


Kemeleon Encodings

Making ML-KEM public keys and ciphertexts indistinguishable from random

6 November 2024


IETF121: CFRG

Previously: Elligator



- Mapping elliptic curve points to uniform random bytestrings.
- Used in protocols for censorship-resistance, hash-to-curve, VRFs, OPRFs, and more...
- Specified in RFC 9380.

Previously: Elligator



- Mapping elliptic curve points to uniform random bytestrings.
- Used in protocols for censorship-resistance, hash-to-curve, VRFs, OPRFs, and more...
- Specified in RFC 9380.

- As we transition to post-quantum algorithms, there is a need for a replacement.
- ML-KEM public keys and ciphertexts are distinguishable from random.

Kameleon

ML-KEM/Kyber public keys

- vector of polynomials with coefficients mod q

$[a_1] [a_2] [a_3] \dots [a_b] \quad a_i \in \mathbb{Z}_q$ ($q=3329$, each a_i requires 12 bits)

↑ ↑ ↑ ↑

most sig. bit of each value biased towards 0

Kameleon

ML-KEM/Kyber public keys

- vector of polynomials with coefficients mod q

$$[a_1] [a_2] [a_3] \dots [a_b] \quad a_i \in \mathbb{Z}_q \text{ (} q=3329, \text{ each } a_i \text{ requires 12 bits)}$$

↑ ↑ ↑ ↑

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

$$[A = a_1 + a_2 \cdot q + a_3 \cdot q^2 + \dots + a_b \cdot q^{b-1}] \quad A \text{ is a number mod } q^{b-1}$$

↑

most sig. bit still biased towards 0

Kameleon

ML-KEM/Kyber public keys

- vector of polynomials with coefficients mod q

$$[a_1] [a_2] [a_3] \dots [a_b] \quad a_i \in \mathbb{Z}_q \text{ (} q=3329, \text{ each } a_i \text{ requires 12 bits)}$$

↑ ↑ ↑ ↑

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

$$[A = a_1 + a_2 \cdot q + a_3 \cdot q^2 + \dots + a_b \cdot q^{b-1}] \quad A \text{ is a number mod } q^{b-1}$$

↑

most sig. bit still biased towards 0

Kemeleon

ML-KEM/Kyber ciphertexts

- vector of polynomials with coefficients mod q

22: $c_1 \leftarrow \text{ByteEncode}_{d_u} (\text{Compress}_{d_u} (\mathbf{u}))$

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

24: **return** $c \leftarrow (c_1 \| c_2)$

- compressed before being returned by Encap
- compressed ciphertexts are no longer uniformly distributed (compress performs some rounding)

Kemeleon

ML-KEM/Kyber ciphertexts

- vector of polynomials with coefficients mod q

22: $c_1 \leftarrow \text{ByteEncode}_{d_u} (\text{Compress}_{d_u} (\mathbf{u}))$

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

24: **return** $c \leftarrow (c_1 \| c_2)$

- compressed before being returned by Encap
- compressed ciphertexts are no longer uniformly distributed (compress performs some rounding)
- Encoding for ciphertexts:
 1. Decompress and "recover" randomness lost from compression
 2. Apply previous technique for public keys

Kemeleon Variants

1. No rejection sampling:
 - Use a technique from Elligator squared (Tibouchi) to avoid rejection (100% success probability), at a cost of slightly bigger output

Kemeleon Variants

1. No rejection sampling:
 - Use a technique from Elligator squared (Tibouchi) to avoid rejection (100% success probability), at a cost of slightly bigger output
2. Deterministic:
 - Public keys: Store randomness in secret key
 - Ciphertexts: Derive from the shared secret key K , a "new" key K' (to be used as the shared secret) and a seed to be used to generate randomness for Kemeleon encoding

Kemeleon Variants

1. No rejection sampling:
 - Use a technique from Elligator squared (Tibouchi) to avoid rejection (100% success probability), at a cost of slightly bigger output
2. Deterministic:
 - Public keys: Store randomness in secret key
 - Ciphertexts: Derive from the shared secret key K , a "new" key K' (to be used as the shared secret) and a seed to be used to generate randomness for Kemeleon encoding
3. Faster:
 - Kemeleon encoding requires computing with large (~ 750 to 1500 byte) integers
 - At the cost of a higher rejection rate, we can compute with coefficients mod $q-1 = 13 \cdot 2^8$.
 - Improved efficiency by using only ~ 237 to 475 byte integers
 - (same output size)

Which variant(s) are you interested in?

	Output size (bytes)	Success probability	Notes
	ML-KEM-512 / 768 / 1024 Public keys: 800 / 1184 / 1568 Ciphertexts: 768 / 1088 / 1568	Elligator: 0.50	
Kemeleon (original)	ML-KEM-512 / 768 / 1024 Public keys: 781 / 1156 / 1530 Ciphertexts: 877 / 1252 / 1658	ML-KEM-512 / 768 / 1024 Public keys: 0.56 / 0.83 / 0.62 Ciphertexts: 0.51 / 0.77 / 0.57	<i>Larger integer arithmetic</i> (~750 to 1500 byte integers)
No rejection	<i>Kemeleon + 16 in all cases</i>	1.00	<i>Larger integer arithmetic</i> (~775 to 1515 byte integers)
Faster	<i>Same as Kemeleon (original)</i>	ML-KEM-512 / 768 / 1024 Public keys: 0.49 / 0.29 / 0.53 Ciphertexts: 0.45 / 0.25 / 0.47	<i>Smaller integer arithmetic</i> (~237 to 475 byte integers)

WIP Draft: <https://github.com/ssveitch/draft-kemeleon>

Contact: shannon.veitch@inf.ethz.ch