

Asynchronous Remote Key Generation for Post-Quantum Cryptosystems from Lattices

Nick Frymann

Surrey Centre for Cyber Security
University of Surrey
Guildford, United Kingdom
n.frymann@surrey.ac.uk

Daniel Gardham

Surrey Centre for Cyber Security
University of Surrey
Guildford, United Kingdom
daniel.gardham@surrey.ac.uk

Mark Manulis

Research Institute CODE
Universität der Bundeswehr München
Munich, Germany
mark@manulis.eu

Abstract—Asynchronous Remote Key Generation (ARKG), introduced by Frymann et al. at CCS 2020, allows for the generation of unlinkable public keys by third parties, for which corresponding private keys may be later learned only by the key pair’s legitimate owner. These key pairs can then be used in common public-key cryptosystems, including signatures, PKE, KEMs, and schemes supporting delegation, such as proxy signatures. The only known instance of ARKG generates discrete-log-based keys.

In this paper, we introduce new ARKG constructions for lattice-based cryptosystems. The key pairs generated using our ARKG scheme can be applied to lattice-based signatures and KEMs, which have recently been selected for standardisation in the NIST PQ process, or as alternative candidates.

In particular, we address challenges associated with the noisiness of lattice hardness assumptions, which requires a new generalised definition of ARKG correctness, whilst preserving the security and privacy properties of the former instantiation. Our ARKG construction uses key encapsulation techniques by Brendel et al. (SAC 2020) coined *Split KEMs*. As an additional contribution, we also show that Kyber (Bos et al., EuroS&P 2018) can be used to construct a Split KEM. The security of our protocol is based on standard LWE assumptions. We also discuss its use with selected candidates from the NIST process and provide an implementation and benchmarks.

1. Introduction

1.1. ARKG and its applications

Asynchronous Remote Key Generation (ARKG) [1] is a key generation protocol that allows for creation of derived public keys, that are associated with original public key. A corresponding derived secret key can be computed later in time, or asynchronously, using only the derived public key, the original secret key and some auxiliary information. All derived public keys remain unlinkable with respect to the original public key, and derivation of a derived secret cannot happen without knowledge of the original secret key. Moreover, only users for whom public keys were derived may compute corresponding private keys when provided with auxiliary data, which may be public. Key pairs computed with ARKG are designed to have statistically-similar distributions to randomly generated

key pairs, and thus can be securely used in public key cryptosystems.

ARKG was proposed for an application in WebAuthn [2], forming part of FIDO2, which is a W3C¹ standard for challenge-response authentication on the web using digital signatures. Its envisioned use is to enable recovery and delegation of WebAuthn accounts. In WebAuthn, authenticators, such as Yubico’s YubiKey or Windows Hello, manage and store private keys on behalf of users, which means account recovery becomes problematic after loss or damage to an authenticator may render a user’s online accounts inaccessible. ARKG produces unlinkable public keys that satisfy WebAuthn’s strong unlinkability requirement, which means users can use a single authenticator when registering for different accounts, yet remain unlinkable across services. Furthermore, the remote feature of this protocol means that the ‘backup’ authenticator does not need to be present at account creation. More recently, ARKG was used in a proxy signature scheme that supports unlinkability across delegations, in a protocol called Proxy Signatures with Unlinkable Warrants [3]. It offers a delegator the ability to delegate signing rights to the proxy signer using a delegation-by-warrant approach—in which a public key for the proxy is signed by the delegator. These proxy signatures were used to enable delegation for WebAuthn accounts. The use of ARKG as a building block results in warrants that are unlinkable with respect to the proxy signer, and the asynchronicity allows the delegation process to be non-interactive.

1.2. Post-quantum ARKG and main challenges

The only known² construction for ARKG is for discrete-log key pairs, i.e. of the form $(sk, pk) = (x, g^x)$ for some group generator g , which are known to be susceptible to quantum attackers [5]. This threat has motivated huge efforts to create protocols that are resistant to quantum attacks. As such, NIST³ have recently selected primitives to standardise post-quantum cryptography based on a range of differing hard problems including isogenies, multivariate polynomials and code-based problems. Many of the candidates submitted to the process are based on

1. <https://www.w3.org/>

2. We note new ARKG constructions for pairing-based cryptosystems introduced concurrently by Frymann et al. at ACNS 2023 [4].

3. <https://csrc.nist.gov/Projects/post-quantum-cryptography>

lattices, such as the signatures in [6]–[8] and KEMs and PKEs in [9]–[11]. Moreover, 3 of the 4 candidates selected for standardisation fall into this category, which is in part due to their relatively more mature security analysis and wide versatility. The focus of our work is, therefore, in the design of suitable ARKG constructions for lattice-based cryptosystems, including those from the current NIST PQ standardisation process.

However, lattices introduce new problems that are not inherent in group-based cryptography. Namely, one of the major security assumptions, Learning with Errors (LWE) [12] critically relies on introducing noise for its security. This noise can affect correctness properties of protocols, and requires careful parameter selection to balance correctness with security. This adds additional complexity to our work as we aim to derive keys compatible with other cryptosystems for which the structure and parameter values of the keys may differ. For example, KEMs and PKE typically rely on LWE based problems where as signatures tend to use the assumed hardness of the Short Integer Solution problem, or SIS. In this case, the distribution of the secret key is not so important (provided there is enough entropy) but the size now plays a critical role. Balancing these two requirements is a new challenge for ARKG in the lattice-based setting.

1.3. Contributions

In this work, we give a new generic construction of ARKG allowing for the use of the primitive with lattice-based public key cryptosystems. We aim to capture a wide range of lattice protocols despite differing requirements on their keys. Lattice schemes based on Short Integer Solution (SIS) [13] often use bounded short vectors for their keys, whereas KEMs/PKEs based on Learning with Errors (LWE) [12] typically require keys with a specific and proper distribution (typically discretised Gaussian). Our ARKG is the first to support LWE and SIS based-schemes, encompassing a wide range of potential applications. We build this general lattice ARKG (LARKG) using Split KEM (sKEM) [14] functionality to achieve the required asynchronicity.

Whilst maintaining ARKG’s original security and privacy properties, we overcome the inherent challenges with lattice assumptions’ noisiness. The need to select cryptographically hard parameters for composed protocols means consideration must be given to the correctness versus security trade-off—this can be challenging for signature schemes using SIS because, as the norm of the key increases, security is weakened. Similarly, PKE typically requires well distributed secret and error vectors. To address this, we introduce a new definition of correctness, that we call τ -correctness that captures the probability that some derived keys may fail or be too large. This parameter is related to the *repetition rate* of the rejection sampling techniques inherent in our generic construction. These parameters can be varied to trade-off correctness against efficiency.

Additionally, we provide two instantiations of our LARKG based on Frodo [9] and Kyber [10]. Since Kyber has previously not known to have sKEM functionality, we also state how to construct this and give proof it has security properties of an sKEM. We give a discussion on potential

applications in lattice-based KEMs, signatures, and encryption schemes. In particular, we discuss integration with Dilithium [6] and Kyber, which have recently been selected for standardisation by the NIST PQC process. We provide benchmarking of both of our instantiations using a publicly available implementation in Python. This performance analysis shows that our implementations achieve practical performance when generating derived key pairs, compared to the standard key generation algorithms.

1.4. Organisation

We start in Section 2 by recalling the ARKG primitive, with a new correctness definition to make it compatible with lattices. In Section 3, we define the necessary preliminaries, building blocks and security assumptions before we present our general lattice-based ARKG construction and security analysis. Section 4 presents two instantiations of LARKG and discusses its usage in lattice-based KEMs, PKE, and signatures. In this section we also show how to construct an sKEM based on Kyber, and give proof it meets standard security properties of an sKEM. We conclude this section with implementation details and benchmarks of our protocol. Our conclusion and a future work discussion is presented in Section 5.

2. Asynchronous Remote Key Generation

In this section we recall Asynchronous Remote Key Generation based on the model introduced by Frymann et al [1]. We give the syntax and security properties, and modify the definition of correctness to account for some failure probability.

2.1. ARKG Model

For the key pair (sk, pk) and credential $cred$, ARKG allows arbitrary public keys pk' to be derived from an original pk , with corresponding sk' being calculated at a later time—requiring private key sk .

Definition 1 (ARKG [1]). *The asynchronous remote key generation scheme ARKG $:=$ (Setup, KeyGen, DerivePK, DeriveSK, Check) consists of the following five algorithms:*

- *Setup(λ) generates and outputs public parameters pp of the scheme for the security parameter $\lambda \in \mathbb{N}$. We assume pp is taken as implicit input to all algorithms.*
- *KeyGen() computes and returns a key pair (sk, pk) .*
- *DerivePK(pk, aux) probabilistically returns a new public key pk' together with the link $cred$ between pk and pk' , for the inputs pp, pk and auxiliary data aux . The input aux is always required but may be empty.*
- *DeriveSK($sk, cred$), computes and outputs either the new private key sk' , corresponding to the public key pk' using $cred$, or \perp .*
- *Check(sk', pk'), on input (sk', pk') , returns 1 if (sk', pk') forms a valid private-public key pair, where sk' is the corresponding private key to public key pk' , otherwise 0.*

Correctness. We modify the correctness property from Frymann et al. [1] to allow for some probability of failure that DeriveSK does not output a valid derived secret

key after correct execution of a corresponding DerivePK. We capture this failure probability with the correctness parameter τ . Note that our definition is consistent with Frymann et al. when $\tau \rightarrow \infty$. In this case we say the scheme has perfect correctness.

An ARKG scheme is τ -correct if, $\forall \lambda \in \mathbb{N}$, $\text{pp} \leftarrow \text{Setup}(\lambda)$, the probability $\Pr[(\text{Check}(\text{pp}, \text{sk}', \text{pk}') = 1)] \geq 1 - 2^{-\tau}$ given

$$\begin{aligned} (\text{sk}, \text{pk}) &\leftarrow \text{KeyGen}(\text{pp}); \\ (\text{pk}', \text{cred}) &\leftarrow \text{DerivePK}(\text{pp}, \text{pk}, \cdot); \\ \text{sk}' &\leftarrow \text{DeriveSK}(\text{pp}, \text{sk}, \text{cred}). \end{aligned}$$

2.1.1. Security Properties. We recall the security properties for ARKG. The first, PK-unlinkability, is an anonymity property whereas unforgeability is provided by SK-security. This property comes in 4 flavours to capture a wide range of adversarial capabilities. Furthermore, we introduce generalised security definitions to accommodate the possibility that DeriveSK returns an error \perp instead of a derived secret key. This can be seen in the challenge and private key oracles. *Adversaries and Oracles*

An adversary \mathcal{A} , used in our security experiments, is modelled as a probabilistic polynomial time (PPT) algorithm. The adversary \mathcal{A} may make a polynomial number of queries to the following oracles:

- Derived public key oracle $\mathcal{O}_{\text{pk}'}(\text{pk}, \cdot)$: $\mathcal{O}_{\text{pk}'}$ is parameterised with public key pk . This oracle returns the result of calling $\text{DerivePK}(\text{pp}, \text{pk}, \text{aux})$ on input aux . It records the resulting $(\text{pk}', \text{cred})$ in PKList: $\text{PKList} \leftarrow \text{PKList} \cup (\text{pk}', \text{cred})$. PKList is initialised as $\text{PKList} \leftarrow \emptyset$.
- Challenge oracle $\mathcal{O}_{\text{pk}'}^b(b, \text{sk}_0, \text{pk}_0)$: $\mathcal{O}_{\text{pk}'}$ is parameterised with a bit b and fixed key pair $(\text{sk}_0, \text{pk}_0)$, and takes no inputs. If $b = 0$, the key-pair (sk', pk') derived using the initial pk_0 using DerivePK and DeriveSK . If DeriveSK returns \perp then a DerivePK is invoked again to create a new derived key, and this is repeated until success. If $b = 1$, then a freshly-generated key pair sampled from a distribution \mathcal{D} .
- Private key oracle $\mathcal{O}_{\text{sk}'}(\text{sk}, \cdot)$: on input cred , where $(\cdot, \text{cred}) \in \text{PKList}$, $\mathcal{O}_{\text{sk}'}$ outputs the result of $\text{DeriveSK}(\text{pp}, \text{sk}, \text{cred})$ and updates $\text{SKList} \leftarrow \text{SKList} \cup \text{cred}$ if the output is not \perp . SKList is initialised as $\text{SKList} \leftarrow \emptyset$. If $(\cdot, \text{cred}) \notin \text{PKList}$, the oracle aborts, otherwise it returns sk' without giving access to sk .

Remark 1. We note that our security model does capture the possibility of timing attacks that might be present due to the variable run time of the challenge oracle. However, in the envisioned use case, such a timing attack would imply adversarial access to back up authenticators – which is beyond the security model of ARKG.

SK-security. This security property ensures that an adversary \mathcal{A} cannot derive a valid key pair $(\text{sk}^*, \text{pk}^*)$ along with corresponding cred^* , for a fixed challenge key pk . We recall four variants of private-key security defined in Fryman et al., modelled using the experiment $\text{Exp}_{\text{LARKG}, \mathcal{A}}^{\text{ks}}(\lambda)$ in Figure 1 with $\text{ks} \in \{\text{mwKS}, \text{hwKS}, \text{msKS}, \text{hsKS}\}$. Adversary \mathcal{A} is always given access to $\mathcal{O}_{\text{pk}'}$ and must find a $(\text{sk}^*, \text{pk}^*, \text{cred}^*)$ triple for a provided pk . We have

altered the original definition to include a check that the derived secret key sk' is not \perp , this is to account for the fact that there is some probability that DeriveSK fails. The malicious (m) and honest (h) variants result from the omission or presence of the PKList check on line 8, respectively, which ensures that the triple is for an honestly-generated pk (modelled using $\mathcal{O}_{\text{pk}'}$) if present. The weak (W) and strong (S) variants depend on whether \mathcal{A} has access to the private key derivation oracle $\mathcal{O}_{\text{sk}'}$. If \mathcal{A} has access to $\mathcal{O}_{\text{sk}'}$, trivially querying it with cred^* is prevented through the SKList check on line 7. It was shown by Frymann et al. that strong implies weak security, and malicious implies honest security. We defer to the original work for the reduction.

Definition 2 (SK-security). *An ARKG scheme provides private-key security, or ks-security with $\text{ks} \in \{\text{mwKS}, \text{hwKS}, \text{msKS}, \text{hsKS}\}$, if the following advantage is negligible in λ :*

$$\text{Adv}_{\text{LARKG}, \mathcal{A}}^{\text{ks}}(\lambda) := \Pr[\text{Exp}_{\text{LARKG}, \mathcal{A}}^{\text{ks}}(\lambda) = 1].$$

PK-unlinkability. This property ensures that derived key pairs cannot be distinguished from a sample of a distribution \mathcal{D} , which also implies an adversary cannot link a derived public key to a long-term public key or second derived public key. This property is formally defined in $\text{Exp}_{\text{LARKG}, \mathcal{A}}^{\text{pku}}(\lambda)$. Intuitively, the game chooses a bit $b \in \{0, 1\}$ and generates a key pair $(\text{sk}_0, \text{pk}_0)$. \mathcal{A} is given access to oracle $\mathcal{O}_{\text{pk}'}^b$, public parameters pp , and pk_0 . When called, $\mathcal{O}_{\text{pk}'}^b$ returns a derived key pair (sk', pk') , which is derived from pk_0 if $b = 0$, otherwise, for $b = 1$, it samples and returns key pair (sk', pk') according to a distribution \mathcal{D} . It is able to corrupt any derived key, for which corresponding sk' is output. The adversary \mathcal{A} wins the game if it is able to determine whether $\mathcal{O}_{\text{pk}'}^b$ is instantiated with $b = 0$ or $b = 1$.

Definition 3 (PK-unlinkability). *An ARKG scheme provides PK-unlinkability if the following advantage is negligible in λ :*

$$\text{Adv}_{\text{LARKG}, \mathcal{A}}^{\text{pku}}(\lambda) := \left| \Pr[\text{Exp}_{\text{LARKG}, \mathcal{A}}^{\text{pku}}(\lambda) = 1] - \frac{1}{2} \right|.$$

Remark 2. As previously mentioned, Frymann et al. [1] give a construction for discrete-log key pairs. They achieve weak variants of the scheme (i.e. *mwKS* and *hwKS*) under the discrete log assumption. For stronger flavours, *msKS* and *hsKS*, security is reduced to the *snPRF-ODH* assumption by Brendel et al. [15], which they introduce to study the security of *TLS1.3*.

3. General Lattice-based ARKG Construction from Split KEMS

In this section we define sKEMs and the other necessary building blocks for our general ARKG construction. We present our construction in Figure 5 and prove its security.

3.1. Assumptions and Building Blocks

We now introduce the technical background, hardness assumptions and the building blocks used in this work.

$\text{Exp}_{\text{LARKG},\mathcal{A}}^{\text{ks}}(\lambda)$

```

1: pp ← Setup(λ)
2: (sk, pk) ← KeyGen(pp)
3: (sk*, pk*, cred*) ←  $\mathcal{A}^{\mathcal{O}_{\text{pk}'}, \overline{\mathcal{O}_{\text{sk}'}}}$ (pp, pk)
4: sk' ← DeriveSK(pp, sk, cred*)
5: return Check(sk*, pk*)  $\stackrel{?}{=} 1$ 
6:   ∧ Check(sk', pk*)  $\stackrel{?}{=} 1$ 
7:   ∧ sk' ≠ ⊥
8:   ∧ cred* ∉ SKList
9:   ∧ (pk*, cred*) ∈ PKList

```

(a) SK-security experiment.

$\text{Exp}_{\text{LARKG},\mathcal{A}}^{\text{pku}}(\lambda)$

```

1: pp ← Setup(1λ)
2: (sk0, pk0) ← KeyGen(pp)
3: b ← $\mathcal{S}$  {0, 1}
4: b' ←  $\mathcal{A}^{\mathcal{O}_{\text{pk}'}}(\text{pp}, \text{pk}_0)$ 
5: return b  $\stackrel{?}{=} b'$ 

```

(b) PK-unlinkability experiment.

Figure 1: Security experiments for LARKG. The boxes denote the four variants of the $\text{ks} \in \{\text{mwKS}, \text{hwKS}, \text{msKS}, \text{hsKS}\}$ experiment. Presence of the dashed boxes gives the strong variants of ks (msKS and hsKS), the presence of the dotted box gives the honest variants (hwKS and hsKS), and the exclusion of all boxes gives mwKS .

On notation, we use $\chi_{\alpha, \gamma}^{n \times m}$ to denote an $n \times m$ matrix that is element-wise sampled from a distribution χ with parameters (α, γ) . In this work, we utilise Gaussian distributions $\chi_{\alpha, \gamma}$ where α is the standard deviation and γ is the centre. If we omit γ then it is understood to be 0. We use $\chi(X)$ to denote the probability of X occurring under the distribution χ . We reserve the notation $\mathcal{U}[a, b]$ to indicate uniform sampling from the closed interval $[a, b]$. Matrices are denoted in bold capital letters, e.g. \mathbf{C} , and vectors are notated as bold lowercase letters such as \mathbf{x} . In this work we consider a polynomial ring $\mathcal{R}_q := \mathbb{Z}_q[X]/(X^n + 1)$ where $n = 2^{\hat{n}-1}$ and \hat{n} is the $2^{\hat{n}}$ -th cyclotomic polynomial. We use $x \leftarrow_{\mathcal{S}} X$ to denote that x is uniformly sampled from X , whereas Gaussian sampling is denoted $x \leftarrow_G X$.

Definition 4. *Decisional Learning with Errors (LWE).* For LWE parameters n, q, χ and a matrix \mathbf{A} sampled uniformly from $\mathbb{Z}_q^{n \times m}$, sample $\mathbf{e} \leftarrow \chi$ and a vector $\mathbf{s} \in \mathbb{Z}_q^n$. In the decisional variant of LWE, an adversary must decide, for a tuple (\mathbf{A}, \mathbf{b}) , whether $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \pmod q$ or \mathbf{b} is a uniform sample from \mathbb{Z}_q^m . This can be extended to the Matrix-LWE by taking $\mathbf{S}, \mathbf{E} \leftarrow_G \chi^{m \times k}$.

We do not distinguish between these two variants since hardness of LWE implies hardness of Matrix-LWE.

Definition 5. *Decisional Module Learning with Errors (MLWE)* [16]. For parameters m, k, χ the MLWE problem is stated for a matrix \mathbf{A} sampled uniformly from $\mathcal{R}_q^{k \times k}$ and $\mathbf{e}, \mathbf{s} \leftarrow \chi_{\eta}^k$. The problem requires an adversary to decide, for a tuple (\mathbf{A}, \mathbf{b}) , whether $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$ or \mathbf{b} is a uniform sample from \mathcal{R}_q^k . MLWE becomes Ring-LWE (RLWE) when the dimension of the ring is set to 1, i.e.

when the parameter $k = 1$.

Definition 6. *Short Integer Solution (SIS)* [13]. The SIS problem $\text{SIS}_{n, m, q, \beta}$ requires an adversary, given a uniformly sampled matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, to find a non-zero vector $\mathbf{z} \in \mathbb{Z}_q^m$ such that $\|\mathbf{z}\| \leq \beta$ and $\mathbf{A}\mathbf{z} = \mathbf{0} \pmod q$.

Definition 7. *Module Short Integer Solution (MSIS)* [16]. The MSIS problem $\text{MSIS}_{n, m, q, \beta}$ requires an adversary who, given a uniformly sampled matrix $\mathbf{A} \in \mathcal{R}_q^{k \times m}$ to find a non-zero vector $\mathbf{z} \in \mathcal{R}_q^k$ such that $\|\mathbf{z}\| \leq \beta$ and $\mathbf{A}\mathbf{z} = \mathbf{0} \in \mathcal{R}_q^k$. MSIS becomes Ring-SIS (RSIS) when the parameter k is set to 1.

Rejection Sampling. This is a technique that ensures we can control the shape of distributions, first used in lattice cryptography by Lyubashevsky [17]. Let f and g be probability distributions and $M \in \mathbb{R}$ a scalar constant, such that for all x we have $f(x) \leq Mg(x)$. Then, one can sample y from g and output it with probability $f(y)/Mg(y)$, and the resulting distribution of output values is precisely f . The repetition rate M is the expected amount of time needed to output a sample.

Split Key Encapsulation Mechanism. The fundamental building block for our ARKG protocol is a *Split Key Encapsulation Mechanism*, or sKEM. It is related to standard notion of a KEM, where the encapsulation algorithm can be divided into two phases for key generation and shared-key computation. This primitive was introduced by Brendel et al. [14] and used to propose quantum secure key exchange for the Signal protocol [18]. We recall their model and security definitions.

Definition 8 (sKEM [14]). *An sKEM consists of four algorithms KeyGenEnc, KeyGenDec, sEncaps and sDecaps, where KeyGenEnc and sEncaps are executed by the encapsulator, and KeyGenDec and sDecaps by the decapsulator.*

- KeyGenEnc is the probabilistic key generation algorithm for the encapsulator $(E, e) \leftarrow \text{KeyGenEnc}(1^\lambda)$.
- KeyGenDec is the probabilistic key generation algorithm for the decapsulator $(D, d) \leftarrow \text{KeyGenDec}(1^\lambda)$.
- sEncaps($e; D$) is a probabilistic algorithm executed by the encapsulator. It takes as input $e \in \mathcal{K}_{\text{enc}}$, the secret key of the encapsulator, and $D \in \mathcal{K}_{\text{dec}}$, the public key of the decapsulator. Algorithm sEncaps then outputs the shared secret $K \in \mathcal{K}$ along with its encapsulation $c \in \mathcal{C}$, sometimes referred to as the ciphertext.
- sDecaps($d; E; c$) is a deterministic algorithm executed by the decapsulator. On input a ciphertext c , the decapsulator's secret key d , and encapsulator's public key E , it outputs either the decapsulation K of c or \perp , if it fails.

Security Properties. sKEMs can offer a range of indistinguishability properties, that are denoted as lr-IND-CCA where the parameters l, r are in the set $\{n, s, m\}$. Here, l indicates whether the adversary is allowed to make no ($l = n$), a single ($l = s$), or polynomially many ($l = m$) queries to the decapsulation oracle $\mathcal{O}_{\text{sDecaps}}$. Analogously, r indicates the number of queries the adversary is allowed to make to the encapsulation oracle $\mathcal{O}_{\text{sEncaps}}$. The case that $r = s$ is excluded since the adversary cannot make

the encapsulator encapsulate only once more under the secret key e used for challenge generation. The key pair of the encapsulator is used either solely for the challenge generation or may be queried up to a bound which is $\text{poly}(\lambda)$. The ‘many’ case models key re-use attacks and has historically been a challenge for post-quantum cryptography. As noted by Brendel et al. [14], nn-IND-CCA security corresponds to the notion of IND-CPA security for KEMs, i.e. the adversary is only able to learn the public keys and challenge ciphertext and key, but remains passive. It is also restricted from querying the decapsulation oracle.

Definition 9. Let $\text{sKEM} = (\text{KeyGenEnc}; \text{KeyGenDec}; \text{sEncaps}; \text{sDecaps})$ be a split KEM with key space \mathcal{K} . Let $l, r \in \{n, s, m\}$. We say sKEM provides lr -indistinguishability under chosen-ciphertext attacks, or for short, sKEM is lr -IND-CCA secure, if for every PPT adversary \mathcal{A} the following advantage is negligible in λ .

$$\text{Adv}_{\text{sKEM}, \mathcal{A}}^{lr\text{-IND-CCA}}(\lambda) := \left| \Pr \left[\text{Exp}_{\text{sKEM}, \mathcal{A}}^{lr\text{-IND-CCA}}(\lambda) = 1 \right] - \frac{1}{2} \right|$$

In Figure 3, we recall a generic lattice-based construction from Brendel et al. [14], presented in Figure 3. Many lattice KEM candidates fit this generic 2-phase structure of generating ephemeral keys and then combing and reconciling a shared key [9], [19]–[21].

Key Derivation Function (KDF) [22] A key generation function $\text{KDF}(k, l)$ takes source key k and label l and returns a new key k' . It is secure if the advantage $\text{Adv}_{\mathcal{A}}^{\text{KDF}}(\lambda)$ is negligible in λ for a PPT adversary \mathcal{A} to distinguish outputs of the KDF from random bitstrings of the same length. In this work we consider uniform outputs of KDF_1 with small Gaussian outputs (parameterised by χ) and KDF_2 with uniform outputs.

Public Key Encryption (PKE) A PKE scheme consists of three algorithms KeyGen , Enc and Dec . The former outputs a key-pair based on a security parameters λ . The encryption algorithm takes as input a key k and plaintext m and returns a ciphertext c . Decrypt then outputs the plaintext m on input of a ciphertext and corresponding secret key. The game is formally defined in Figure 4.

It is IND-CPA secure if the advantage $\text{Adv}_{\mathcal{A}}^{\text{PKE}}(\lambda)$ is negligible in λ for a PPT adversary \mathcal{A} to decide which of two chosen messages a ciphertext contains. Precisely, PKE is IND-CPA secure if the following advantage is negligible in λ .

$$\text{Adv}_{\mathcal{A}, \text{PKE}}^{\text{IND-CPA}}(\lambda) := \left| \Pr \left[\text{Exp}_{\text{PKE}}^{\text{IND-CPA}}(\lambda) = 1 \right] - \frac{1}{2} \right|.$$

3.2. Our LARKG scheme based on Split KEMs

We now present our ARKG construction based on lattices. In our construction, key pairs (pk, sk) are given by the vectors of matrices $(\mathbf{A}, \mathbf{B} = \mathbf{A}\mathbf{S} + \mathbf{E}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^{n \times k}$ which form an LWE sample. The derived key pair (pk', sk') , i.e. the outputs of DerivePK and DeriveSK respectively, are given by \mathbf{P} and \mathbf{S}'' also form an LWE sample. For this construction, we focus on plain LWE and the simple case where the secret and error matrices are taken from the same distribution χ . This is for simplicity in presentation, but we

emphasise that our techniques readily extend to Ring-LWE and Module-LWE settings. Furthermore, the restricted distribution is not a requirement of our scheme and can be chosen independently according to the application.

The core building block of our construction is the sKEM primitive. The main benefit of using this over standard KEMs is that the flavours of security that allow for oracle queries capture the notion of key reuse. As our security theorems will show, security in the presence of key reuse allows us to achieve all, including the strongest, security flavours of SK security for ARKG.

Intuitively, we use a split KEM to create a shared secret K with recovery information c . The format of sKEMs allows us to assume the structure of the KEM mechanism, in particular, that the encapsulation and decapsulation algorithms both take in public and secret keys. This enables an asynchronous approach to decapsulating a shared key, which forms the basis of our DerivePK and DeriveSK algorithms. The KDFs are used in one case to create a shared secret with appropriate distribution, i.e. a small Gaussian distribution χ . The other allows DeriveSK to abort since an incorrect long-term key or auxiliary information was used, and thus Check would fail.

Our protocol makes use of rejection sampling to ensure that output of DeriveSK follows a desired distribution. We note that flooding techniques, e.g. [23], could also be implemented here but at cost of increased parameter size. Instead, rejection sampling allows us to keep the parameters fixed and instead vary the correctness probability with repetition rate. That is, the protocol will need to be executed an expected M times to obtain a derived secret key.

We give a high-level description of the algorithms and present the full algorithms in Figure 5.

Setup This algorithm takes as input a security parameter λ and outputs public parameters pp that contains LWE parameters n, q, χ and $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$.

KeyGen On input of public parameters, this algorithm outputs a key-pair as (sk, pk) such that $\text{sk} = \mathbf{S}$ is a small gaussian and $\text{pk} = \mathbf{B} = \mathbf{A} \cdot \mathbf{S} + \mathbf{E}$. That is, (\mathbf{A}, \mathbf{B}) form an LWE sample.

DerivePK This algorithm generates a derived public key pk' from input of a public key $\text{pk} = \mathbf{B}$. Intuitively, it samples an ephemeral key pair as $(\text{sk}', \text{pk}') = (\mathbf{S}', \mathbf{P} = \mathbf{A} \cdot \mathbf{S}' + \mathbf{E}')$. It computes a shared key K by using a split KEM with inputs $(\text{sk}', \text{pk}) = (\mathbf{S}', \mathbf{B})$, and uses this to derive a credential key \mathbf{K} . The derived public key is computed as $\text{pk}' = \mathbf{B} + \mathbf{A} \cdot \mathbf{K} + \mathbf{E}'$. This step also outputs ciphertext c . A second key derivation function KDF_2 is used to generate μ over the inputs \mathbf{P}, aux using KDF_2 . It builds the recovery information as $\text{cred} = (\mathbf{D}, c, \mu)$.

DeriveSK This algorithm takes as input sk and cred . It recomputes the key \mathbf{K} by executing the decapsulation algorithm on its sk and with \mathbf{P} . A candidate $\mu^* \leftarrow \text{KDF}_2(K)$ is derived from K , and it is checked that $\mu^* \stackrel{?}{=} \mu$. If this holds, then it proceeds to recover its derived secret key as $\mathbf{K} \leftarrow \text{KDF}_1(K)$ and setting $\text{sk}' = \mathbf{S} + \mathbf{K}$. Finally, rejection sampling is performed, so that the key is only output with probability $\chi_{b, \mathbf{S}}(\mathbf{S}'') / (M \cdot \chi_a(\mathbf{S}))$, for some M .

Check Takes as input a candidate key pair $(\text{sk}, \text{pk}) = (\mathbf{S}, \mathbf{P})$ and returns 1 if $\|\mathbf{P} - \mathbf{A} \cdot \mathbf{S}\| \leq \beta$, where the

$\text{Exp}_{\text{sKEM}, \mathcal{A}}^{\text{lr-IND-CCA}}(\lambda)$	$\mathcal{O}_{\text{sDecaps}}(E', c)$	$\mathcal{O}_{\text{sEncaps}}(D')$
1: $n_l, n_r \leftarrow 0$	1: if $n_l \geq l$	1: if $n_r \geq r$
2: $(D, d) \leftarrow \text{KeyGenDec}(\text{pp})$	2: return \perp	2: return \perp
3: $(E, e) \leftarrow \text{KeyGenEnc}(\text{pp})$	3: $n_l = n_l + 1$	3: $n_r = n_r + 1$
4: $(c^*, K_0^*) \leftarrow \text{sEncaps}(e, D)$	4: if $(E', c) \stackrel{?}{=} (E, c^*)$	4: $(c, K) \leftarrow \text{sEncaps}(e, D')$
5: $K_1^* \leftarrow \mathcal{K}$	5: return \perp	5: if $(D', c) \stackrel{?}{=} (D, c^*)$
6: $b \leftarrow \mathcal{S}\{0, 1\}$	6: $L \leftarrow \text{sDecaps}(d, E', c)$	6: return \perp
7: $b' \leftarrow \mathcal{A}^{\mathcal{O}_{\text{sDecaps}}, \mathcal{O}_{\text{sEncaps}}}(D, E, c^*, K_b^*)$	7: return K	7: else
8: return $b' \stackrel{?}{=} b$		8: return (c, K)

Figure 2: Security experiments and oracles for sKEMs. The values of n_l and n_r are set to as either 0, 1 or k (which is $\text{poly}(\lambda)$) according to the $\{n, s, m\}$ -IND-CPA security, respectively.

KeyGenEnc(pp)	KeyGenDec(pp)	sEncaps(pp, pk, sk')	sDecaps(pp, pk', sk, c)
1: $\mathbf{S}' \leftarrow \mathcal{S}\chi^{m \times k}, \mathbf{E}' \leftarrow \mathcal{S}\chi^{n \times k}$	1: $\mathbf{S} \leftarrow \mathcal{S}\chi^{m \times k}, \mathbf{E} \leftarrow \mathcal{S}\chi^{n \times k}$	1: $\mathbf{E}'' \leftarrow \mathcal{S}\chi^{n \times k}$	1: $\mathbf{V}' \leftarrow \mathbf{B}'\mathbf{S}$
2: $\mathbf{B}' \leftarrow \mathbf{A}\mathbf{S}' + \mathbf{E}'$	2: $\mathbf{B} \leftarrow \mathbf{S}\mathbf{A} + \mathbf{E}$	2: $\mathbf{V} \leftarrow \mathbf{B}\mathbf{S}' + \mathbf{E}''$	2: $K \leftarrow \text{Rec}(\mathbf{V}', c)$
3: return $(\text{sk}', \text{pk}') := (\mathbf{S}', \mathbf{B}')$	3: return $(\text{sk}, \text{pk}) := (\mathbf{S}, \mathbf{B})$	3: $c \leftarrow \text{HelpRec}(\mathbf{V})$	3: return K
		4: $K \leftarrow \text{Rec}(\mathbf{V}, c)$	
		5: return (K, c)	

Figure 3: Lattice-based sKEM with generalised Rec and HelpRec algorithms by Brendel et al [14].

$\text{Exp}_{\text{PKE}, \mathcal{A}}^{\text{IND-CPA}}(\lambda)$
1: $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(\lambda)$
2: $m_0, m_1 \leftarrow \mathcal{A}_1(\text{pk})$
3: $b \leftarrow \mathcal{S}\{0, 1\}$
4: $c \leftarrow \text{Enc}(\text{sk}, m_b)$
5: $b' \leftarrow \mathcal{A}_2(c)$
6: return $b' \stackrel{?}{=} b$

Figure 4: IND-CPA security experiment for PKE.

distribution χ_a is bounded by β .

3.3. τ -Correctness

We observe that τ -correctness follows from correctness of the split KEM and that \mathbf{P} is well formed:

$$\begin{aligned} \mathbf{P} &= \mathbf{B} + \mathbf{E} + \mathbf{A}\mathbf{K} + \mathbf{E}' = \mathbf{A}\mathbf{S} + \mathbf{A}\mathbf{K} + \mathbf{E} + \mathbf{E}' \\ &= \mathbf{A}(\mathbf{S} + \mathbf{K}) + \mathbf{E}'' = \mathbf{A}\mathbf{S}'' + \mathbf{E}'' \end{aligned}$$

We note that our use of rejection sampling in DeriveSK ensures that the sum $\mathbf{S} + \mathbf{K}$ is statistically close to discrete Gaussian distribution χ_a , i.e. the distribution which \mathcal{S} follows. If this is chosen to be the distribution for a compatible protocol, then the correctness is perfect and $t = \infty$. However, this may require a high repetition rate M . We can trade τ off against M to be most suitable for an application, the exact relationship for which is described by an appropriate choice of parameters and distributions. If we instantiate the error vectors to be sampled from Gaussian distributions as χ_a, χ_b , then the resultant distribution of the derived key \mathbf{S}'' is $\chi_{b, \mathbf{S}}$, i.e. a Gaussian distribution with variance b centred at \mathbf{S} . If χ_a is bounded by a value β_1 , and $\sigma_b = \alpha\beta_1$ for some positive value α , then following [17], we can compute the repetition rate M as follows:

$$M = e^{\frac{12}{\alpha} + \frac{1}{2\alpha^2}}$$

KeyGen(pp)
1: return $(\text{sk}, \text{pk}) \leftarrow \text{sKEM.KeyGenDec}(\text{pp})$
DerivePK(\mathbf{B})
1: $(\mathbf{S}', \mathbf{B}') \leftarrow \text{sKEM.KeyGenEnc}(\text{pp})$
2: $(K, c) \leftarrow \text{sKEM.sEncaps}(\mathbf{B}, \mathbf{S}')$
3: $\mathbf{K} \leftarrow \text{KDF}_1(K)$
4: $\mu \leftarrow \text{KDF}_2(K)$
5: $\mathbf{E}' \leftarrow \mathcal{S}\chi_b^{n \times k}$
6: $\mathbf{P} \leftarrow \mathbf{B} + \mathbf{A}\mathbf{K} + \mathbf{E}' \pmod{q}$
7: return $(\mathbf{P}, \mathbf{B}', c, \mu)$
Check(\mathbf{P}, \mathbf{S}'')
1: return $\ \mathbf{P} - \mathbf{A}\mathbf{S}''\ \leq \beta'$
DeriveSK($\mathbf{B}', \mathbf{S}, \mu, c$)
1: $\mathbf{S}'' \leftarrow \perp$
2: $K \leftarrow \text{sKEM.sDecaps}(\mathbf{B}', \mathbf{S}, c)$
3: $\mathbf{K} \leftarrow \text{KDF}_1(K)$
4: $\mu^* \leftarrow \text{KDF}_2(K)$
5: if $\mu^* \stackrel{?}{=} \mu$ then
6: $\mathbf{S}'' \leftarrow \mathbf{K} + \mathbf{S}$
7: $u \leftarrow \mathcal{U}[0, 1]$
8: if $u < \left(\frac{M \cdot \chi_{b, \mathbf{S}}(\mathbf{S}'')}{\chi_a(\mathbf{S}'')}\right)$, return \mathbf{S}''
9: return \perp

Figure 5: General Lattice-based ARKG from sKEM

If we consider the scenario where the error \mathbf{E}'' and secret \mathbf{S}'' are derived from the distinct distributions, then a repetition rate would need to be computed for each component, denoted M_S and M_E . Then M , the realised repetition rate of the protocol, would be set to $\max\{M_S, M_E\}$.

Since the distribution of the derived secret key \mathbf{S}'' is χ_a , then for an appropriate choice of parameters, it forms a valid LWE instance. For composable protocols that use Gaussian distributed keys, all that is required is to select a and b such that χ_a and χ_b are cryptographically hard distributions. If these conditions are met, then we can use the composability theorem in Frymann et al. [1] to generate keys for protocols that include many of the NIST PQC lattice KEM and encryption schemes such as Frodo [9] and Kyber [10], but also some signatures such as qTESLA [7].

For schemes based on the SIS assumption [13], where the bound β' on the norm plays an important factor in the security of the scheme it may be necessary or desirable to set $\beta' < 2\beta$. However, the wider distribution of the derived key \mathbf{S}'' may be too large and thus Check could fail. Nonetheless, the trade off in lower repetition rate versus a larger value of β' may be desirable if the composed protocol has parameter values chosen to accommodate this. The exact values in these cases are dependent on protocol specific parameters, see Section 4.4 for an example using the lattice signature Dilithium [6].

3.4. Security Analysis

Theorem 1. LARKG satisfies PK-unlinkability if sKEM is nn-IND-CCA secure and LWE_{n,q,χ_a} is hard.

Proof. Game₀: Defined to be the PK-unlinkability experiment $\text{Exp}_{\text{LARKG}}^{\text{pk}}(\lambda)$.

Game₁ is defined as Game₀ with the exception that line 6 of DerivePK is replaced with “ $\mathbf{S}', \mathbf{E} \leftarrow_G \chi_a, \mathbf{B}' \leftarrow \mathbf{A}\mathbf{S}' + \mathbf{E}, \mathbf{P} \leftarrow \mathbf{B}' + \mathbf{A}\mathbf{K} + \mathbf{E}''$ ”. It follows that \mathbf{B} and \mathbf{B}' are statistically indistinguishable from decisional LWE, since rejection sampling ensures that the derived secret key \mathbf{S}'' is distributed according to χ_a . Hence the probability that an adversary wins is preserved.

$$|\Pr[\text{Game}_0 = 1] - \Pr[\text{Game}_1 = 1]| \leq \text{Adv}_{\mathcal{B}_1}^{\text{LWE}_{n,q,\chi_a}}(\lambda)$$

Game₂ is defined as Game₁ with the exception that line 4 of the oracle is replaced with “ $\mathcal{K}^*, \mathbf{c}^* \leftarrow \text{sEncaps}(\mathbf{B}, \mathbf{S}'), \mathbf{V} \leftarrow \mathbf{B}\mathbf{S}' + \mathbf{E}'', K_0^* \leftarrow \$_{\mathcal{K}}$. That is, we replace the shared-key K_0^* with a uniformly random sample from the key space \mathcal{K} . We argue that the nn-IND-CCA property of sKEM implies that the difference in success probability for the adversary is negligible between games Game₁ and Game₀. To see this, we construct an adversary \mathcal{B} against the nn-IND-CCA property of sKEM using Game₂ as a distinguisher. Precisely, \mathcal{A} plays the role of challenger against \mathcal{B} , and initiates the experiment as described in Figure 2, where it sets the public parameters \mathbf{A}, \mathbf{B} and \mathbf{B}' to be those from its nn-IND-CCA experiment. We note that \mathcal{A} does not have access to any oracles in aid of it, other than the challenge oracle. Instead of executing the oracle as described, it uses the values it has obtained from its own game. Since all values are known to \mathcal{B} , it perfectly

simulates the game for \mathcal{A} , in particular, if nn-IND-CCA has been instantiated with challenge bit $b = 0$ then it coincides with Game₁ and if $b = 1$ then it coincides with Game₂. Therefore, if \mathcal{A} can distinguish between Game₂ and Game₁, \mathcal{B} can construct a correct response to the nn-IND-CCA experiment, that succeeds if \mathcal{A} succeeds.

$$|\Pr[\text{Game}_2 = 1] - \Pr[\text{Game}_1 = 1]| \leq \text{Adv}_{\mathcal{B}_2, \text{sKEM}}^{\text{nn-IND-CCA}}(\lambda).$$

Now, the output of the challenge oracle is independent of the challenge bit b , and thus the probability that \mathcal{A} outputs a correct responses is $\frac{1}{2}$.

$$\Pr[\text{Game}_2 = 1] = \frac{1}{2} \implies \text{Adv}_{\mathcal{B}_2, \text{sKEM}}^{\text{Game}_2}(\lambda) = 0.$$

The claimed results of the theorem then follow from Game₀ to Game₂, and the observation that this argument requires KDF_2 to be secure. We conclude the advantage of an adversary against $\text{Exp}_{\text{LARKG}}^{\text{pk}}(\lambda)$ is bound by the advantage of an adversary against nn-IND-CCA of sKEM.

$$\begin{aligned} \text{Adv}_{\text{LARKG}}^{\text{pk}}(\lambda) &\leq \text{Adv}_{\mathcal{B}_3}^{\text{KDF}_2}(\lambda) + \text{Adv}_{\mathcal{B}_2, \text{sKEM}}^{\text{nn-IND-CCA}}(\lambda) \\ &\quad + \text{Adv}_{\mathcal{B}_1}^{\text{LWE}_{n,q,\chi_a}}(\lambda) \end{aligned}$$

Hence we have shown that LARKG has PK-unlinkability if LWE is hard and sKEM is nn-IND-CCA secure. \square

Theorem 2. LARKG satisfies honest strong SK-secrecy if LWE_{n,q,χ_a} is hard.

Proof. We directly embed an LWE challenge into the protocol as follows. Line 2 of the key generation algorithm in the experiment is replaced with “ $(\text{pk}, \text{sk}) \leftarrow (\mathbf{B}^*, \perp)$ ” where $(\mathbf{A}^*, \mathbf{B}^*)$ is an LWE challenge. Since \mathcal{A} has access to no oracles, \mathcal{A} can simulate the experiment to \mathcal{B} . After some time, \mathcal{B} outputs a forgery \mathbf{S}'' . \mathcal{A} then computes $\mathbf{S}^* = \mathbf{S}'' - \mathbf{K}$, which it can do so since \mathbf{K} was computed by \mathcal{A} during the DerivePK oracle call that created \mathbf{P} . If $\|\mathbf{P} - \mathbf{A}\mathbf{S}^*\|_{\infty} = \|\mathbf{E}^*\|_{\infty} \leq 2\beta$ then \mathcal{A} knows that the challenge \mathbf{B}^* is an LWE sample, and thus wins the decisional LWE_{n,q,χ_a} game, else it was a uniformly sampled matrix. In either case, if \mathcal{A} wins its game against HSKS, then \mathcal{B} wins its LWE game. This we have:

$$\text{Adv}_{\text{LARKG}}^{\text{hsk}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{LWE}_{n,q,\chi_a}}(\lambda).$$

We conclude the advantage of an adversary against SK-security is bound by the advantage of an adversary against LWE, which proves the theorem. \square

Theorem 3. LARKG satisfies malicious strong SK-secrecy if sKEM is sn-IND-CCA secure and LWE_{n,q,χ_a} is hard.

Proof. Let Game₀ be defined as the MSKS experiment in Figure 1. Then, Game₁ is defined as Game₀ with the exception that line 7 of DerivePK is replaced with “ $(\hat{\mathbf{S}}, \hat{\mathbf{B}}) \leftarrow \text{LARKG.KeyGen}(), \mathbf{P} \leftarrow \hat{\mathbf{B}} + \mathbf{A}\mathbf{K} + \mathbf{E}' \bmod q$. This change is undetectable by the adversary since $\mathbf{A}\mathbf{K} + \mathbf{E}'$ is uniformly distributed by LWE. Thus \mathbf{P} is uniformly distributed in both Game₀ and in Game₁ and hence the winning probability of \mathcal{A} is unchanged.

$$|\Pr[\text{Game}_0 = 1] - \Pr[\text{Game}_1 = 1]| \leq \text{Adv}_{\mathcal{B}_1}^{\text{LWE}_{n,q,\chi_a}}(\lambda)$$

We now argue that \mathcal{A} has negligible advantage in Game₁ if the sKEM is sn-IND-CCA secure. To see this, let \mathcal{B} set up the challenge as described in Game₁. It obtains challenge

$(\mathbf{B}, \mathbf{B}^*, c, K)$ and embeds \mathbf{B}^* into the LARKG keys in line 2 of the experiment, i.e. replace with (\perp, \mathbf{B}^*) . We note that \mathcal{B} is able to answer all oracle queries made by \mathcal{A} since it uses the public challenge key \mathbf{B}^* to compute the sKEM output $\mathbf{K} \leftarrow \text{sKEM.KeyGenEnc}(\mathbf{B}^*, \mathbf{E})$ (where \mathbf{E} is the ephemeral key created by \mathcal{B} in the DerivePK oracle). Eventually, the adversary outputs a forgery of the form $(\mathbf{S}^*, \mathbf{P}^*, (c^*, \mathbf{B}^*))$. It extracts the reconciliation data c^* from aux^* and uses this in its single oracle query. It receives back the corresponding shared key K which it used to compute \mathbf{K} . \mathcal{A} then constructs the response as $\mathbf{S}^* \leftarrow \mathbf{S}'' - \mathbf{K}$. It is trivial for \mathcal{B} to then compute $\mathbf{K} \leftarrow \text{KDF}_2(K)$. If $\mathbf{K} \stackrel{?}{=} \mathbf{K}^*$ then set $b = 0$, otherwise set $b = 1$. The simulation only fails if \mathcal{A} queries \mathbf{B}^* which happens with low probability q_o/p^n . Thus the advantage of \mathcal{A} against Game_1 is bound by \mathcal{B} against sn-IND-CCA of sKEM.

$$\text{Adv}_{\text{LARKG}}^{\text{msks}}(\lambda) \leq \text{Adv}_{\mathcal{B}_2, \text{sKEM}}^{\text{sn-IND-CCA}}(\lambda) + \text{Adv}_{\mathcal{B}_1}^{\text{LWE}_{n,q,\chi_a}}(\lambda).$$

We conclude that LARKG has malicious key secrecy if sKEM is sn-IND-CCA secure and LWE_{n,q,χ_a} is hard. \square

4. LARKG Instantiations, Usage and Performance

In this section we instantiate our generic scheme with the sKEMs based on passively-secure versions of Frodo and Kyber. That is, we do not perform the FO transforms and thus are limited to IND-CPA security. Before we can use our general results for (IND-CPA-secure) Kyber, we first show that Kyber can be used to construct an nn-IND-CCA secure sKEM. We discuss composability with other cryptographic primitives and give benchmark timings for our implementation.

The constructions for sKEM by Brendel et al. [14] are limited to providing nn-IND-CCA security, which is also the security property we achieve for our Kyber-based approach. This is inherent in both designs since they are based on KEMs that do not have adaptive security. Indeed, as specified in their NIST submissions, both Frodo and Kyber achieve IND-CCA security via use of variants of the FO-transform [24], [25]. However, in the construction by Brendel et al. and in this work, we use the IND-CPA versions. This limits our LARKG constructions to honest weak and honest strong key security only. In the discrete-log setting, a similar restriction is observed where a stronger sn-PRF-ODH assumption [14] is required to achieve malicious security [1]. Nonetheless, malicious security may not be required—since it models the instance an attacker can perform (potentially arbitrary) secret key derivation queries, which may be too strong for many applications. This is the case for the original ARKG-based recovery [1] and delegation mechanisms [3] for WebAuthn, where honest strong key secrecy is sufficient.

4.1. LARKG Instantiation from FrodoKEM

We instantiate our protocol with passively-secure FrodoKEM, i.e. with IND-CPA security, which was shown to be an nn-IND-CCA-secure sKEM by Brendel et al. [14]. We call our instantiation LARKG-Frodo and present it in Figure 6. Illustrated in Figure 3, lines 1 and 2 of the

```

KeyGen(pp)
-----
1:  $\mathbf{A} \leftarrow_{\mathcal{S}} \mathbb{Z}_q^{n \times m}$ 
2:  $\mathbf{S}, \mathbf{E} \leftarrow_{\mathcal{S}} \chi$ 
3:  $\mathbf{B} \leftarrow \mathbf{A}\mathbf{S} + \mathbf{E}$ 
4: return (sk, pk) = ( $\mathbf{S}, (\mathbf{A}, \mathbf{B})$ )

Decapsulate( $\mathbf{B}', \mathbf{S}, \mu, \mathbf{C}$ )
-----
1: return rec( $\mathbf{B}'\mathbf{S}, \mathbf{C}$ )

Encapsulate( $\mathbf{A}, \mathbf{B}$ )
-----
1:  $(\mathbf{S}', \mathbf{E}') \leftarrow_{\mathcal{S}} \chi$ 
2:  $\mathbf{B}' \leftarrow \mathbf{A}\mathbf{S}' + \mathbf{E}'$ 
3:  $\mathbf{E}'' \leftarrow_{\mathcal{S}} \chi$ 
4:  $\mathbf{V} \leftarrow \mathbf{S}'\mathbf{B} + \mathbf{E}''$ 
5:  $\mathbf{C} \leftarrow \langle \mathbf{V} \rangle_{2^B}$ 
6:  $K \leftarrow [\mathbf{V}]_{2^B}$ 
7: return (ss,  $C$ ) = ( $K, \mathbf{C}$ )

```

Figure 6: IND-CPA FrodoKEM [9] based on LWE parameters (n, m, χ) . This implies an nn-IND-CCA-secure sKEM as detailed in Figure 3.

Encapsulate algorithm that generates an LWE key pair $(\mathbf{S}', \mathbf{B}')$ are lifted into a separate algorithm KeyGenEnc. The encapsulate algorithm must now take in the secret key \mathbf{S}' to correctly compute the shared secret. FrodoKEM has some protocol specific parameters \bar{n}, \bar{m} and \bar{B} , which are defined as $\bar{B} = (\log q) - B$ and where \bar{n}, \bar{m} and B are chosen such that $\bar{n} \cdot \bar{m} \cdot B > \lambda$, the security parameter. The function $[\cdot]_{2^B} : \mathbb{Z}_q^{n \times m} \rightarrow \mathbb{Z}_{2^{B-1}}$ is a rounding function defined as $[v]_{2^B} : v \rightarrow \lfloor 2^{-\bar{B}} v \rfloor \bmod 2^B$, which partitions \mathbb{Z}_q into 2^B intervals of integers with matching B significant bits. Similarly, the function $\langle \cdot \rangle_{2^B} : \mathbb{Z}_q^{n \times m} \rightarrow \{0, 1\}$ is a cross-rounding function defined as $\langle v \rangle_{2^B} : v \rightarrow \lfloor 2^{-\bar{B}+1} v \rfloor \bmod 2$, which divides \mathbb{Z}_q into 2 partitions according to the value of the bit at position $B + 1$. We can now define the reconciliation function rec, which takes as input $\mathbf{B}'\mathbf{S}$ and recovery information \mathbf{C} . It element-wise computes a matrix \mathbf{V}' such that each element is the closest which satisfies $\langle \mathbf{V}'_{i,j} \rangle_{2^B} = \mathbf{C}_{i,j}$. It outputs the shared secret as $[\mathbf{V}'_{i,j}]_{2^B}$. Crucially, leaking $\langle \mathbf{V} \rangle_{2^B}$ has been shown to not reveal any information about the shared secret, we refer to Bos et al. [9] for further details. We have chosen our notation carefully so that parameters and variables in FrodoKEM (Figure 6) align those used to describe our generic LARKG construction in Figure 5. We instantiate KDF with HKDF [22] with double length output, i.e. of length 2λ , parsed as $\rho || \mu$, each a bit string of length λ . To sample a Gaussian distributed key K from ρ , we use Frodo's algorithm SampleG, that takes as input a uniform seed ρ , a centre α and a standard deviation η and outputs $K \sim \chi_\eta$ centred at α .

The key pairs derived by our LARKG-Frodo instance from Figure 5 are compatible with LWE and SIS-based KEM/PKE and signature schemes. We note that the construction of sKEM using FrodoKEM by Brendel et al. only meets nn-IND-CCA security, our ARKG instance is at most hSKS-secure. This property also implies LARKG-

KeyGen(pp)
1: $\mathbf{A} \leftarrow \mathcal{R}_q^{k \times k}$
2: $\mathbf{s}, \mathbf{e} \leftarrow \chi_\eta^k$
3: $\mathbf{b} \leftarrow \mathbf{A}\mathbf{s} + \mathbf{e}$
4: return $(\text{sk}, \text{pk}) = (\mathbf{s}, (\mathbf{A}, \mathbf{b}))$
Enc(pp, \mathbf{b}, m)
1: $\mathbf{r}, \mathbf{e} \leftarrow \chi_\eta^k$
2: $\mathbf{u} \leftarrow \mathbf{A}^T \mathbf{r} + \mathbf{e}$
3: $e' \leftarrow \chi_\eta$
4: $v \leftarrow \mathbf{b}^T \mathbf{r} + e' + K$
5: return $c = (\mathbf{u}, v)$
Dec($\mathbf{s}, \mathbf{u}, v$)
1: $m \leftarrow \lfloor v - \mathbf{s}^T \mathbf{u} \rfloor_{q/2}$
2: return m

Figure 7: Kyber PKE algorithms with MLWE parameters (m, k, χ) . This gives an nn-IND-CCA-secure sKEM shown in Figure 8.

Frodo is PK-unlinkable.

Remark 3. We note a small error in the sKEM construction described in Brendel et al. [14] that prevents correctness of the protocol in the plain-LWE setting. We have presented a correction in our protocol described in Figure 3. Their analogous RLWE-based sKEM is unaffected due to commutativity in the ring.

4.2. LARKG Instantiation from Kyber

In this section we show how to construct LARKG based on Kyber, which we call LARKG-Kyber. We cannot directly use the generic construction of sKEMs as given by Brendel et al. [14] since Kyber uses a PKE scheme instead of the reconciliation methods used by e.g. Frodo [9]. We first show that we can construct an sKEM based on Kyber, the security of which relies on the IND-CPA property of the underlying PKE scheme. Intuitively, our sKEM is constructed by extracting some computations (which can be viewed as creating an LWE key) in the encapsulation algorithm into a key generation algorithm, KeyGenEnc, for sKEM.

In particular, the encapsulation algorithm works by generating some randomness for key generation, using Kyber’s PKE to encrypt the shared key. In our scheme, we separate out part of the Encryption algorithm into a key generation algorithm for sKEM, namely, KeyGenEnc. We base our construction on Kyber PKE which we have recalled in Figure 7. Note that we have abstracted away the use of $\text{compress}_q, \text{decompress}_q, \text{NTT}$ and NTT^{-1} algorithms present in the specification of Kyber [10] for conciseness. Compression is mainly used to discard low-order bits (which has limited impact on correctness), which allows for smaller parameter choices. The only exception to this during the decryption algorithm where compress_q is used to ‘round’ the decryption to 0 or $\lfloor \frac{q}{2} \rfloor$. We have modelled this behaviour with a rounding function, which we denote by $\lfloor \cdot \rfloor_{q/2}$. Looking ahead, our implementation

will use compression and NTT techniques since this offers fair comparisons with Kyber implementations. The security of Kyber is based on MLWE, and so our construction is also presented using module lattices.

Our algorithms are given in Figure 8. In Theorem 4, we show that our construction is nn-IND-CCA secure based on the security of Kyber’s IND-CPA secure PKE. We can then use the generic construction in Figure 5 to construct an ARKG scheme based on Kyber.

The correctness of our sKEM is implied by the correctness of the underlying PKE. Intuitively, we have split the encapsulation algorithm into two algorithms for sKEM but not changed the underlying computations. Hence, we choose parameters to ensure PKE is correct, and correctness of our sKEM follows.

Theorem 4. *Our sKEM construction based on Kyber in Figure 8 is nn-IND-CCA secure.*

Proof. We consider the PK-Unlinkability experiment $\text{Exp}_{\text{sKEM}}^{\text{nn-IND-CCA}}(\lambda)$ and give a direct reduction to the IND-CPA property of the Kyber’s PKE scheme. To do so, we construct an adversary \mathcal{B} that sets up the sKEM experiment as described with the following exceptions. It invokes the IND-CPA game to receive the matrix \mathbf{A} , the challenge key pk^* and a challenge ciphertext (\mathbf{u}^*, v^*) based on one of two distinct messages $K_b, b \in \{0, 1\}$. Note that here, \mathbf{u}^* takes the form of the keys created by KeyGenEnc. \mathcal{B} forwards the following sKEM challenge to \mathcal{A} : $(\text{pk}^*, \mathbf{u}^*, v^*, K_1)$. Note that the adversary cannot distinguish which key queries are not genuine since \mathbf{u}^* is distributed properly in \mathcal{R}_q^k . Note that \mathcal{B} does not need to respond to any oracle queries for nn-IND-CCA.

\mathcal{A} guesses a bit b' , which \mathcal{B} forwards as a response to its own game. It wins its IND-CPA experiment whenever \mathcal{A} wins its nn-IND-CCA game. Thus, we have:

$$\Pr \left[\text{Exp}_{\text{sKEM}}^{\text{nn-IND-CCA}}(\lambda) = 1 \right] = \Pr \left[\text{Exp}_{\text{PKE}}^{\text{IND-CPA}}(\lambda) = 1 \right]$$

which implies

$$\text{Adv}_{\mathcal{A}, \text{sKEM}}^{\text{nn-IND-CCA}}(\lambda) = \text{Adv}_{\mathcal{B}, \text{PKE}}^{\text{IND-CPA}}(\lambda).$$

Hence, we have shown that sKEM nn-IND-CCA secure if Kyber PKE is IND-CPA secure. \square

In this instantiation we use our result from Theorem 4. The distribution of χ is chosen to be a centered binomial distribution (CBD) instead of small discretised Gaussian, so that our LARKG protocol outputs derived keys that are also distributed according to a CBD. We emphasise this is not an incompatibility with Kyber, rather the decision to use CBD is due to the relative inefficiencies of sampling Gaussian distributions. LWE-based encryption does not depend on the exact distribution but rather the standard deviation [10], thus we do not lose security of guarantees of LWE when error and secrets are sampled from a CBD. We instantiate KDF with HKDF [22] with double length output, i.e. of length 2λ , parsed as $\rho || \mu$, each a bit string of length λ . To sample a centered binomial distributed key K from ρ , we use Kyber’s algorithm CBD, that takes as input a uniform seed ρ , a centre α and a standard deviation η and outputs $K \sim \chi_\eta$ centred at α .

Setup(pp)	KeyGenDec(pp)	KeyGenEnc(pp)										
1: $\mathbf{A} \leftarrow \mathcal{R}_q^{k \times k}$	1: $\mathbf{s}, \mathbf{e} \leftarrow \mathcal{X}_\eta^k$	1: $\mathbf{r}, \mathbf{e} \leftarrow \mathcal{X}_\eta^k$										
2: return pp = \mathbf{A}	2: $\mathbf{b} \leftarrow \mathbf{A}\mathbf{s} + \mathbf{e}$	2: $\mathbf{u} \leftarrow \mathbf{A}^T \mathbf{r} + \mathbf{e}$										
	3: return (sk, pk) = (s, (A, b))	3: return (sk, pk) = (r, u)										
<table border="0" style="width: 100%;"> <tr> <td style="border-bottom: 1px solid black; padding: 5px;">Encapsulate(pp, b)</td> <td style="border-bottom: 1px solid black; padding: 5px;">Decapsulate(u, s, C)</td> </tr> <tr> <td style="padding: 5px;">1: $K \leftarrow \mathcal{R}_q$</td> <td style="padding: 5px;">1: Parse \mathbf{C} as v</td> </tr> <tr> <td style="padding: 5px;">2: $e' \leftarrow \mathcal{X}_\eta$</td> <td style="padding: 5px;">2: $K \leftarrow v - \mathbf{s}^T \mathbf{u}$</td> </tr> <tr> <td style="padding: 5px;">3: $v \leftarrow \mathbf{b}^T \mathbf{r} + e' + K$</td> <td style="padding: 5px;">3: return K</td> </tr> <tr> <td style="padding: 5px;">4: return (ss, C) = (K, v)</td> <td></td> </tr> </table>			Encapsulate(pp, b)	Decapsulate(u, s, C)	1: $K \leftarrow \mathcal{R}_q$	1: Parse \mathbf{C} as v	2: $e' \leftarrow \mathcal{X}_\eta$	2: $K \leftarrow v - \mathbf{s}^T \mathbf{u}$	3: $v \leftarrow \mathbf{b}^T \mathbf{r} + e' + K$	3: return K	4: return (ss, C) = (K, v)	
Encapsulate(pp, b)	Decapsulate(u, s, C)											
1: $K \leftarrow \mathcal{R}_q$	1: Parse \mathbf{C} as v											
2: $e' \leftarrow \mathcal{X}_\eta$	2: $K \leftarrow v - \mathbf{s}^T \mathbf{u}$											
3: $v \leftarrow \mathbf{b}^T \mathbf{r} + e' + K$	3: return K											
4: return (ss, C) = (K, v)												

Figure 8: sKEM based on Kyber with MLWE parameters (m, k, χ) .

Remark 4. We observe that our approach could be applied to other lattice-based KEMs that use the PKE method of computing a shared secret. In particular, the team behind Frodo made changes in their submission to Round 3 of the NIST PQC standardisation process. They shifted away from reconciliation techniques in favour of a PKE based solution⁴, which we distinguish from Frodo by calling this variant FrodoKEM. Thus both Frodo (proof by Brendel et al. [14]) and FrodoKEM (analogous to Theorem 4) are sKEMs and therefore can be used to construct LARKG.

4.3. LARKG usage in KEMs and PKE

KEMs are an important primitive used in many cryptographic protocols, including for challenge-response for authentication. Instead of requiring a user to sign a challenge string, the user must decrypt a challenge KEM issued by the relying party. The user sends back the decrypted plaintext, and the sever accepts if the plaintexts match. This is the approach proposed in KEMTLS [26], which is a post-quantum replacement for TLS1.3. The benefit over more traditional signature-based solutions is that lattice KEMs typically offer more efficient constructions. For this reason, we conjecture that KEM based challenge-response protocols (such as WebAuthn) are likely to become more wide spread as post-quantum protocols are adopted.

Our lattice ARKG has been designed to produce key pairs that are compatible with lattice-based KEMs and signature schemes. Furthermore, security of the composed protocol is maintained when used with keys generated by LARKG. This follows from the general composability theorem in [1] which states that ARKG can be securely composed with arbitrary public-key protocols provided the key distributions match, and ARKG is PK-unlinkable.

We first consider LARKG-Frodo, presented in Section 4.1, with Frodo [9] as a KEM. This composition of protocols follows from the fact that our LARKG-Frodo outputs derived keys that have a small Gaussian distribution (if χ is also a small Gaussian), which is exactly the structure of secret keys of Frodo. All that remains is to select parameters to meet cryptographic requirements. It is sufficient to instantiate the protocol with the parameters from Frodo directly. This follows because the long term and ephemeral public key forms a cryptographically hard LWE instances, and the derived key has the same dimensions

with an error distribution that has standard deviation (s.d.) $\bar{\sigma} = 2\sigma$, where σ is the standard deviation from the Frodo parameters. Therefore the derived keys are secure since $\bar{\sigma} > 2\sqrt{n}$ if $\sigma > 2\sqrt{n}$, which ensures cryptographic hardness [12]. Note that this does not impact correctness since the KEM correctly decrypts the ciphertext since $\|\mathbf{E}''\| \ll 2^{\bar{B}-2}$.

We next consider the LARKG-Kyber instance in Section 4.2. To select parameters to meet cryptographic requirements, we follow the parameter levels included in the round 3 Kyber specification⁵. Intuitively, because the security of our construction for an sKEM based on Kyber relies the underlying PKE, parameters that ensure a secure PKE instance also imply a secure sKEM. Thus we are assured the keys used in sKEM form cryptographically hard LWE instances if we instantiate with secure Kyber parameters. The resulting LARKG instance outputs keys with a wider error distribution. However this does not affect correctness provided $\|\mathbf{e}^T \mathbf{r} + e_2 + \mathbf{c}_v + \mathbf{c}_t^T \mathbf{r} - \mathbf{s}^T \mathbf{e}_1 - \mathbf{s}^T \mathbf{c}_u\|_\infty \leq \lceil q/4 \rceil$, where \mathbf{e} and \mathbf{r} are CBD error vectors, e_2 is a small scalar, $\mathbf{c}_t, \mathbf{C}_u$ are (small Gaussian) rounding error vectors, and \mathbf{s} is the derived secret key. The only term that changes from the original correctness analysis is that \mathbf{s} is now distributed according to a centered binomial distribution with variance η , rather than with variance $\eta/2$ that \mathbf{s} would typically have. The exact amount this changes will depend on the selected parameter choices. To give an example, if we consider Kyber Level 5 parameter sets then the probability of failure is 2^{-174} . When composed with LARKG-Kyber also instantiated with Kyber’s level 5 parameter set, then we estimate using a security estimator for Kyber [27]. This gives a value of $M = 9$ (at the expense of additional derived public keys) we are able to obtain correctness values of $1 - 2^{-174}$, or equivalently $\tau = 274$, when the distribution χ_b is instantiated with standard deviation $\sigma_b = 5\sigma_u$. Further optimisation of parameter choices may improve correctness or reduce M yet retain the desired security level may be possible, but we leave identifying such parameters future work as it will be application and security level dependent.

Furthermore, we note that since many lattice KEMs can be extended to PKE algorithms by replacing the sampling of a shared key with the message to be encrypted, we conclude that our LARKG construction is also compatible

4. <https://frodokem.org/files/FrodoKEM-specification-20210604.pdf>

5. <https://www.pq-crystals.org/kyber/data/kyber-specification-round3-20210804.pdf>

Table 1: Parameter trade-offs for our Lattice-ARKG based on Dilithium level 5 parameters.

β'	$\Pr[\ \mathbf{S}\ _\infty \leq \beta']$	τ
5	0	0
6	6×10^{-5}	0.0009
7	0.1	0.15
8	0.63	1.43
9	0.92	3.64
10	1	∞

with corresponding lattice PKE schemes. Concrete examples of this include FrodoPKE and KyberPKE. Finally, our LARKG construction is compatible with many lattice KEMs that achieve chosen ciphertext security via (variants of) the FO-transform [24], [25], again such examples include Frodo, FrodoKEM or Kyber.

4.4. LARKG usage in Signatures

Some lattice digital signatures, such as qTESLA [7] also support small Gaussian keys, and are compatible with the previous techniques described for KEMs. However, other schemes based on SIS use short vectors as secret keys (provided they meet some min-entropy [28]). We consider Dilithium [29], which was selected for standardisation in the NIST post quantum process and takes this form. The key structure consists of two short vectors \mathbf{s}_1 and \mathbf{s}_2 , bounded by a value β that satisfy $\mathbf{A}\mathbf{s}_1 + \mathbf{s}_2 = \mathbf{b}$. Since our protocol outputs keys with Gaussian distribution $\chi_{2\sigma}$, we need to ensure that the parameters selected for Dilithium remain secure, despite the wider error distribution. To achieve this, one could select a tighter distribution for the long term keys. This corresponds to a potentially higher-security level, and therefore larger key sizes, to ensure the derived keys meet their desired minimum security level.

Alternatively, another approach would be to choose $\beta' < 2\beta$ for the composed signature, where there is some probability that a derived secret key has norm $> \beta'$ and therefore would be an invalid key. However, as this process is interactive, it is not possible to simply repeat the procedure as required until a ‘good’ key is found. Instead, it must be somehow accounted for during execution of DerivePK and before the values of $(\mathbf{S}_1, \mathbf{S}_2)$ are known. To achieve this one could provide t candidate keys, that probabilistically, at least one of t will create a valid derived key pair. We illustrate this trade off in Table 1, based on parameters for Dilithium level 5 (where $\beta = 3$). We vary the value of β' to show how this affects correctness of our protocol and corresponding τ value. Intuitively, increased failure rate allows for higher security since the norm bound is smaller. We emphasise that the value of β' must be selected based on security analysis of the composed signature.

4.5. LARKG Implementation, Parameters and Performance

In our code repository [30], we provide implementations of our LARKG-Frodo and LARKG-Kyber schemes, and we present the benchmarking of these in Table 2, as well as respective sKEM instantiations. We taken as the mean of 100 invocations on a Intel i7-8700 processor,

Table 2: Performance of Frodo- and Kyber-based KEM and LARKG primitives for a single call, in seconds, calculated as the mean of 10 invocations.

Primitive	Algorithm	Instantiation	
		Frodo	Kyber
KEM	KeyGen	3.412	0.011
	Encaps	3.624	0.016
	Decaps	3.546	0.024
LARKG	KeyGen	1.146	0.010
	DerivePK	2.536	0.027
	DeriveSK	1.641	0.014

with a clock speed of 3.20GHz. Note that the libraries we used are not optimised or side-channel secure, instead we illustrate here the relative performance of LARKG versus the underlying KEMs.

For LARKG-Frodo, we instantiate the scheme with parameter set FrodoKEM-976-AES, i.e., $n = 976$, $m = k = 8$, $q = 2^{16}$, with the error distribution χ having $\sigma = 2.3$. This matches parameter choices for NIST Level 3 of FrodoKEM. We used a Python implementation of FrodoKEM⁶. Compared to FrodoKEM’s KeyGen, the figures in Table 2 give an overall performance cost of 56% for LARKG-Frodo when generating a derived key pair, i.e., calling KeyGen, then DerivePK, and finally DeriveSK. Note that our LARKG-Frodo.KeyGen is 298% faster than FrodoKEM’s KeyGen as we fix the matrix \mathbf{A} as part of our public parameters, whereas FrodoKEM must recompute this matrix from a seed on each invocation.

For LARKG-Kyber, we instantiated Kyber with the Level 5 parameter set, Kyber1024, using a Python library⁷. Precisely, the parameters we use are: $n = 256$, $k = 4$, $q = 3329$, $n_1 = n_2 = 2$, $d_u = 11$, $d_v = 5$. Note that we also used the number theoretic transform (NTT) to perform polynomial multiplication to give a fair comparison to the base Kyber implementation. We also demonstrate the correctness of our implementation in our code repository [30] by running an instance of Kyber’s IND-CPA PKE using LARKG-Kyber keys, as well as give a Kyber with split KEM functionality based on the required changes presented in Figure 8.

We present timings given in Table 2 that do not account for repeat executions required by rejection sampling, and an example whereby we do consider the probability of failure. We do so because the rejection sampling parameters are highly dependent on scheme and security level, whereas a single run demonstrates the relative performance of the underlying computations with the understanding that it would most likely require repeated executions. By setting the distribution χ_b to have standard deviation $\sigma_b = 5\sigma_a$, then we get $M = 9$ (which is similar to other works [17]), then we instead measure a full correctness value of 2^{-174} which corresponds to $\tau = 174$, still targeting 256-bit security. In particular, the wider distribution of LARKG-Kyber keys has increased the difficulty in solving MLWE, where the optimal block size of dual and primal attacks has increased from 868 to 953, which gives an estimated classical bit security of 278 versus the 253 of Kyber1024 [27].

6. github.com/microsoft/PQCrypto-LWEKE/tree/master/python3

7. github.com/jack4818/kyber-py

As an *illustrative example*, we measure the average time for a successful run of the complete LARKG primitive (i.e., KeyGen, followed by DerivePK, then DeriveSK), when using rejection sampling and instantiating χ_a and χ_b to be the same CBD distributions with $\eta = 2$. The total run time of LARKG-Kyber is then 0.19s with an acceptance rate of 0.32—measured over 1000 runs on an i7-8700. This is an increase in execution time of 272.5%, compared to the average without rejection sampling—which, for Kyber, is 0.051s. This means that, for these particular parameters, a successful run of DeriveSK requires around three executions of LARKG. These measurements vary considerably based on parameter selection.

Looking at Table 2, we note that, similarly to LARKG-Frodo, LARKG-Kyber’s DeriveSK takes 92% less time than DerivePK. Without any optimisations, we see an overall cost of 51ms for LARKG-Kyber when generating derived key pairs compared to Kyber’s 11ms for standalone key generation. We note that currently existing Python libraries behind Frodo and Kyber which we implement upon, differ in that the Kyber library is optimised whereas Frodo library not. This explains why in our measurements in Table 2 Frodo appear much slower (with values being in s range) in comparison to Kyber (in ms range) than one would expect. We anticipate that an optimised underlying Python library for Frodo might yield a somewhat slower implementation (based on similar security levels) than Kyber albeit not at the same magnitude. We emphasise this does not affect our claim of practicality since we compare relative performance.

5. Conclusion

In this paper we have given the first Asynchronous Remote Key Generation scheme for use with lattice-based public key cryptosystems that use either SIS or LWE-based keys. We have given a general lattice ARKG (LARKG) using Split KEMs which we prove is PK-unlinkable under the nn-IND-CCA property of the sKEM. We showed that if LWE is hard, then our construction has hsKS security, which is extended to msKS security under the sn-IND-CCA property of the underlying sKEM. We gave two instantiations of our hsKS-secure protocol, LARKG-Frodo and LARKG-Kyber, which are based on popular KEMs. To use Kyber in LARKG-Kyber, we first had to show that it could be used to realise an sKEM and showed that it satisfies the nn-IND-CCA security property. Other KEMs based on an IND-CPA PKE could also be used to create sKEMs in a similar way.

To capture the need to restrict the norm of the derived secret key, we have defined a relaxed notion of ARKG correctness to account for inherent noise of our underlying lattice building blocks. Our constructions support lattice schemes that use both LWE and SIS style keys, and thus are compatible with several KEMs, PKE and signatures selected for standardisation by NIST, or that appear in round 3 as alternative candidate. We discussed integration details with selected primitives Frodo, Kyber and Dilithium. We highlighted the necessary the parameter selections needed to ensure that ARKG keys are compatible with the selected primitive, without compromising security or correctness.

Additionally, we have presented publicly-accessible implementations in Python of our LARKG based on both FrodoKEM and Kyber, along with timings and a discussion of our results. We note that, despite our unoptimised implementations, our performance metrics show that our protocols are practical.

Future Work. A future direction for this work could be to develop sKEMs that are secure against adaptive adversaries, in particular, to satisfy the sn-IND-CCA security. This would allow LARKG schemes with stronger key secrecy properties, i.e. msKS and mwKS, that would allow for use in wider applications with stronger adversarial models. The difficulty in existing techniques to construct sKEMs lies in the use of the FO transform to achieve IND-CCA security in the corresponding KEM, e.g. Kyber or Frodo. In particular, the decapsulation requires reconstruction of the ciphertext using the same randomness to ensure consistency. However, when translated to the sKEM setting using either the techniques by Brendel et al. [14] or Theorem 4, some of this ‘randomness’ forms the secret key of the encapsulator and thus is not known during decapsulation. This prevents use of the FO-like transform to achieve adaptive security in sKEMs, i.e. sn-IND-CCA security. Overcoming these challenges would be interesting future work.

We also observe that the original discrete-log construction of ARKG was previously used to enable anonymous payments in cryptocurrencies. Coined stealth addresses [31], [32], a payer can generate an ephemeral address for the receiver based on its long term key. This corresponds to the DerivePK, and DeriveSK corresponds to the receiver computing the secret key for the ephemeral address. The only known constructions are based on the group-theoretic cryptography, and we speculate that lattice ARKG could be used to give a post-quantum construction for stealth addresses.

References

- [1] N. Frymann, D. Gardham, F. Kiefer, E. Lundberg, M. Manulis and D. Nilsson, ‘Asynchronous remote key generation: An analysis of yubico’s proposal for w3c webauthn’, *Proceedings of the 2020 ACM SIG-SAC Conference on Computer and Communications Security*, 2020.
- [2] D. Balfanz, A. Czeskis, J. Hodges *et al.*, ‘Web authentication: An API for accessing public key credentials level 1’, Tech. Rep., 2019. [Online]. Available: <https://www.w3.org/TR/webauthn>.
- [3] N. Frymann, D. Gardham and M. Manulis, ‘Unlinkable delegation of webauthn credentials’, in *Computer Security – ESORICS 2022*, V. Atluri, R. Di Pietro, C. D. Jensen and W. Meng, Eds., Cham: Springer Nature Switzerland, 2022, pp. 125–144, ISBN: 978-3-031-17143-7.
- [4] N. Frymann, D. Gardham, M. Manulis and H. Nartz, *Generalised asynchronous remote key generation for pairing-based cryptosystems*, Cryptology ePrint Archive, Paper 2023/456, <https://eprint.iacr.org/2023/456>, 2023. [Online]. Available: <https://eprint.iacr.org/2023/456>.
- [5] P. W. Shor, ‘Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer’, *SIAM J. Comput.*, vol. 26, no. 5,

- pp. 1484–1509, 1997, ISSN: 0097-5397. DOI: 10.1137/S0097539795293172. [Online]. Available: <https://doi.org/10.1137/S0097539795293172>.
- [6] L. Ducas, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler and D. Stehle, *Crystals – dilithium: Digital signatures from module lattices*, Cryptology ePrint Archive, Report 2017/633, <https://eprint.iacr.org/2017/633>, 2017.
- [7] E. Alkim, P. S. L. M. Barreto, N. Bindel, J. Krämer, P. Longa and J. E. Ricardini, ‘The lattice-based digital signature scheme qtesla’, in *Applied Cryptography and Network Security: 18th International Conference, ACNS 2020, Rome, Italy, October 19–22, 2020, Proceedings, Part I*, Rome, Italy: Springer-Verlag, 2020, pp. 441–460, ISBN: 978-3-030-57807-7.
- [8] P.-A. Fouque, J. Hoffstein, P. Kirchner *et al.*, *Falcon: Fast-fourier lattice-based compact signatures over ntru*. 2017.
- [9] J. Bos, C. Costello, L. Ducas *et al.*, ‘Frodo: Take off the ring! practical, quantum-secure key exchange from lwe’, in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’16, Vienna, Austria: Association for Computing Machinery, 2016, pp. 1006–1018, ISBN: 9781450341394. DOI: 10.1145/2976749.2978425. [Online]. Available: <https://doi.org/10.1145/2976749.2978425>.
- [10] J. W. Bos, L. Ducas, E. Kiltz *et al.*, ‘CRYSTALS -kyber: A cca-secure module-lattice-based KEM’, in *2018 IEEE European Symposium on Security and Privacy, EuroS&P 2018, London, United Kingdom, April 24-26, 2018*, IEEE, 2018, pp. 353–367. DOI: 10.1109/EuroSP.2018.00032. [Online]. Available: <https://doi.org/10.1109/EuroSP.2018.00032>.
- [11] E. Alkim, L. Ducas, T. Pöppelmann and P. Schwabe, ‘Post-quantum key exchange: A new hope’, in *Proceedings of the 25th USENIX Conference on Security Symposium*, ser. SEC’16, Austin, TX, USA: USENIX Association, 2016, pp. 327–343, ISBN: 9781931971324.
- [12] O. Regev, ‘On lattices, learning with errors, random linear codes, and cryptography’, in *STOC*, 2005.
- [13] M. Ajtai, ‘Generating hard instances of lattice problems (extended abstract)’, in *STOC ’96*, 1996.
- [14] J. Brendel, M. Fischlin, F. Günther, C. Janson and D. Stebila, ‘Towards post-quantum security for signal’s x3dh handshake’, in *Proceedings of the Selected Areas in Cryptography, 27th International Conference.*, ser. SAC 2020, Springer, 2020.
- [15] J. Brendel, M. Fischlin, F. Günther and C. Janson, ‘Prf-odh: Relations, instantiations, and impossibility results’, in *Advances in Cryptology – CRYPTO 2017*, J. Katz and H. Shacham, Eds., Cham: Springer International Publishing, 2017, pp. 651–681, ISBN: 978-3-319-63697-9.
- [16] Z. Brakerski, C. Gentry and V. Vaikuntanathan, ‘(leveled) fully homomorphic encryption without bootstrapping’, in *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ser. ITCS ’12, Cambridge, Massachusetts: Association for Computing Machinery, 2012, pp. 309–325, ISBN: 9781450311151. DOI: 10.1145/2090236.2090262. [Online]. Available: <https://doi.org/10.1145/2090236.2090262>.
- [17] V. Lyubashevsky, ‘Lattice signatures without trapdoors’, in *Advances in Cryptology – EUROCRYPT 2012*, D. Pointcheval and T. Johansson, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 738–755, ISBN: 978-3-642-29011-4.
- [18] J. Brendel, R. Fiedler, F. Günther, C. Janson and D. Stebila, ‘Post-quantum asynchronous deniable key exchange and the signal handshake’, in *Public-Key Cryptography – PKC 2022*, G. Hanaoka, J. Shikata and Y. Watanabe, Eds., Cham: Springer International Publishing, 2022, pp. 3–34, ISBN: 978-3-030-97131-1.
- [19] J. Ding, X. Xie and X. Lin, ‘A simple provably secure key exchange scheme based on the learning with errors problem’, Cryptology ePrint Archive, Paper 2012/688, <https://eprint.iacr.org/2012/688>, 2012. [Online]. Available: <https://eprint.iacr.org/2012/688>.
- [20] R. Lindner and C. Peikert, ‘Better key sizes (and attacks) for lwe-based encryption’, in *Topics in Cryptology – CT-RSA 2011*, A. Kiayias, Ed., Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 319–339, ISBN: 978-3-642-19074-2.
- [21] C. Peikert, ‘Lattice cryptography for the internet’, in *Post-Quantum Cryptography*, Cham: Springer International Publishing, 2014, pp. 197–219, ISBN: 978-3-319-11659-4.
- [22] H. Krawczyk, ‘Cryptographic extraction and key derivation: The hkdf scheme’, in *Advances in Cryptology – CRYPTO 2010*, T. Rabin, Ed., Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 631–648, ISBN: 978-3-642-14623-7.
- [23] C. Gentry, ‘A fully homomorphic encryption scheme’, crypto.stanford.edu/craig, Ph.D. dissertation, Stanford University, 2009.
- [24] E. Fujisaki and T. Okamoto, ‘Secure integration of asymmetric and symmetric encryption schemes’, in *Proceedings of the 19th Annual International Cryptology Conference on Advances in Cryptology*, ser. CRYPTO ’99, Berlin, Heidelberg: Springer-Verlag, 1999, pp. 537–554, ISBN: 3540663479.
- [25] E. Fujisaki and T. Okamoto, ‘Secure integration of asymmetric and symmetric encryption schemes’, *J. Cryptol.*, vol. 26, no. 1, pp. 80–101, 2013. DOI: 10.1007/s00145-011-9114-1. [Online]. Available: <https://doi.org/10.1007/s00145-011-9114-1>.
- [26] P. Schwabe, D. Stebila and T. Wiggers, ‘Post-quantum tls without handshake signatures’, in *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. New York, NY, USA: Association for Computing Machinery, 2020, pp. 1461–1480, ISBN: 9781450370899. [Online]. Available: <https://doi.org/10.1145/3372297.3423350>.
- [27] *Github - pq-crystals/security-estimates: Security estimation scripts for kyber and dilithium*. <https://github.com/pq-crystals/security-estimates>, Accessed: 2022-08-01.
- [28] S. Agrawal, C. Gentry, S. Halevi and A. Sahai, ‘Discrete gaussian leftover hash lemma over infinite domains’, in *Advances in Cryptology - ASIACRYPT 2013*, K. Sako and P. Sarkar, Eds., Berlin, Heidel-

- berg: Springer Berlin Heidelberg, 2013, pp. 97–116, ISBN: 978-3-642-42033-7.
- [29] L. Ducas, E. Kiltz, T. Lepoint *et al.*, ‘Crystals-dilithium: A lattice-based digital signature scheme’, *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2018, no. 1, pp. 238–268, 2018. DOI: 10.13154/tches.v2018.i1.238-268. [Online]. Available: <https://tches.iacr.org/index.php/TCHES/article/view/839>.
 - [30] *Source code for FrodoKEM- and Kyber-based LARKG implementations*, <https://gitlab.surrey.ac.uk/sccs/larkg>.
 - [31] P. Todd, *Stealth addresses*, <https://lists.linuxfoundation.org/pipermail/bitcoin-dev/2014-January/004020.html>, 2014.
 - [32] N. van Saberhagen, *Cryptonote v2.0*, <https://cryptonote.org/whitepaper.pdf>, 2013.