

Advanced KEM Concepts

(Hybrid) Obfuscation and Verifiable Decapsulation

IBM **Research** Security

Felix Günther

IBM Research Europe – Zurich

based on work with

Lewis Glabush
EPFL

Kathrin Hövelmanns
TU Eindhoven

Michael Rosenberg
Cloudflare


Douglas Stebila
U Waterloo

Shannon Veitch
ETH Zurich

Protocol Obfuscation


Internet protocols hide **metadata** to protect user privacy, dissuade protocol fingerprinting, and prevent network ossification

- TLS 1.3 Encrypted Client Hello, QUIC, pseudorandom cTLS, obfs4, Shadowsocks, ...
- “Fully-encrypted” protocols, with **obfuscated** key exchange



Quantum-safe transition?

☒ ML-KEM public keys and ciphertexts **don't look random!**



ML-KEM/Kyber public keys

- vector of coefficients mod $q = 3329$

$$\begin{matrix} [a_1] & [a_2] & [a_3] & \dots & [a_b] \\ \uparrow & \uparrow & \uparrow & & \uparrow \end{matrix} \quad a_i \in \mathbb{Z}_q = \{0, \dots, 3328\} \text{ -- each } a_i \text{ represented in 12 bits}$$

most sig. bit of each value biased towards 0

- Encoding for public keys:

1. accumulate into one big number
2. rejection sampling: reject if msb is 1

Encoded public keys **~2.5% smaller** than regular
(-19/28/38 bytes for ML-KEM-512/768/1024)

ML-KEM-768 **likelihood of rejection is ~17%**


$$\begin{matrix} [A = a_1 + a_2 \cdot q + a_3 \cdot q^2 + \dots + a_b \cdot q^{b-1}] \\ \uparrow \end{matrix}$$

most sig. bit still biased towards 0

ML-KEM/Kyber ciphertexts

- vector of **compressed** coefficients – need to first “decompress”
- encoded ciphertexts larger than regular (6–15%)


Obfuscated KEMs



ML-KEM

+ Kemeleon public key and ciphertext encoding

✓ = Obfuscated KEM: **ML-Kemeleon**



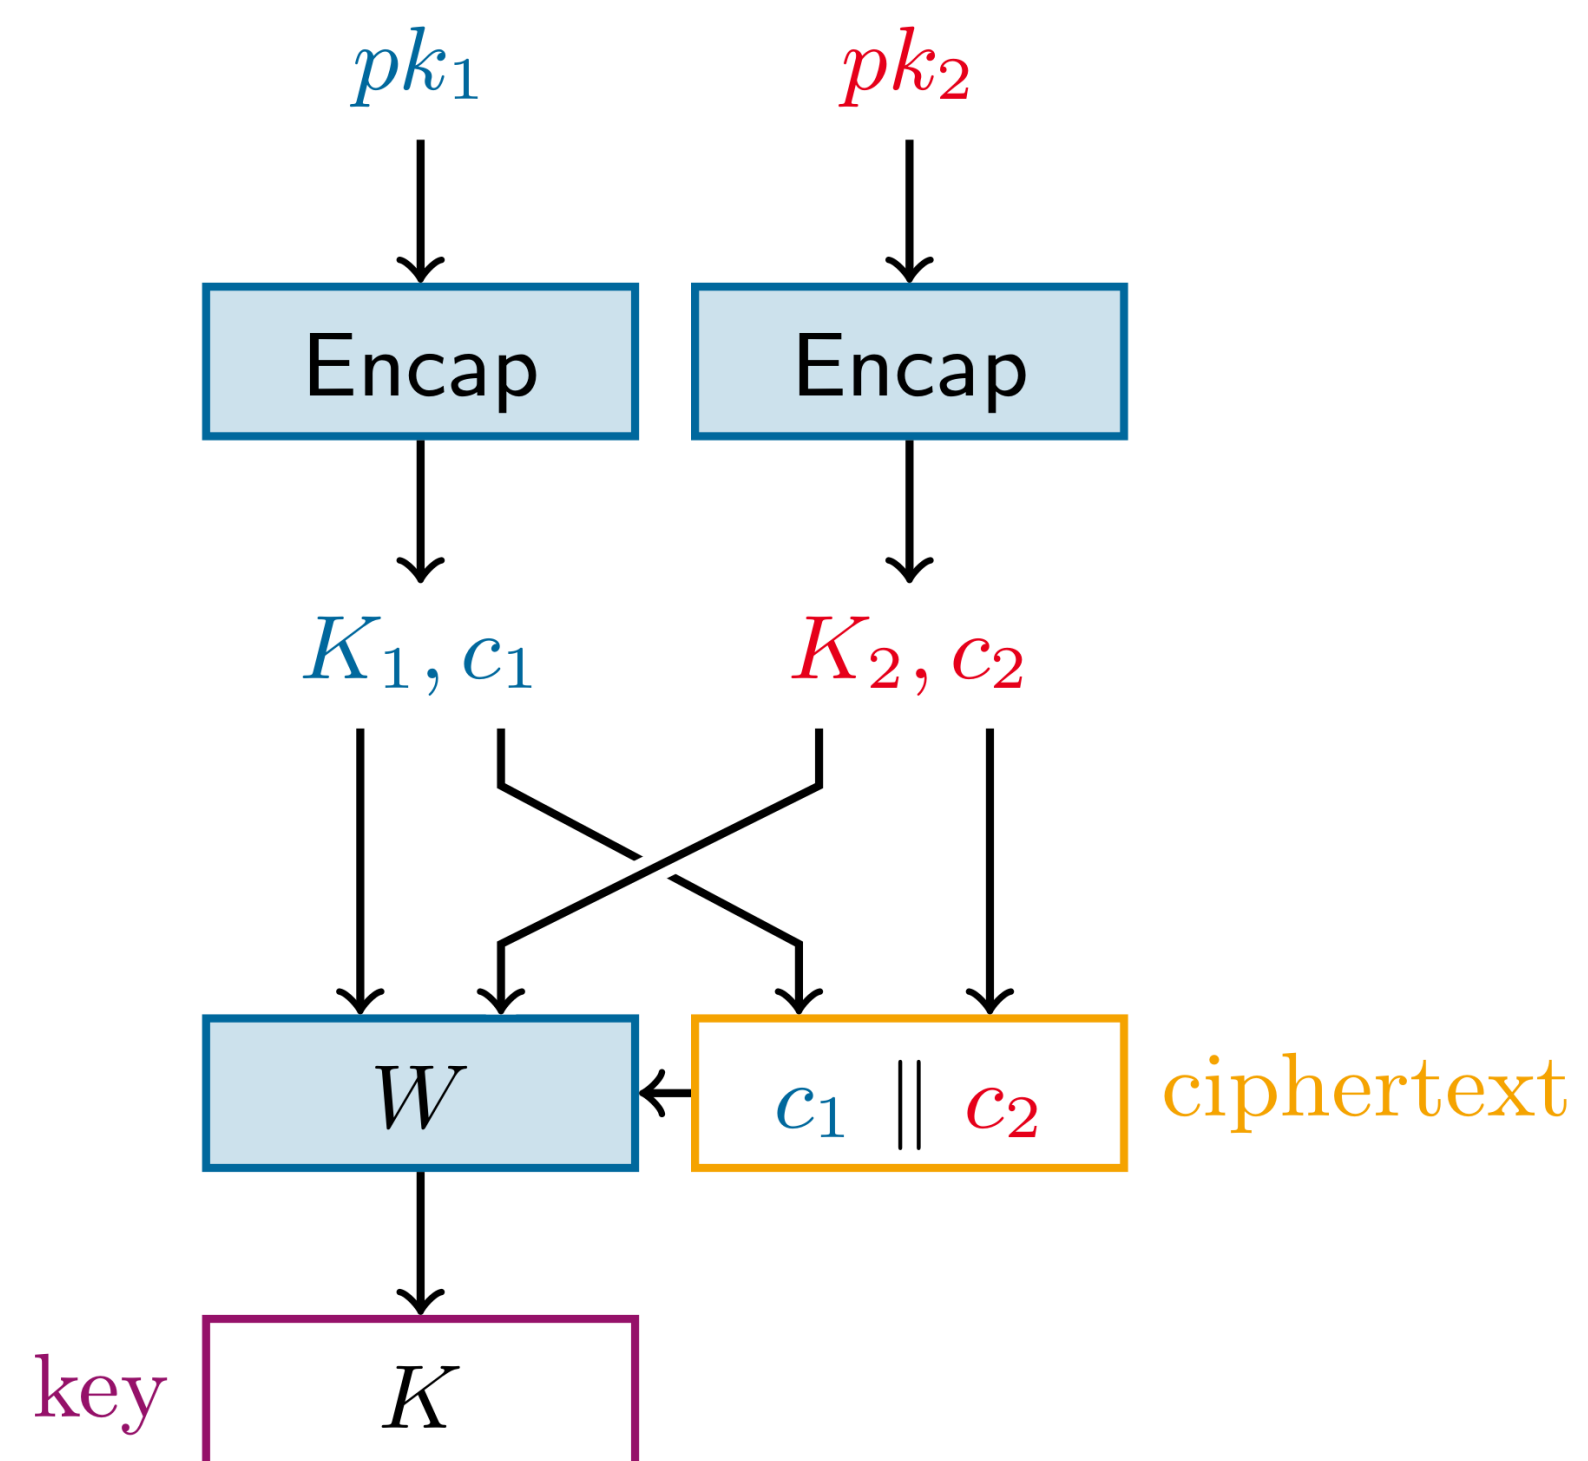
Kemeleon considered for adoption by CFRG

<https://datatracker.ietf.org/doc/draft-veitch-kemeleon/>

(more variants: no rejection, deterministic, ...)

Hybrid KEMs


Parallel Combiner



TLS 1.3 hybrid, HPKE Xyber, XWing, ...

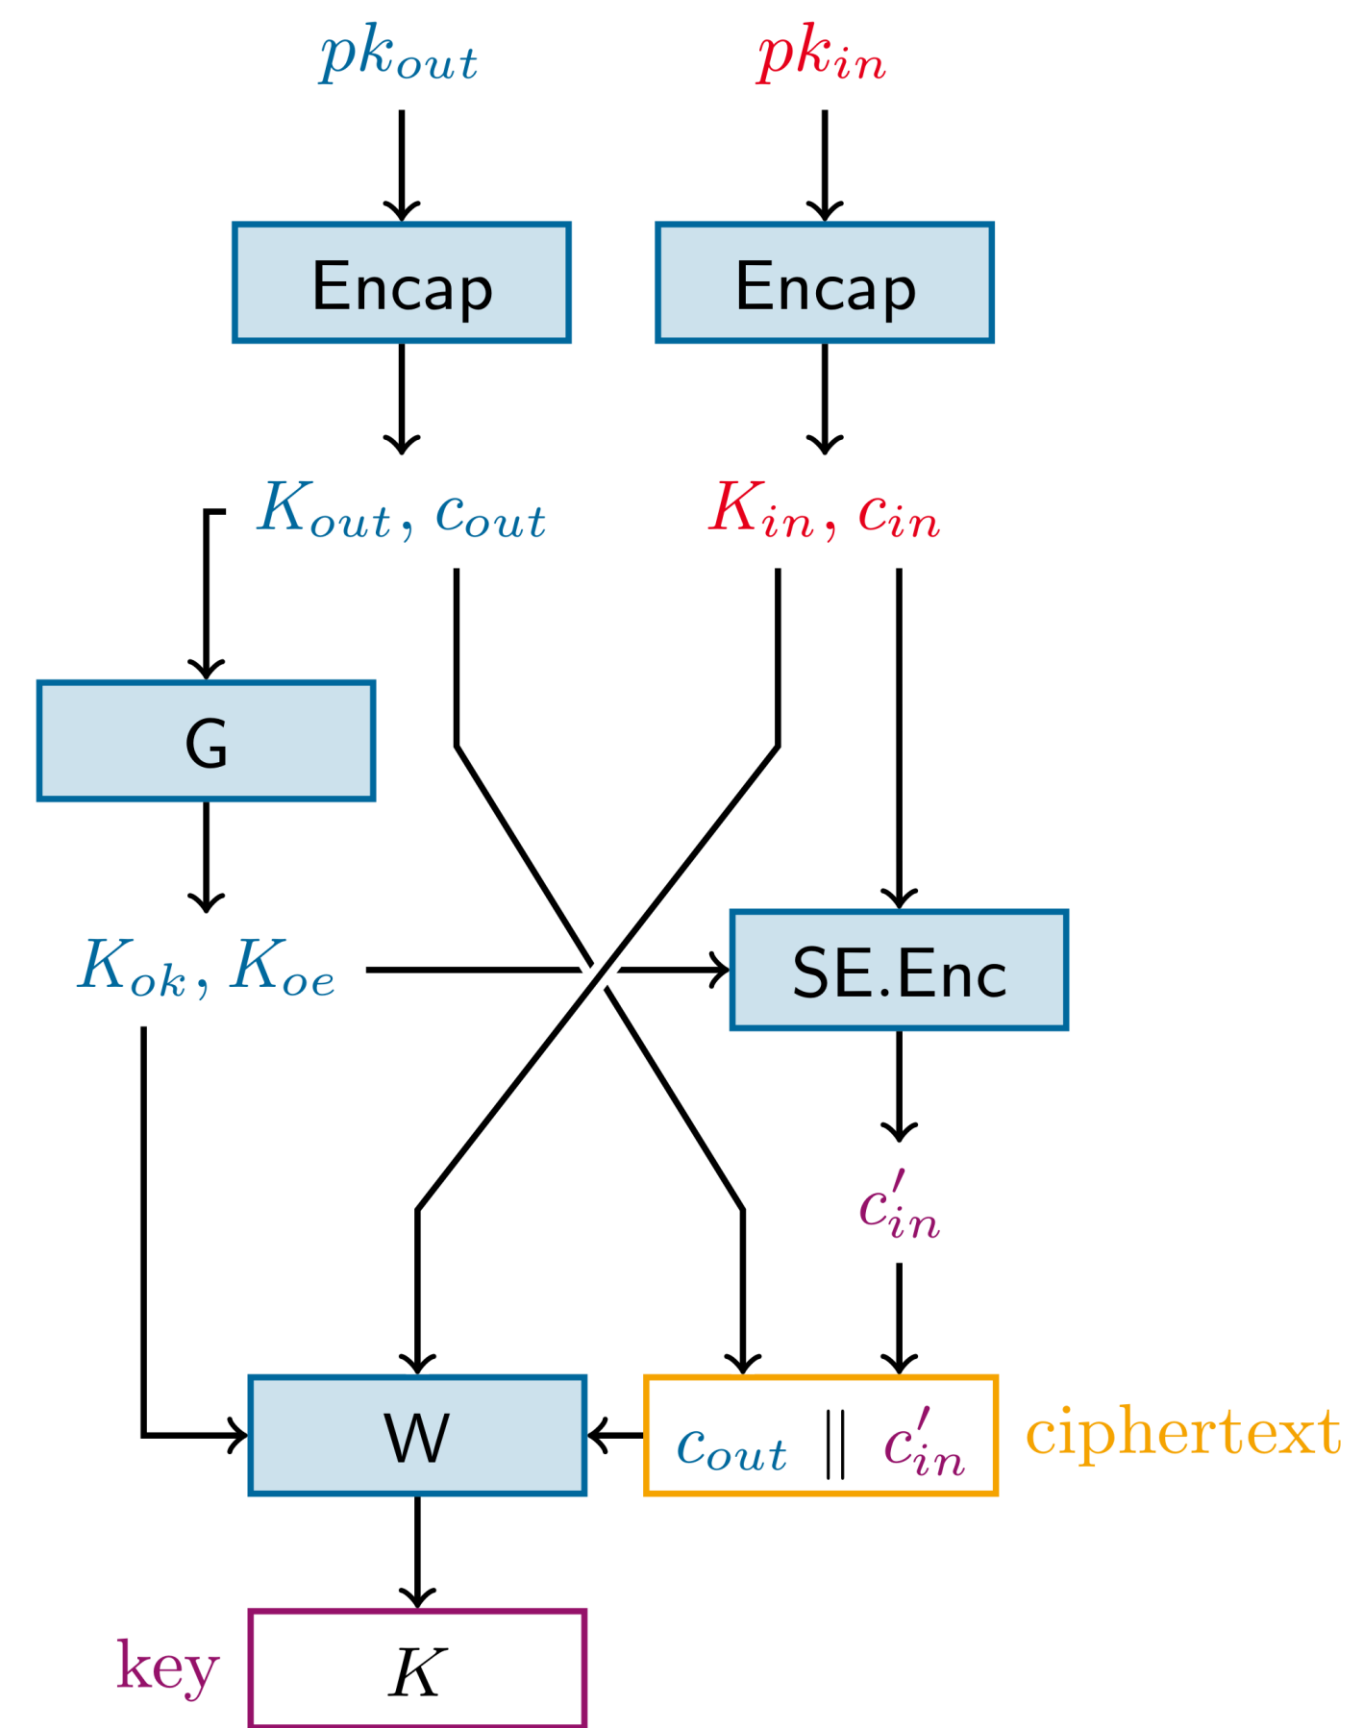
- ✓ Hybrid IND-CCA security
- ✗ Hybrid Obfuscation

Hybrid Obfuscated KEMs



More details?
→ Shannon & Michael's talk @ RWC
(Thu morning, PQ Deployment session)

OEINC



Outer-encrypts-inner nested combiner

- ✓ Hybrid IND-CCA security
- ✓ Hybrid Obfuscation
- ✓ Low overhead: 1 PRG + 1 XOR

example: outer = DH-Elligator (statistical)
 inner = ML-Kameleon (computational)

Use OEINC to build

- hybrid obfuscated key exchange
- hybrid PAKE (w/ adaptive corruptions)

Cryptography Is Brittle

functionality




security

```
static OSStatus
SSLVerifySignedServerKeyExchange(SSLContext *ct
    SSLBuffer signedParams, uint8_t *signature,
{
    OSStatus      err;
    ...

    if ((err = SSLHashSHA1.update(&hashCtx, &se
        goto fail;
    if ((err = SSLHashSHA1.update(&hashCtx, &si
        goto fail;
    if ((err = SSLHashSHA1.final(&hashCtx, &has
        goto fail;
    ...

fail:
    SSLFreeBuffer(&signedHashes);
    SSLFreeBuffer(&hashCtx);
    return err;
}
```



Cryptography Is Brittle

functionality




security

```
static OSStatus
SSLVerifySignedServerKeyExchange(SSLContext *ctx,
SSLBuffer signedParams, uint8_t *signature,
{
    OSStatus      err;
    ...

    if ((err = SSLHashSHA1.update(&hashCtx, &signedParams)) != 0)
        goto fail;
    if ((err = SSLHashSHA1.update(&hashCtx, &signature)) != 0)
        goto fail;
    if ((err = SSLHashSHA1.final(&hashCtx, &hash)) != 0)
        goto fail;
    ...

fail:
    SSLFreeBuffer(&signedHashes);
    SSLFreeBuffer(&hashCtx);
    return err;
}
```



```
Algorithm 18 ML-KEM.Decaps_internal(dk, c)
5: m' ← K-PKE.Decrypt(dk_PKE, c)
6: (K', r') ← G(m' || h)
7: K̄ ← J(z || c)
8: c' ← K-PKE.Encrypt(ek_PKE, m', r')
9: if c ≠ c' then
10:   K' ← K̄
11: end if
12: return K'
```



FO transform


Cryptography Is Brittle

functionality

```
static OSStatus
SSLVerifySignedServerKeyExchange(SSLContext *ct
    SSLBuffer signedParams, uint8_t *signature,
{
    OSStatus      err;
    ...

    if ((err = SSLHashSHA1.update(&hashCtx, &se
        goto fail;
    if ((err = SSLHashSHA1.update(&hashCtx, &si
        goto fail;
    if ((err = SSLHashSHA1.final(&hashCtx, &has
        goto fail;
    ...

fail:
    SSLFreeBuffer(&signedHashes);
    SSLFreeBuffer(&hashCtx);
    return err;
}
```



Can we tie **security** to
basic functionality?

security

```
85  int PQCLEAN_HQC128_CLEAN_crypto_kem_dec(uint8_t *ss, const uint8_t *ct, const uint8_t *sk) {
86
87      uint8_t result;
88      uint64_t u[VEC_N_SIZE_64] = {0};
89      uint64_t v[VEC_N1N2_SIZE_64] = {0};
90      const uint8_t *pk = sk + SEED_BYTES;
91      uint8_t sigma[VEC_K_SIZE_BYTES] = {0};
92      uint8_t theta[SHAKE256_512_BYTES] = {0};
93      uint64_t u2[VEC_N_SIZE_64] = {0};
94      uint64_t v2[VEC_N1N2_SIZE_64] = {0};
95      uint8_t mc[VEC_K_SIZE_BYTES + VEC_N_SIZE_BYTES + VEC_N1N2_SIZE_BYTES] = {0};
96      uint8_t tmp[VEC_K_SIZE_BYTES + PUBLIC_KEY_BYTES + SALT_SIZE_BYTES] = {0};
97      uint8_t *m = tmp;
98      uint8_t *salt = tmp + VEC_K_SIZE_BYTES + PUBLIC_KEY_BYTES;
99      shake256incctx shake256state;
100
101      // Retrieving u, v and d from ciphertext
102      PQCLEAN_HQC128_CLEAN_hqc_ciphertext_from_string(u, v, salt, ct);
103
104      // Decrypting
105      result = PQCLEAN_HQC128_CLEAN_hqc_pke_decrypt(m, sigma, u, v, sk);
106
107      // Computing theta
108      memcpy(tmp + VEC_K_SIZE_BYTES, pk, PUBLIC_KEY_BYTES);
109      PQCLEAN_HQC128_CLEAN_shake256_512_ds(&shake256state, theta, tmp, VEC_K_SIZE_BYTES + PUBLIC_
110
111      // Encrypting m'
112      PQCLEAN_HQC128_CLEAN_hqc_pke_encrypt(u2, v2, m, theta, pk);
113
114      // Check if c != c'
115      result |= PQCLEAN_HQC128_CLEAN_vect_compare((uint8_t *)u, (uint8_t *)u2, VEC_N_SIZE_BYTES);
116      result |= PQCLEAN_HQC128_CLEAN_vect_compare((uint8_t *)v, (uint8_t *)v2, VEC_N1N2_SIZE_BYTE
117
118      result = (uint8_t) (-((int16_t) result) >> 15);
119
120      for (size_t i = 0; i < VEC_K_SIZE_BYTES; ++i) {
121          mc[i] = (m[i] & result) ^ (sigma[i] & ~result);
122      }
```

Verifiable Decapsulation

Decaps(sk, c)

05 $m' \leftarrow \text{Dec}(\text{sk}, c)$

06 $(c' \quad) \leftarrow \text{Enc}(\text{pk}, m')$

07 check $c' = c$

08 $K' \leftarrow \text{KDF}(m', \text{pk} \quad)$

09 **return** K'

Verifiable Decapsulation

Enter: **Confirmation Codes**

building on ideas from [\[Fischlin-G'23\]](#)

Decaps(sk, c)

05 $m' \leftarrow \text{Dec}(\text{sk}, c)$

06 (c', cd') $\leftarrow \text{Enc}(\text{pk}, m')$

07 check $c' = c$

08 $K' \leftarrow \text{KDF}(m', \text{pk}, \text{cd}')$

09 **return** K'

Idea: **faulty implementation** of re-encryption \rightarrow noticeable KEM **correctness failure**

ML-KEM with Confirmation Codes

ML-KEM ciphertext compression → lost entropy



leverage lost entropy for confirmation code

Using **12-20 bytes** of confirmation code

detect **faulty re-encryption** in ML-KM-512/768/1024

by **single test** w/ probability **~1/3**

at **≤ 3.4%** performance overhead

Algorithm 14 **K-PKE.Encrypt**($\text{ek}_{\text{PKE}}, m, r$)

Uses the encryption key to encrypt a plaintext message using the randomness r .

Input: encryption key $\text{ek}_{\text{PKE}} \in \mathbb{B}^{384k+32}$.

Input: message $m \in \mathbb{B}^{32}$.

Input: randomness $r \in \mathbb{B}^{32}$.

Output: ciphertext $c \in \mathbb{B}^{32(d_u k + d_v)}$.

1: $N \leftarrow 0$

2: $\hat{\mathbf{t}} \leftarrow \text{ByteDecode}_{12}(\text{ek}_{\text{PKE}}[0 : 384k])$ ▷ run **ByteDecode**₁₂ k times to decode $\hat{\mathbf{t}} \in (\mathbb{Z}_q^{256})^k$

3: $\rho \leftarrow \text{ek}_{\text{PKE}}[384k : 384k + 32]$ ▷ extract 32-byte seed from ek_{PKE}

4: **for** $(i \leftarrow 0; i < k; i++)$ ▷ re-generate matrix $\hat{\mathbf{A}} \in (\mathbb{Z}_q^{256})^{k \times k}$ sampled in Alg. 13

5: **for** $(j \leftarrow 0; j < k; j++)$

6: $\hat{\mathbf{A}}[i, j] \leftarrow \text{SampleNTT}(\rho \| j \| i)$ ▷ j and i are bytes 33 and 34 of the input

7: **end for**

8: **end for**

9: **for** $(i \leftarrow 0; i < k; i++)$ ▷ generate $\mathbf{y} \in (\mathbb{Z}_q^{256})^k$

10: $\mathbf{y}[i] \leftarrow \text{SamplePolyCBD}_{\eta_1}(\text{PRF}_{\eta_1}(r, N))$ ▷ $\mathbf{y}[i] \in \mathbb{Z}_q^{256}$ sampled from CBD

11: $N \leftarrow N + 1$

12: **end for**

13: **for** $(i \leftarrow 0; i < k; i++)$ ▷ generate $\mathbf{e}_1 \in (\mathbb{Z}_q^{256})^k$

14: $\mathbf{e}_1[i] \leftarrow \text{SamplePolyCBD}_{\eta_2}(\text{PRF}_{\eta_2}(r, N))$ ▷ $\mathbf{e}_1[i] \in \mathbb{Z}_q^{256}$ sampled from CBD

15: $N \leftarrow N + 1$

16: **end for**

17: $\mathbf{e}_2 \leftarrow \text{SamplePolyCBD}_{\eta_2}(\text{PRF}_{\eta_2}(r, N))$ ▷ sample $\mathbf{e}_2 \in \mathbb{Z}_q^{256}$ from CBD

18: $\hat{\mathbf{y}} \leftarrow \text{NTT}(\mathbf{y})$ ▷ run **NTT** k times

19: $\mathbf{u} \leftarrow \text{NTT}^{-1}(\hat{\mathbf{A}}^\top \circ \hat{\mathbf{y}}) + \mathbf{e}_1$ ▷ run **NTT**⁻¹ k times

20: $\mu \leftarrow \text{Decompress}_{d_u}(\text{ByteDecode}_1(m))$

21: $\mathbf{v} \leftarrow \text{NTT}^{-1}(\hat{\mathbf{t}}^\top \circ \hat{\mathbf{y}}) + \mathbf{e}_2 + \mu$ ▷ encode plaintext m into polynomial v

22: $c_1 \leftarrow \text{ByteEncode}_{d_u}(\text{Compress}_{d_u}(\mathbf{u}))$ ▷ run **ByteEncode** _{d_u} and **Compress** _{d_u} k times

23: $c_2 \leftarrow \text{ByteEncode}_{d_v}(\text{Compress}_{d_v}(v))$

24: **cd** $\leftarrow (\mathbf{u}[1][S], \dots, \mathbf{u}[k][S], v[S])$


return $(c = c_1 \| c_2, \text{cd})$

Summary

(HYBRID) OBFUSCATION

Kameleon: obfuscate ML-KEM pk/ctxt

– pk even 2.5% smaller



OEINC: hybrid KEM obfuscation

full versions @ IACR ePrint:

– Kameleon: ia.cr/2024/1086

<https://datatracker.ietf.org/doc/draft-veitch-kameleon/>

– hybrid OKEMs: ia.cr/2025/408

– Verifiable Decaps: ia.cr/2025/450

VERIFIABLE DECAPSULATION

functionality  security
conf. code

ML-KEM: 12-20B → detect prob. ~1/3

HQC: 1B → basic tests catch bug

Thank You!

mail@**felixguenther.info**