beta$, $\mathbf{x}_2$ must be a non-zero short vector. Rewriting Eqn. (2) gives

$$\mathbf{A} \cdot \mathbf{x}_2 = \mathbf{y} + \Delta \mathbf{y} + \mathbf{z} - \mathbf{x}_1(\text{mod } q),$$

so $\mathbf{x}_2$ is a valid solution to Approx. ISIS $_{n,m,q,2(\alpha+\beta),\beta}$.

IV. APPROXIMATE GADGET TRAPDOOR

[0155] We present an instantiation of an approximate trapdoor based on the gadget-based trapdoor generation and preimage sampling algorithms of [MP12] (without the tag matrices).

10 In short, we show how to generate a pseudorandom $\mathbf{A}$ with entries modulo q along with an approximate trapdoor $\mathbf{R}$ with small integer entries.

[0156] In the rest of this section, we first recall the exact G-trapdoor from [MP12], then present the approximate trapdoor generation algorithm and the approximate preimage sampling algorithm. Finally we analyze the preimage distribution. The analyses make
15 extensive use of Theorem 2.7 (linear transformations of discrete Gaussians).

A. Recall the G-trapdoor from MP12

[0157] Let $b \geq 2$ be the base for the G-lattice. Let q be the modulus, $k = \lceil \log_b q \rceil$. b is typically chosen to be 2 for simplicity, but often a higher base b is used for efficiency trade-offs in lattice-based schemes.

20 [0158] Recall the MP12 gadget-lattice trapdoor technique: the public matrix is

$$\mathbf{A} = [\bar{\mathbf{A}} | \mathbf{G} - \bar{\mathbf{A}}\mathbf{R}]$$

where $\mathbf{G}$ is the commonly used gadget matrix, $\mathbf{G} =: \mathbf{I}_n \otimes \mathbf{g}_k^t$, $\mathbf{g}_k^t := (1, b, \dots, b^{k-1})$, and $\mathbf{R}$ is a secret, trapdoor matrix with small, random entries. $\mathbf{A}$ is either statistically close to uniformly random or pseudorandom, depending on the structure of $\bar{\mathbf{A}}$ and the choice of χ (in
25 the pseudorandom case $\chi \subseteq \mathbb{Z}$ is chosen to be a distribution such that $\text{LWE}_{n,n,q,\chi,U(\mathbb{Z}_q),\chi}$ is hard). In this paper we focus on the pseudorandom case since the resulting public matrix $\mathbf{A}$ and preimage have smaller dimensions.

[0159] In order to sample a short element in $\Lambda_{\mathbf{u}}^{\perp}(\mathbf{A})$, we use the trapdoor to map short coset representatives of $\Lambda_q^{\perp}(\mathbf{G})$ to short coset representatives of $\Lambda_q^{\perp}(\mathbf{A})$ by the relation

$$\mathbf{A} \begin{bmatrix} \mathbf{R} \\ \mathbf{I} \end{bmatrix} = \mathbf{G}.$$

[0160] Using the trapdoor as a linear transformation alone leaks information about the trapdoor. Therefore, we perturb the sample to statistically hide the trapdoor. Let Σ_p be a positive definite matrix defined as $\Sigma_p := s^2 \mathbf{I} - \sigma^2 \begin{bmatrix} \mathbf{R}\mathbf{R}^t & \mathbf{R}^t \\ \mathbf{R} & \mathbf{I} \end{bmatrix}$ where σ is at least $\eta_{\varepsilon}(\Lambda_q^{\perp}(\mathbf{G}))$. The perturbation can be computed offline as $\mathbf{p} \leftarrow D_{\mathbb{Z}^m, \sqrt{\Sigma_p}}$. We then sample a $\mathbf{G}$ -lattice vector in a coset dependent on $\mathbf{p}$ as $\mathbf{z} \leftarrow D_{\Lambda_{\mathbf{v}}^{\perp}(\mathbf{G}), \sigma}$ and $\mathbf{v} = \mathbf{u} - \mathbf{A}\mathbf{p} \in \mathbb{Z}_q^n$. Finally, the preimage is set to be

$$\mathbf{y} := \mathbf{p} + \begin{bmatrix} \mathbf{R} \\ \mathbf{I} \end{bmatrix} \mathbf{z}.$$

B. The algorithms of the approximate $\mathbf{G}$ -trapdoor

[0161] As mentioned in the introduction, one objective of obtaining an approximate trapdoor is to adapt the MP12 algorithms with a gadget matrix without the lower-order entries. Let $0 < l < k$ be the number of lower-order entries dropped from the gadget vector $\mathbf{g} \in \mathbb{Z}^k$. Define the resulting approximate gadget vector as $\mathbf{f} := (b^l, b^{l+1}, \dots, b^{k-1})^t \in \mathbb{Z}^{(k-l)}$. Let $w = n(k-l)$ be the number of columns of the approximate gadget $\mathbf{F} := \mathbf{I}_n \otimes \mathbf{f}^t \in \mathbb{Z}^{n \times w}$. Then the number of columns of $\mathbf{A}$ will be $m := 2n + w$.

[0162] FIG. 7 shows algorithms for generating a trapdoor matrix and verification matrix, and sampling values from a Gaussian gadget lattice distribution.

1. Algorithm 1: Sampling from gadget distribution

[0163] Algorithm 1 shows an algorithm for sampling a vector from a Gaussian distribution based on lattice derived from a gadget matrix. Algorithm 1 takes in a value v from $\mathbb{Z}_q$ (v may be one component of a vector). The gadget sampling algorithm also takes in σ , which defines the width of the distribution. This algorithm may be performed by a sending device, such as sending device 115 of FIG. 1. Additionally, the steps of Algorithm 1 may be performed as part of a signature generation process such as that of FIG. 3.

[0164] In step 701, the sending device can sample a vector $\mathbf{x}$ comprising k values from $\mathbb{Z}$. The values may be pulled from a distribution D . D may be a lattice Gaussian distribution, that is, a Gaussian distribution over a lattice, and the lattice may be defined by a gadget matrix, and the deviation σ . The distribution D is described in further detail in Section II.C.

5 [0165] In step 702, the sending device can drop the first l entries from the vector $\mathbf{x}$ and take the last $k - l$ entries of the vector $\mathbf{x}$ as an intermediate vector $\mathbf{z}$.

[0166] In step 703, the sampling algorithm can return the intermediate vector $\mathbf{z}$.

2. Algorithm 2: Generating trapdoor matrix

10 [0167] Algorithm 2 shows a trapdoor matrix generation algorithm. As input, this algorithm takes in a security parameter λ . Trapdoor generation can be done by a generating device, such as generating device 105 of FIG. 1. Additionally, the steps of Algorithm 2 may be performed as part of a key generation process such as that of FIG. 2.

[0168] Prior to generating the trapdoor matrix, the generating device can generate a gadget matrix $\mathbf{G}$. The gadget matrix $\mathbf{G}$ may be a matrix where each column comprises one value of an ascending power of a base b (e.g., $b = 2$), for k powers ($k = \lceil \log_b q \rceil$). There may be n repeats of the k columns with ascending powers of b , resulting in an $n \times k$ matrix. The generating device can then generate a reduced gadget matrix $\mathbf{F}$ that has dimension $n \times w$ (where $w = n(k - l)$), after dropping the first l columns (the columns with the smallest powers) of each group of k columns of the gadget matrix $\mathbf{G}$. More detail about gadget matrices and the generation of a reduced gadget matrix can be found in section II.D.

20

[0169] In step 711, the generating device can randomly sample an $n \times n$ uniform matrix $\hat{\mathbf{A}}$ from a uniform distribution over $\mathbb{Z}_q$ for a modulus q .

[0170] In step 712, the generating device can form a generation matrix $\bar{\mathbf{A}}$. The generation matrix $\bar{\mathbf{A}}$ can be a $n \times 2n$ matrix formed from concatenating the uniform matrix $\hat{\mathbf{A}}$ and an $n \times n$ identity matrix $\mathbf{I}_n$. Thus $\bar{\mathbf{A}} := [\mathbf{I}_n, \hat{\mathbf{A}}]$.

25

[0171] In step 713, an approximate trapdoor $\mathbf{R}$ can be sampled from a distribution χ . The approximate trapdoor $\mathbf{R}$ can be a $2n \times w$ matrix. The distribution χ may be the distribution $D_{\mathbb{Z}, \sigma}$, the Gaussian distribution over integers with width σ .

[0172] In step 714, the generating device can form the verification matrix $\mathbf{A}$. The verification matrix $\mathbf{A}$ can be $\mathbf{A} := [\bar{\mathbf{A}}|\mathbf{F} - \bar{\mathbf{A}}\mathbf{R}]$. The verification matrix may have dimension $n \times m$, where $m = 2n + w = 2n + n(k - l)$, with elements from $\mathbb{Z}_q$.

3. Algorithm 3: Generating lattice-based signature

5 [0173] Algorithm 3 shows a signature generation algorithm. The signature generation algorithm takes as input the verification matrix $\mathbf{A}$, the trapdoor matrix $\mathbf{R}$, a message vector $\mathbf{u}$, and a distribution width s . The distribution width s may be related to the distribution width σ through the equation in table 1. Signature generation may be done by a sending device, such as sending device 115 of FIG. 1. Additionally, the steps of Algorithm 3 may be performed as
10 part of a signature generation process such as that of FIG. 3.

[0174] In step 721, the generating device samples a perturbation vector $\mathbf{p}$ from a Gaussian distribution. The perturbation vector $\mathbf{p}$ has length m . The covariance of the perturbation is chosen to hide the trapdoor as in the exact gadget-based trapdoor, i.e.

$$\Sigma_p := s^2 \mathbf{I}_m - \sigma^2 \begin{bmatrix} \mathbf{R}\mathbf{R}^t & \mathbf{R} \\ \mathbf{R}^t & \mathbf{I} \end{bmatrix}.$$

15

[0175] In step 722, the generating device can form a lattice vector $\mathbf{v}$ from the message $\mathbf{u}$, the verification matrix $\mathbf{A}$, and the perturbation vector $\mathbf{p}$. The lattice vector $\mathbf{v}$ has length n and is formed as $\mathbf{v} = \mathbf{u} - \mathbf{A}\mathbf{p}$.

[0176] In step 723, the generating device can sample an approximate gadget preimage vector $\mathbf{z}$ (i.e., the intermediate vector), using the gadget sampling algorithm of Algorithm 1.
20 The gadget sampling algorithm can sample a vector $\mathbf{z}$ for each element of the lattice vector $\mathbf{v}$ and concatenate them together, thus there are $n(k - l)$ total elements in the total preimage vector $\mathbf{z}$.

[0177] In step 724, the generating device can form the signature vector $\mathbf{y}$, which is a vector
25 of length $m = 2n + w = 2n + n(k - l)$. The perturbation $\mathbf{p}$ can be added to a product of the trapdoor matrix $\mathbf{R}$ and the preimage vector $\mathbf{z}$. The trapdoor matrix $\mathbf{R}$ has dimension $2n \times w$, and thus it can be stacked on a $w \times w$ identity matrix, in order to have the appropriate dimension of a $2n + w \times w$ matrix. This stack is multiplied by the preimage vector $\mathbf{z}$ which has length w , so the product is a vector of length $2n + w = m$.

[0178] In step 725, the signature generation algorithm can return the signature vector $\mathbf{y}$.

[0179] The results of this section are summarized in the following theorem.

[0180] **Theorem 4.1** *There exists probabilistic, polynomial time algorithms $\text{Approx.TrapGen}(\cdot)$ and $\text{Approx.SamplePre}(\cdot, \cdot, \cdot)$ satisfying the following.*

1. $\text{Approx.TrapGen}(n)$ takes an input a security parameter n and returns a matrix-approximate trapdoor pair $(\mathbf{A}, \mathbf{R}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}^{2n \times n(k-l)}$.
2. Let $\mathbf{A}$ be generated with an approximate trapdoor as above, $\text{approx.}\mathbf{A}^{-1}(\cdot)$ denote the approximate preimage sampling algorithm, $\text{Approx.SamplePre}(\mathbf{A}, \mathbf{R}, s, \cdot)$. The following two distributions are statistically indistinguishable:

$$\{(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{e}) : \mathbf{u} \leftarrow U(\mathbb{Z}_q^n), \mathbf{y} \leftarrow \text{approx.}\mathbf{A}^{-1}(\mathbf{u}), \mathbf{e} = \mathbf{u} - \mathbf{A}\mathbf{y} \bmod q\}$$

and

$$\{(\mathbf{A}, \mathbf{y}, \mathbf{A}\mathbf{y} + \mathbf{e} \bmod q, \mathbf{e} \bmod q) : \mathbf{y} \leftarrow D_{\mathbb{Z}^m, s}, \mathbf{e} \leftarrow D_{\mathbb{Z}^n, \sigma \sqrt{(b^{2l}-1)/(b^2-1)}}\}$$

for any $\sigma \geq \sqrt{b^2 + 1} \omega(\sqrt{\log n})$ and $s \geq \sqrt{\sigma^2 s_1^2(\mathbf{R}) + \omega(\sqrt{\log n})}$. Furthermore, in the second distribution, $\mathbf{A}$ is computationally indistinguishable from random assuming

- 15 $\text{LWE}_{n,n,q,\chi,U(\mathbb{Z}_q),\chi}$. [0181] Let us remark that while the preimage sampling algorithm works for any target $\mathbf{u} \in \mathbb{Z}_q^n$, we are able to prove the approximate preimage sample-coset pair $(\mathbf{y}, \mathbf{u})$ hides the approximate trapdoor, $\mathbf{R}$, over a uniform target $\mathbf{u} \in \mathbb{Z}_q^n$. This is unlike the exact gadget-based trapdoor setting in [MP12] which proves the trapdoor is hidden for each fixed $\mathbf{u}$. In the overview of the proof idea, we will explain where the proof breaks down when
- 20 we try to spell out the preimage distributions of every $\mathbf{u} \in \mathbb{Z}_q^n$.

C. Analyze the preimage and error distributions for a uniformly random target

[0182] This subsection is dedicated to proving Theorem 4.1.

- 25 [0183] Let $\mathbf{x} \in \mathbb{Z}^{nk}$ denote the short preimage of $\mathbf{u} - \mathbf{A}\mathbf{p}(\bmod q)$ under the full gadget matrix $\mathbf{G}$, i.e. $\mathbf{G}\mathbf{x} = \mathbf{u} - \mathbf{A}\mathbf{p}(\bmod q)$. The main idea of the proof is to first show that the joint distribution of $(\mathbf{p}, \mathbf{x})$ produced in Algorithm 4 is statistically close to $D_{\Lambda_{\mathbf{A}}^\perp[\mathbf{A}, \mathbf{G}], \sqrt{\Sigma_p \oplus \sigma^2}}$ for any

$\mathbf{u} \in \mathbb{Z}_q^n$ (this is a stronger theorem than what we need). And then apply the linear transformation theorem on $(\mathbf{p}, \mathbf{x})$ to obtain the distributions of $\mathbf{y}$ and $\mathbf{e}$. However, applying the linear transformation theorem directly on the lattice coset $\Lambda_{\mathbf{u}}^\perp[\mathbf{A}, \mathbf{G}]$ leads to a technical problem. That is, the intermediate lattice intersections $\Lambda_{\mathbf{T}}$ required in Theorem 2.7 have large smoothing parameters. To get around this problem, we use the condition that $\mathbf{u}$ is uniformly random to argue that $(\mathbf{p}, \mathbf{x})$ is statistically close to $D_{\mathbb{Z}^{m+nk}, \sqrt{\Sigma_p \oplus \sigma^2}}$. Then the support of $(\mathbf{p}, \mathbf{x})$ is $\mathbb{Z}^{m+nk}$, so we can apply the linear transformation theorem to prove the final distributions of $\mathbf{y}$ and $\mathbf{e}$ are trapdoor-independent.

[0184] Formally, let $\varepsilon = \text{negl}(\lambda) > 0$. We first prove three lemmas.

10 [0185] **Lemma 4.2** *For any $\sigma \geq \eta_\varepsilon(\Lambda^\perp(\mathbf{G}))$, the random process of first choosing $\mathbf{u} \leftarrow U(\mathbb{Z}_q^n)$ then returning $\mathbf{x} \leftarrow D_{\Lambda_{\mathbf{u}}^\perp(\mathbf{G}), \sigma}$ is statistically close to sampling $\mathbf{x} \leftarrow D_{\mathbb{Z}^{nk}, \sigma}$.*

[0186] *Proof.* The proof follows directly from $\det(\Lambda_q^\perp(\mathbf{G})) = q^n$ and Corollary 2.7.

Alternatively, one can use two applications of the fact $\rho_r(\Gamma + \mathbf{c}) \in (1 \pm \varepsilon)\sigma^n / \det(\Gamma)$ for any $r \geq \eta_\varepsilon(\Gamma)$. The latter yields $\Pr\{\text{Process returns } \mathbf{x}\} \in \left(\frac{1-\varepsilon}{1+\varepsilon}, \frac{1+\varepsilon}{1-\varepsilon}\right) \cdot D_{\mathbb{Z}^{nk}, \sigma}(\mathbf{x})$.

15 [0187] **Lemma 4.3** *The following random processes are statistically close for any $\sigma \geq \sqrt{b^2 + 1}\omega(\sqrt{\log n}) \geq \eta_\varepsilon(\mathbf{g}^t)$: sample a uniformly random coset $u \leftarrow U(\mathbb{Z}_q)$, then return the error $e := u - \text{Gsamp.Cut}(u, \sigma) \bmod q$. Or, return $e \leftarrow D_{\mathbb{Z}, \sigma \sqrt{(b^{2l}-1)/(b^2-1)}} \bmod q$.*

[0188] *Proof.* For a fixed $u \in \mathbb{Z}_q$, the error is distributed as

$$e := \mathbf{L} \cdot D_{\Lambda_{\mathbf{u}}^\perp(\mathbf{g}^t), \sigma} := (1, b, b^2, \dots, b^{l-1}, \mathbf{0}) \cdot D_{\Lambda_{\mathbf{u}}^\perp(\mathbf{g}^t), \sigma}.$$

20 [0189] Randomizing over u gives us $e \sim_\varepsilon \mathbf{L} \cdot D_{\mathbb{Z}^k, \sigma}$ by Lemma 4.2. Next, we apply Theorem 2.7. Let $\mathbf{B}_q = [\mathbf{b}_1, \dots, \mathbf{b}_k]$ be the basis of $\Lambda_q^\perp(\mathbf{g}^t)$ given in Section 2. Then, the kernel of $\mathbf{L}$ is generated by $\{\mathbf{b}_1, \dots, \mathbf{b}_{l-1}, \mathbf{e}_l, \dots, \mathbf{e}_{k-1}\}$ (recall that $\{\mathbf{e}_i\}_{i=1}^n$ denotes the canonical basis $\mathbb{R}^n$) and $\Lambda_{\mathbf{L}} := \text{Ker}(\mathbf{L}) \cap \mathbb{Z}^k$ is the set of all integer combinations of these vectors. $\Lambda_{\mathbf{L}}$ spans the kernel of $\mathbf{L}$ and the smoothing parameter of $\Lambda_{\mathbf{L}}$ is at most $\sqrt{b^2 + 1}\omega(\sqrt{\log n})$ since $\|\mathbf{b}_i\| \leq \sqrt{b^2 + 1}$ for all $i = 1, \dots, l-1$.

[0190] Let $\mathbf{R}' := \begin{bmatrix} \mathbf{R} \\ \mathbf{I}_{n(k-l)} \end{bmatrix}$. Next, we analyze the distribution given by the linear transformation representing truncating $\mathbf{x} \in \mathbb{Z}^{nk}$ in the joint distribution of $(\mathbf{p}, \mathbf{x}) \leftarrow$

$D_{\mathbb{Z}^{m+nk}, \sqrt{\Sigma_p \oplus \sigma^2 \mathbf{I}_{nk}}}$ to $\mathbf{z} \in \mathbb{Z}^{n(k-l)}$ (as in GSAMP.CUT, Algorithm 1) and returning $\mathbf{y} = \mathbf{p} +$

$\mathbf{R}'\mathbf{z}$. For simplicity, we permute the columns of $\mathbf{G}$ to $\mathbf{G}' := [\mathbf{F} | \mathbf{I}_n \otimes (1, b, \dots, b^{l-1})]$. This allows us to express truncation and convolution as a simple linear transformation: $\mathbf{y} = \mathbf{L}(\mathbf{p}, \mathbf{x})$ for $\mathbf{L} := [\mathbf{I}_m | \mathbf{R}' | \mathbf{0}]$.

5 [0191] **Lemma 4.4** For any $\sqrt{\Sigma_p}, \sigma \geq \eta_\varepsilon(\mathbb{Z}^{nk})$, $LD_{\mathbb{Z}^{m+nk}, \sqrt{\Sigma_p \oplus \sigma^2 \mathbf{I}_{nk}}}$ is statistically close to $D_{\mathbb{Z}^m, S}$.

[0192] *Proof.* The range and covariance are immediate. Next, we use Theorem 2.7. The kernel of $\mathbf{L}$ is given by all vectors $(\mathbf{a}, \mathbf{b}, \mathbf{c})$ where $(\mathbf{b}, \mathbf{c}) \in \mathbb{R}^{nk}$ and $\mathbf{a} = -\mathbf{R}'\mathbf{b}$. The integer lattice $\mathbb{Z}^{m+nk}$ contains all such integer vectors so $\Lambda_{\mathbf{L}} := \mathbb{Z}^{m+nk} \cap \ker(\mathbf{L})$ spans $\mathbf{L}$'s kernel.

10 [0193] Now we determine the smoothing parameter of $\Lambda_{\mathbf{L}}$. Add nl zero columns to $\mathbf{R}'$ to form $\mathbf{R}_{nk} := [\mathbf{R}' | \mathbf{0}]$ and rotate the space by

$$\mathbf{Q} := \begin{bmatrix} \mathbf{I} & \mathbf{R}_{nk} \\ \mathbf{0} & \mathbf{I}_{nk} \end{bmatrix}.$$

[0194] This rotation yields $\mathbf{Q}\Lambda_{\mathbf{L}} = \{\mathbf{0}\} \oplus \mathbb{Z}^{nk}$ and since rotating a covariance does not change its eigenvalues, we have $\sqrt{\Sigma_p \oplus \sigma^2 \mathbf{I}_{nk}} \geq \eta_\varepsilon(\mathbb{Z}^{nk})$. This implies $LD_{\mathbb{Z}^{m+nk}, \sqrt{\Sigma_p \oplus \sigma^2 \mathbf{I}_{nk}}}$

15 is statistically close to $D_{\mathbb{Z}^m, S}$.

[0195] We are now ready to prove Theorem 4.1.

[0196] *Proof.* (of Theorem 5.1) We prove Theorem 5.1 via a sequence of hybrid distributions. Let $\mathbf{M} := [\mathbf{I}_n \otimes (1, b, \dots, b^{l-1}, 0, \dots, 0)] \in \mathbb{Z}^{n \times (nk)}$ be the linear transformation corresponding to obtaining the error vector $\mathbf{e} \in \mathbb{Z}^n$ from $\mathbf{x} \in \mathbb{Z}^{nk}$. Let $\mathbf{L}$ be the linear

20 transformation used in Lemma 4.4.

[0197] Note that the equation $\mathbf{u} = \mathbf{v} + \mathbf{A}\mathbf{p} = \mathbf{G}\mathbf{x} + \mathbf{A}\mathbf{p} = \mathbf{F}\mathbf{z} + \mathbf{e} + \mathbf{A}\mathbf{p} = \mathbf{A}\mathbf{L}(\mathbf{p}, \mathbf{x}) + \mathbf{e} = \mathbf{A}\mathbf{y} + \mathbf{e} \pmod{q}$ holds in all the hybrids. The difference lies in the order and the distributions that these vectors are sampled from. In all the hybrids except the last one, $\mathbf{A}$ is generated as $\mathbf{A} := [\bar{\mathbf{A}} | \mathbf{F} - \bar{\mathbf{A}}\mathbf{R}]$ for $\bar{\mathbf{A}} := [\mathbf{I}_n, \mathbf{A}_0] \in \mathbb{Z}_q^{n \times 2n}$.

25 [0198] **Real distribution:** The real distribution of $\{(\mathbf{A}, \mathbf{y}, \mathbf{u}, \mathbf{e})\}$ is:

$\mathbf{A}, \mathbf{u} \leftarrow U(\mathbb{Z}_q^n), \mathbf{p} \leftarrow D_{\mathbb{Z}^m, \sqrt{\Sigma_p}}, \mathbf{v} := \mathbf{u} - \mathbf{A}\mathbf{p}, \mathbf{x} \leftarrow D_{\Lambda_{\mathbf{V}}^\perp(\mathbf{G}), \sigma}, \mathbf{e} = \mathbf{M}(\mathbf{x}),$ and $\mathbf{y} = \mathbf{L}(\mathbf{p}, \mathbf{x})$.

[0199] **Hybrid 1:** Here we swap the order of sampling $\mathbf{u}$ and $\mathbf{v}$ by first sampling $\mathbf{v} \leftarrow U(\mathbb{Z}_q^n)$ and setting $\mathbf{u} := \mathbf{v} + \mathbf{A}\mathbf{p}$: $\mathbf{A}, \mathbf{v} \leftarrow U(\mathbb{Z}_q^n), \mathbf{p} \leftarrow D_{\mathbb{Z}^m, \sqrt{\Sigma_p}}, \mathbf{u} = \mathbf{v} + \mathbf{A}\mathbf{p}$. We keep $\mathbf{x}, \mathbf{e}$, and $\mathbf{y}$ unchanged: $\mathbf{x} \leftarrow D_{\Lambda_{\mathbf{v}}^\perp(\mathbf{G}), \sigma}, \mathbf{e} = \mathbf{M}(\mathbf{x})$, and $\mathbf{y} = L(\mathbf{p}, \mathbf{x})$. Then, the real distribution and Hybrid 1 are the same.

5 [0200] **Hybrid 2:** Instead of sampling a uniform $\mathbf{v} \in \mathbb{Z}_q^n$ and a G-lattice sample $\mathbf{x} \leftarrow D_{\Lambda_{\mathbf{v}}^\perp(\mathbf{G}), \sigma}$, we sample $\mathbf{x} \leftarrow D_{\mathbb{Z}^{nk}, \sigma}$ and $\mathbf{v} = \mathbf{G}\mathbf{x} \in \mathbb{Z}_q^n$. The rest remains the same:

$$\mathbf{A}, \mathbf{x} \leftarrow D_{\mathbb{Z}^{nk}, \sigma}, \mathbf{v} = \mathbf{G}\mathbf{x}, \mathbf{p} \leftarrow D_{\mathbb{Z}^m, \sqrt{\Sigma_p}}, \mathbf{u} = \mathbf{v} + \mathbf{A}\mathbf{p}, \mathbf{e} = \mathbf{M}(\mathbf{x}), \text{ and } \mathbf{y} = L(\mathbf{p}, \mathbf{x}).$$

[0201] Lemma 4.2 implies Hybrid 1 and Hybrid 2 are statistically close.

[0202] **Hybrid 3:** We combine $\mathbf{p}, \mathbf{x}$ into the joint distribution $(\mathbf{p}, \mathbf{x}) \leftarrow D_{\mathbb{Z}^{m+nk}, \sqrt{\Sigma_p \oplus \sigma^2 \mathbf{I}}}$:

10 $\mathbf{A}, (\mathbf{p}, \mathbf{x}) \leftarrow D_{\mathbb{Z}^{m+nk}, \sqrt{\Sigma_p \oplus \sigma^2 \mathbf{I}}}, \mathbf{e} = \mathbf{M}(\mathbf{x}), \mathbf{y} = L(\mathbf{p}, \mathbf{x}), \mathbf{v} = \mathbf{G}\mathbf{x}, \text{ and } \mathbf{u} = \mathbf{v} + \mathbf{A}\mathbf{p}.$

[0203] **Hybrid 4:** Here we apply the linear transformation theorem on $\mathbf{e}, \mathbf{y}$:

$$\mathbf{A}, (\mathbf{p}, \mathbf{x}) \leftarrow D_{\mathbb{Z}^{m+nk}, \sqrt{\Sigma_p \oplus \sigma^2 \mathbf{I}}}, \mathbf{e} \leftarrow D_{\mathbb{Z}^n, \sigma \sqrt{(b^{2l}-1)/(b^2-1)}}, \mathbf{y} = D_{\mathbb{Z}^m, s}, \mathbf{v} = \mathbf{G}\mathbf{x}, \text{ and } \mathbf{u} = \mathbf{v} + \mathbf{A}\mathbf{p}.$$

[0204] Lemmas 4.3 and 4.4 imply Hybrids 3 and 4 are statistically close.

15 [0205] **Final distribution:** Sample $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times m})$ and keep the rest of the vectors from the same distribution as Hybrid 4 (notice that the trapdoor $\mathbf{R}$ of $\mathbf{A}$ are not used to sample $\mathbf{p}, \mathbf{x}, \mathbf{e}$ and $\mathbf{y}$). The final distribution is computationally indistinguishable from Hybrid 4 assuming $\text{LWE}_{n,n,q,\chi,U(\mathbb{Z}_q),\chi}$.

D. From an approximate G-trapdoor to an approximate kernel-trapdoor

20 [0206] Given an approximate G-trapdoor for $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ in the form of $\mathbf{R} \in \mathbb{Z}^{(m-w) \times w}$ where $w = n(k-l)$ and $\mathbf{A} \cdot \begin{pmatrix} \mathbf{R} \\ \mathbf{I}_w \end{pmatrix} = \mathbf{F}(\text{mod } q)$. We transform it to an approximate kernel of m columns (i.e., smaller than $n+m$) that still preserves the trapdoor functionality.

[0207] The transformation is similar to the one from [MP12] that works over the exact trapdoors. Write $\mathbf{A} = [\mathbf{A}_1 | \mathbf{A}_2]$ where $\mathbf{A}_1 \in \mathbb{Z}_q^{n \times (m-w)}$. Let $\mathbf{W} \in \mathbb{Z}^{w \times (m-w)}$ be an approximate short preimage of $\mathbf{A}_1$ under the public matrix $\mathbf{F}$, i.e. $\mathbf{F} \cdot \mathbf{W} = \mathbf{A}_1 + \mathbf{E}_1(\text{mod } q)$. Let $\mathbf{S} \in \mathbb{Z}^{w \times w}$

be an approximated kernel of $\mathbf{F}$, i.e. $\mathbf{F} \cdot \mathbf{S} = \mathbf{E}_2(\text{mod } q)$ (both the norms of $\mathbf{E}_1$ and $\mathbf{E}_2$ are small). Then

$$(\mathbf{A}_1 | \mathbf{A}_2) \cdot \underbrace{\begin{pmatrix} \mathbf{I} & \mathbf{R} \\ \mathbf{I} & \mathbf{I} \end{pmatrix} \cdot \begin{pmatrix} \mathbf{I} & \mathbf{I} \\ \mathbf{W} & \mathbf{S} \end{pmatrix}}_{=: \mathbf{K}} = (\mathbf{A}_1 | \mathbf{F}) \cdot \begin{pmatrix} \mathbf{I} & \mathbf{I} \\ \mathbf{W} & \mathbf{S} \end{pmatrix} = (\mathbf{E}_1 | \mathbf{E}_2)(\text{mod } q)$$

[0208] So $\mathbf{K}$ is an approximate kernel of $\mathbf{A}$ of m columns (i.e. lower than $n + m$) that is at the same time an approximate trapdoor.

V. HASH-AND-SIGN SIGNATURE INSTANTIATED WITH THE APPROXIMATE TRAPDOOR

[0209] We spell out the details of the hash-and-sign signature scheme from [GPV08] instantiated with the approximate G-trapdoor instead of an exact trapdoor.

10 [0210] Recall the parameters from the last section. We set $k = \lceil \log_b q \rceil$, set l to be the number of entries dropped from the G-trapdoor such that $1 \leq l < k$ and $m = n(2 + (k - l))$. Let $\sigma, s \in \mathbb{R}^+$ be the discrete Gaussian widths of the distributions over the cosets of $\Lambda_q^\perp(\mathbf{G})$ and $\Lambda_q^\perp(\mathbf{A})$ respectively. Let χ be the distribution of the entries of the trapdoor $\mathbf{R}$ chosen so that $\text{LWE}_{n,n,q,\chi,U(\mathbb{Z}_q),\chi}$ is hard.

15 [0211] **Construction 5.1** *Given an approximate trapdoor sampler from Theorem 4.1, a hash function $H = \{H_\lambda: \{0,1\}^* \rightarrow R_\lambda\}$ modeled as a random oracle, we build a signature scheme as follows.*

- $\text{Gen}(1^\lambda)$: The key-generation algorithm can sample $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together with its (α, β) -approximate trapdoor $\mathbf{R}$ from *Approx. TrapGen* (1^λ) . Let the range R_λ of H be $\mathbb{Z}_q^n$. It can output $\mathbf{A}$ as the verification matrix and can keep $\mathbf{R}$ as the trapdoor matrix (i.e., the secret signing key).
- $\text{Sig}(\mathbf{R}, m)$: The signing algorithm can check if the message-signature pair $(m, \mathbf{x}_m)$ has been produced before. If so, it outputs $\mathbf{x}_m$ as the signature of m ; if not, it computes $\mathbf{u} = H(m)$, and samples an approximate preimage $\mathbf{x}_m \leftarrow \text{Approx. SamplePre}(\mathbf{A}, \mathbf{R}, \mathbf{u}, s)$. It outputs $\mathbf{x}_m$ as the signature and stores $(m, \mathbf{x}_m)$ in the list.
- $\text{Ver}(\mathbf{A}, m, \mathbf{x})$: The verification algorithm checks if $\|\mathbf{x}\| \leq \beta$ and $\|\mathbf{A} \cdot \mathbf{x} - H(m)\| \leq \alpha$. If so, it outputs accept; otherwise, it outputs reject.

[0212] We provide a proof-of-concept implementation of the signature. Embodiments can use several groups of parameters using different dimensions n , moduli q , bases b , targeting different security level (mainly around 100-bit and 192-bit security). In each group of parameters, we use fixed n , q , b , and compare the use of exact trapdoor (under our reference implementation) versus approximate trapdoor. In FIG. 8 and FIG. 9, 6 groups of parameters are listed.

[0213] Let us first explain how we make the security estimations. The concrete security estimation of lattice-based cryptographic primitive is a highly active research area and more sophisticated methods are proposed recently. Here we use relatively simple methods to estimate the pseudorandomness of the public-key (henceforth “LWE security”), and the hardness of breaking approximate ISIS (henceforth “AISIS security”). Let us remark that our estimations may not reflect the state-of-art, but at least provide a fair comparison of the parameters for the exact trapdoor versus the approximate trapdoor.

[0214] LWE security depends on the choices of q , n , and the Gaussian width τ of the trapdoor $\mathbf{R}$. The estimation of LWE security was done with the online LWE bit security estimator with BKZ as the reduction model [ACD+18].

[0215] For the approximate ISIS problem, the only direct cryptanalysis result we are aware of is the work of Bai et al. [BGLS19], but it is not clearly applicable to the parameters we are interested. Instead we estimate AISIS through $\text{ISIS}_{n,m,q,\alpha+\beta}$ following the reduction in Lemma 3.5, where α and β are the upper-bounds of l_2 norm of the error $\mathbf{z}$ and preimage $\mathbf{x}$. We estimate the security level of $\text{ISIS}_{n,m,q,\alpha+\beta}$ based on how many operations BKZ would take to find a vector in the lattice $\Lambda_{\frac{1}{q}}(\mathbf{A})$ of length $\alpha + \beta$. Further, we can throw away columns in $\mathbf{A}$. We choose to only use $2n$ columns of $\mathbf{A}$ as done in [BFRS18], denoted $\mathbf{A}_{2n}$, since Minkowski’s theorem tells us $\Lambda_{\frac{1}{q}}(\mathbf{A}_{2n})$ has a short enough vector. Following [APS15, ACD+18], we use sieving as the SVP oracle with time complexity $2^{.292k+16.4}$ in the block size, k . BKZ is expected to return a vector of length $\delta^{2n} \det^{1/2n}$ for a lattice of dimension $2n$. Hence, we found the smallest block size k achieving the needed δ corresponding to forging a signature, $\frac{\alpha+\beta}{\sqrt{q}} = \delta^{2n}$. Finally, we used the heuristic $\delta \approx \left(\frac{k}{2\pi e} (\pi k)^{1/k} \right)^{1/2(k-1)}$ to determine the relation between k and δ , and we set the total time complexity of BKZ with block-size k , dimension $2n$ as $\frac{(2n)^3}{k^2} \cdot \log(2n) \cdot \text{time}(SVP) = \frac{(2n)^3}{k^2} \cdot \log(2n) \cdot 2^{.292k+16.4}$ [Che13, APS15].

[0216] For an estimation of 100-bit security, our reference implementation for the exact trapdoor under the modulus $q \approx 2^{24}$ and base $b = 2$ matches the parameters reported in [BB13] (the parameters in the other implementation [BFRS18, GPR+18] are possibly measured in different ways). We also use smaller moduli and bigger bases to reduce the size and increase the security level. The parameters in FIG. 8 and FIG. 9 suggest that for all the choices of q and b , using the approximate gadget trapdoor by setting $l = \lceil (\log_b q)/2 \rceil$ saves about half of the sizes in the public key and signatures comparing to using the exact trapdoor, with even a slight increase in the security estimation.

[0217] Our implementation shows that the sizes of the public-key & signature can be reduced to 5 kB & 4.45 kB for an estimation of 100-bit security, and 11.25 kB & 9.38 kB for an estimation of 192-bit security. Those are still larger than, but much closer to the sizes for the signatures based on the rejection sampling approach [Lyu12, BG14, DKL+18, ABB+19]. As a reference, the sizes of the public-key & signature for qTESLA [ABB+19] are 4.03 kB & 3.05 kB for an estimation of 128-bit security, and 8.03 kB & 6.03 kB for an estimation of 192-bit security.

[0218] In the security analysis we use the following properties on the distributions produced by *Approx.SamplePre* proven in Theorem 4.1:

1. An approximate preimage of a uniformly random image is statistically close to a distribution D_{pre} , which is independent of $\mathbf{A}$ and $\mathbf{R}$.
2. For a uniformly random image, the error vector $\mathbf{e} := \mathbf{u} - \mathbf{A}\mathbf{y} \pmod{q}$ is statistically close to a distribution D_{err} , which is also independent of $\mathbf{A}$ and $\mathbf{R}$.

[0219] To prove that the signature satisfies the strong EU-CMA security, we can make use of an additional “near-collision-resistance” property for Ajtai’s function, which can be based on the standard SIS assumption. Let us remark that without this property, we can still prove the signature scheme satisfies static security based on the hardness of the approximate ISIS problem.

[0220] **Lemma 5.2 (The near-collision-resistance of Ajtai’s function)** *For any $n, m, q \in \mathbb{N}$ and $\alpha, \beta \in \mathbb{R}$. If there is an efficient adversary A that given $\mathbf{A} \leftarrow U(\mathbb{Z}_q^{n \times m})$, finds $\mathbf{x}_1 \neq \mathbf{x}_2 \in \mathbb{Z}^m$ such that*

$$\|\mathbf{x}_1\| \leq \beta \text{ and } \|\mathbf{x}_2\| \leq \beta \text{ and } \|\mathbf{A}\mathbf{x}_1 - \mathbf{A}\mathbf{x}_2 \pmod{q}\| \leq 2\alpha$$

[0221] Then there is an efficient adversary B that solves $SIS_{n,n+m,q,2\sqrt{\alpha^2+\beta^2}}$.

[0222] *Proof.* Suppose B gets an HNF. $SIS_{n,n+m,q,2\sqrt{\alpha^2+\beta^2}}$ challenge (which is as hard as $SIS_{n,n+m,q,2\sqrt{\alpha^2+\beta^2}}$) with the public matrix $[\mathbf{I}_n|\mathbf{A}]$, B sends $\mathbf{A}$ to A , gets back $\mathbf{x}_1 \neq \mathbf{x}_2 \in \mathbb{Z}^m$ such that

$$5 \quad \|\mathbf{x}_1\| \leq \beta \text{ and } \|\mathbf{x}_2\| \leq \beta \text{ and } \|\mathbf{y} := \mathbf{A}\mathbf{x}_1 - \mathbf{A}\mathbf{x}_2 \pmod{q}\| \leq 2\alpha$$

B then sets $\mathbf{z} := [-\mathbf{y}^t | (\mathbf{x}_1 - \mathbf{x}_2)^t]^t$ as the solution. $\mathbf{z}$ is then non-zero and satisfies $\|\mathbf{z}\| \leq 2\sqrt{\alpha^2 + \beta^2}$ and $[\mathbf{I}_n|\mathbf{A}]\mathbf{z} = \mathbf{0} \pmod{q}$.

[0223] **Theorem 5.3** Construction 6.1 is strongly existentially unforgeable under a chosen-message attack in the random oracle model assuming the hardness of $SIS_{n,n+m,q,2\sqrt{\alpha^2+\beta^2}}$ and

$$10 \quad LWE_{n,n,q,\chi,U(\mathbb{Z}_q),\chi}.$$

[0224] *Proof.* The correctness follows the functionality of the approximate trapdoor.

[0225] Suppose there is a polynomial time adversary A that breaks the strong EU-CMA of the signature scheme, we construct a polynomial time adversary B that breaks the near-collision-resistance of Ajtai's function, which is as hard as $SIS_{n,n+m,q,2\sqrt{\alpha^2+\beta^2}}$ due to Lemma

15 5.2.

[0226] To start, B sends Ajtai's function $\mathbf{A}$ to A as the public key for the signature scheme. Once A makes a random oracle query w.r.t. a message m , B samples $\mathbf{y} \leftarrow D_{\text{pre}}$, computes $\mathbf{u} := \mathbf{A}\mathbf{y} + D_{\text{err}} \pmod{q}$ as the random oracle response on m . B then replies $\mathbf{u}$ to A and stores $(m, \mathbf{u})$ in the random oracle storage, $(m, \mathbf{y})$ in the message-signature pair storage. Once A makes a signing query on the message m (wlog assume m has been queried to the random oracle before, since if not B can query it now), B finds $(m, \mathbf{y})$ in the storage and reply $\mathbf{y}$ as the signature. The signatures and the hash outputs produced by B are indistinguishable from the real ones due to the properties of the distributions D_{pre} and D_{err} , and the assumption that a real public key is indistinguishable from random under $LWE_{n,n,q,\chi,U(\mathbb{Z}_q),\chi}$.

25 [0227] Without loss of generality, assume that before A tries to forge a signature on m^* , A has queried H on m^* . Denote the pair that B prepares and stores in the random oracle storage as $(m^*, \mathbf{u}^*)$, and the pair in the signature storage as $(m^*, \mathbf{y}^*)$. Finally A outputs $\mathbf{y}$ as the forged signature on m^* . So we have $\|\mathbf{A}(\mathbf{y} - \mathbf{y}^*) \pmod{q}\| \leq 2\alpha$. It remains to prove that $\mathbf{y} \neq$

y^* so as to use them as a near-collision-pair. If m^* has been queried to the signing oracle before, then $y \neq y^*$ by the definition of a success forgery; if m^* has not been queried to the signing oracle before, then y^* is with high min-entropy by the settings of the parameter, so $y \neq y^*$ with overwhelming probability.

5 VI. COMPUTER SYSTEM

[0228] Any of the computer systems mentioned herein may utilize any suitable number of subsystems. Examples of such subsystems are shown in FIG. 10 in computer apparatus 700. In some embodiments, a computer system includes a single computer apparatus, where the subsystems can be components of the computer apparatus. In other embodiments, a
10 computer system can include multiple computer apparatuses, each being a subsystem, with internal components. A computer system can include desktop and laptop computers, tablets, mobile phones and other mobile devices.

[0229] The subsystems shown in FIG. 10 are interconnected via a system bus 75. Additional subsystems such as a printer 74, keyboard 78, storage device(s) 79, monitor 76,
15 which is coupled to display adapter 82, and others are shown. Peripherals and input/output (I/O) devices, which couple to I/O controller 71, can be connected to the computer system by any number of means known in the art such as input/output (I/O) port 77 (e.g., USB, FireWire®). For example, I/O port 77 or external interface 81 (e.g. Ethernet, Wi-Fi, etc.) can be used to connect computer system 10 to a wide area network such as the Internet, a mouse
20 input device, or a scanner. The interconnection via system bus 75 allows the central processor 73 to communicate with each subsystem and to control the execution of a plurality of instructions from system memory 72 or the storage device(s) 79 (e.g., a fixed disk, such as a hard drive, or optical disk), as well as the exchange of information between subsystems. The system memory 72 and/or the storage device(s) 79 may embody a computer readable
25 medium. Another subsystem is a data collection device 85, such as a camera, microphone, accelerometer, and the like. Any of the data mentioned herein can be output from one component to another component and can be output to the user.

[0230] A computer system can include a plurality of the same components or subsystems, e.g., connected together by external interface 81, by an internal interface, or via removable
30 storage devices that can be connected and removed from one component to another component. In some embodiments, computer systems, subsystem, or apparatuses can communicate over a network. In such instances, one computer can be considered a client and

another computer a server, where each can be part of a same computer system. A client and a server can each include multiple systems, subsystems, or components.

[0231] Aspects of embodiments can be implemented in the form of control logic using hardware circuitry (e.g. an application specific integrated circuit or field programmable gate array) and/or using computer software with a generally programmable processor in a modular or integrated manner. As used herein, a processor can include a single-core processor, multi-core processor on a same integrated chip, or multiple processing units on a single circuit board or networked, as well as dedicated hardware. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will know and appreciate other ways and/or methods to implement embodiments of the present invention using hardware and a combination of hardware and software.

[0232] Any of the software components or functions described in this application may be implemented as software code to be executed by a processor using any suitable computer language such as, for example, Java, C, C++, C#, Objective-C, Swift, or scripting language such as Perl or Python using, for example, conventional or object-oriented techniques. The software code may be stored as a series of instructions or commands on a computer readable medium for storage and/or transmission. A suitable non-transitory computer readable medium can include random access memory (RAM), a read only memory (ROM), a magnetic medium such as a hard-drive, or an optical medium such as a CD (compact disk) or DVD (digital versatile disk), flash memory, and the like. The computer readable medium may be any combination of such storage or transmission devices.

[0233] Such programs may also be encoded and transmitted using carrier signals adapted for transmission via wired, optical, and/or wireless networks conforming to a variety of protocols, including the Internet. As such, a computer readable medium may be created using a data signal encoded with such programs. Computer readable media encoded with the program code may be packaged with a compatible device or provided separately from other devices (e.g., via Internet download). Any such computer readable medium may reside on or within a single computer product (e.g. a hard drive, a CD, or an entire computer system), and may be present on or within different computer products within a system or network. A computer system may include a monitor, printer, or other suitable display for providing any of the results mentioned herein to a user.

[0234] Any of the methods described herein may be totally or partially performed with a computer system including one or more processors, which can be configured to perform the steps. Thus, embodiments can be directed to computer systems configured to perform the steps of any of the methods described herein, potentially with different components

5 performing a respective step or a respective group of steps. Although presented as numbered steps, steps of methods herein can be performed at a same time or in a different order. Additionally, portions of these steps may be used with portions of other steps from other methods. Also, all or portions of a step may be optional. Additionally, and of the steps of any of the methods can be performed with modules, circuits, or other means for performing these

10 steps.

[0235] The specific details of particular embodiments may be combined in any suitable manner without departing from the spirit and scope of embodiments of the invention. However, other embodiments of the invention may be directed to specific embodiments relating to each individual aspect, or specific combinations of these individual aspects.

15 [0236] The above description of exemplary embodiments of the invention has been presented for the purpose of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form described, and many modifications and variations are possible in light of the teaching above. The embodiments were chosen and described in order to best explain the principles of the invention and its practical applications to thereby

20 enable others skilled in the art to best utilize the invention in various embodiments and with various modifications as are suited to the particular use contemplated.

[0237] A recitation of “a”, “an” or “the” is intended to mean “one or more” unless specifically indicated to the contrary. The use of “or” is intended to mean an “inclusive or,” and not an “exclusive or” unless specifically indicated to the contrary.

25 [0238] All patents, patent applications, publications and description mentioned herein are incorporated by reference in their entirety for all purposes. None is admitted to be prior art.

VII. REFERENCES

- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108, 1996.
- 30 [Ajt99] Miklós Ajtai. Generating hard instances of the short basis problem. In Jir Wiedermann, Peter van Emde Boas, and Mogens Nielsen, editors, *Automata*,

Languages and Programming, 26th International Colloquium, ICALP '99, Prague, Czech Republic, July 11-15, 1999, Proceedings, volume 1644 of LNCS, pages 1–9. Springer, 1999.

- [AAAS+19] Gorjan Alagic, Gorjan Alagic, Jacob Alperin-Sheriff, Daniel Apon, David
5 Cooper, Quynh Dang, Yi-Kai Liu, Carl Miller, Dustin Moody, Rene Peralta, et al. *Status Report on the First Round of the NIST Post-Quantum Cryptography Standardization Process*. US Department of Commerce, National Institute of Standards and Technology, 2019.
- [APS15] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of
10 learning with errors. *J. Mathematical Cryptology*, 9(3):169–203, 2015.
- [ACD+18] Martin R. Albrecht, Benjamin R. Curtis, Amit Deo, Alex Davidson, Rachel
Player, Eamonn W. Postlethwaite, Fernando Virdia, and Thomas Wunderer. Estimate all the {LWE, NTRU} schemes! In *Security and Cryptography for Networks - 11th International Conference, SCN 2018, Amalfi, Italy, September*
15 *5-7, 2018, Proceedings*, pages 351–367, 2018.
- [ABB+19] Erdem Alkim, Paulo S. L. M. Barreto, Nina Bindel, Patrick Longa, and
Jefferson E. Ricardini. The lattice-based digital signature scheme qtesla. *IACR Cryptology ePrint Archive*, 2019:85, 2019.
- [AP11] Joël Alwen and Chris Peikert. Generating shorter bases for hard random
20 lattices. *Theory of Computing Systems*, 48(3):535–553, 2011.
- [ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast
cryptographic primitives and circular-secure encryption based on hard learning
problems. In Shai Halevi, editor, *Advances in Cryptology - CRYPTO 2009, 29th*
25 *Annual International Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2009. Proceedings*, volume 5677 of LNCS, pages 595–618. Springer, 2009.
- [BGLS19] Shi Bai, Steven D. Galbraith, Liangze Li, and Daniel Sheffield. Improved
combinatorial algorithms for the inhomogeneous short integer solution problem.
J. Cryptology, 32(1):35–83, 2019.

- [BB13] Rachid El Bansarkhani and Johannes A. Buchmann. Improvement and efficient implementation of a lattice-based signature scheme. In *Selected Areas in Cryptography, volume 8282 of Lecture Notes in Computer Science*, pages 48–67. Springer, 2013.
- 5 [BFRS18] Pauline Bert, Pierre-Alain Fouque, Adeline Roux-Langlois, and Mohamed Sabt. Practical implementation of ring-sis/lwe based signature and IBE. In *PQCrypto, volume 10786 of Lecture Notes in Computer Science*, pages 271–291. Springer, 2018.
- 10 [CC17] Ran Canetti and Yilei Chen. Constraint-hiding constrained prfs for nc^1 from LWE. In *EUROCRYPT (1)*, volume 10210 of *Lecture Notes in Computer Science*, pages 446–476, 2017.
- [Che13] Yuanmi Chen. *Réduction de réseau et sécurité concrète du chiffrement complètement homomorphe*. PhD thesis, Paris 7, 2013.
- 15 [dPLS18] Rafaël del Pino, Vadim Lyubashevsky, and Gregor Seiler. Lattice-based group signatures and zero-knowledge proofs of automorphism stability. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, ON, Canada, October 15-19, 2018*, pages 574–591, 2018.
- 20 [DKL+18] Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):238–268, 2018.
- 25 [FHK+18] Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Prest, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. Falcon: Fast-fourier lattice-based compact signatures over ntru, 2018.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.
- 30 [GGH97] Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Public-key cryptosystems from lattice reduction problems. In *Advances in Cryptology - CRYPTO '97*,

17th Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 1997, Proceedings, pages 112–131, 1997.

- [GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, pages 498–527, 2015.
- [GKW17] Rishab Goyal, Venkata Koppula, and Brent Waters. Lockable obfuscation. In *FOCS*, pages 612–621. IEEE Computer Society, 2017.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
- [GS02] Craig Gentry and Mike Szydlo. Cryptanalysis of the revised ntru signature scheme. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 299–320. Springer, 2002.
- [GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In *STOC*, pages 545–554. ACM, 2013.
- [HHSS17] Shai Halevi, Tzipora Halevi, Victor Shoup, and Noah Stephens-Davidowitz. Implementing bp-obfuscation using graph-induced encoding. In *ACM Conference on Computer and Communications Security*, pages 783–798. ACM, 2017.
- [Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, volume 7237 of *Lecture Notes in Computer Science*, pages 738–755. Springer, 2012.
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, pages 700–718, 2012.
- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measure. *SIAM Journal on Computing*, 37(1):267–302, 2007.

- [PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Theory of Cryptography*, pages 145–166. Springer, 2006.
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
- [SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 617–635. Springer, 2009.

WHAT IS CLAIMED IS:

1 1. A method for use in verifying digital signatures, the method
2 comprising performing, by a generating device:
3 generating a gadget matrix G , wherein columns of the gadget matrix G
4 comprise ascending powers of a base b ;
5 generating a reduced gadget matrix F , wherein the reduced gadget matrix F is
6 generated by dropping l columns of the gadget matrix G , wherein a value of l is
7 determined based on a desired level of security and a key size;
8 sampling a probability distribution to form a trapdoor matrix R that functions
9 as a secret key for generating a digital signature;
10 generating a verification matrix A using the reduced gadget matrix F and the
11 trapdoor matrix R ; and
12 transmitting, to a verification device, the verification matrix A , causing the
13 verification device to verify the digital signature generated by a sending device, using
14 the verification matrix A .

1 2. The method of claim 1, wherein the base b of the gadget matrix G is 2.

1 3. The method of claim 1, wherein the trapdoor matrix R is sampled from
2 a Gaussian distribution.

1 4. The method of claim 1, wherein generating the verification matrix A
2 comprises:

3 sampling a uniform distribution to form a uniform matrix $\hat{A}$;
4 forming a generation matrix $\bar{A}$ by concatenating an identity matrix I and the
5 uniform matrix $\hat{A}$, wherein $\bar{A} = [I \mid \hat{A}]$; and
6 forming the verification matrix A , wherein $A = [\bar{A} \mid F - \bar{A}R]$.

7 5. The method of claim 1, wherein l is in a range 0 to $(\log_b q)/2$ for a
8 modulus q , and wherein the modulus q is between 2^{16} and 2^{24} .

1 6. The method of claim 1, wherein the gadget matrix G comprises n
2 repeats of k columns with ascending powers of the base b , and where generating the reduced
3 gadget matrix F further comprises dropping the l columns with a smallest powers of the base
4 b from each group of k columns of the gadget matrix G .

1 7. The method of claim 1, wherein the probability distribution is a
2 Gaussian distribution.

1 8. The method of claim 1, wherein the verification matrix A is transmitted
2 to the sending device.

1 9. The method of claim 1, wherein the verification matrix A is made
2 public.

1 10. A computer product comprising a computer readable medium storing a
2 plurality of instructions for controlling a computer system to perform the method of any one
3 of claims 1 to 9.

1 11. A system comprising:
2 the computer product of claim 10; and
3 one or more processors for executing instructions stored on the computer
4 readable medium.