The Fujisaki-Okamoto Transform in the Quantum Random Oracle Model

Based on

V. Kuchta, et al. "Measure-Rewind-Measure: Tighter Quantum Random Oracle Model Proofs for One-Way to Hiding and CCA Security." EUROCRYPT 2020, LNCS 12107, pp. 703-728.

August 25, 2020

Carl A. Miller

The Big Picture

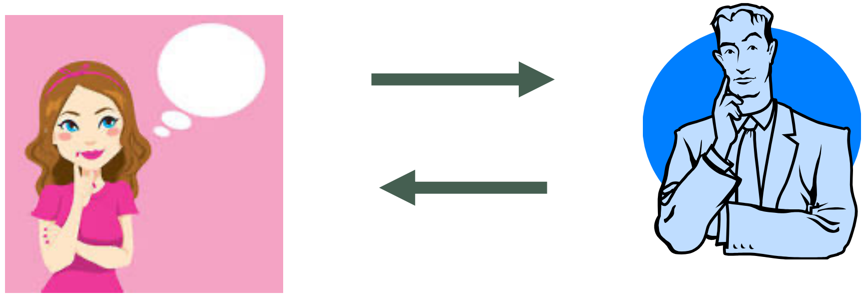

A public-key encryption scheme secure against chosen **plaintext** attacks

\+

H

A hash function

Fujisaki-Okamoto

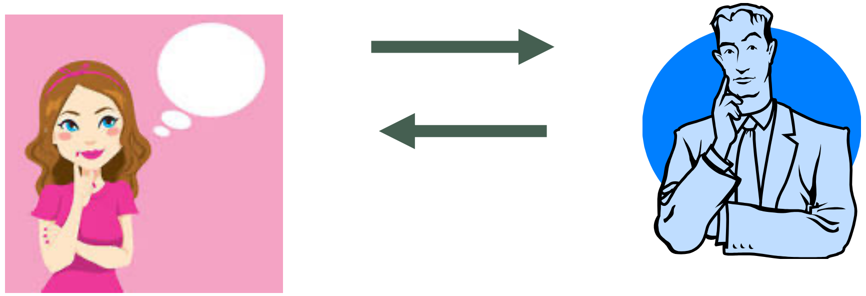

A KEM secure against chosen **ciphertext** attacks

The Big Picture

Fujisaki-Okamoto is known to be secure in the **random oracle model** for H.

(Uniform output for every new input)

The Big Picture

Fujisaki-Okamoto is known to be secure in the **random oracle model** for H.

What about the **quantum** random oracle model?

$$\frac{|x_1\rangle + |x_2\rangle}{\sqrt{2}}$$

$$\frac{|x_1, f(x_1)\rangle + |x_2, f(x_2)\rangle}{\sqrt{2}}$$

The Big Picture

This paper shows a **tighter** QROM proof of security for Fujisaki-Okamoto, under some conditions.

In this talk I'll give a (heavily simplified) overview of the proof and the main result.

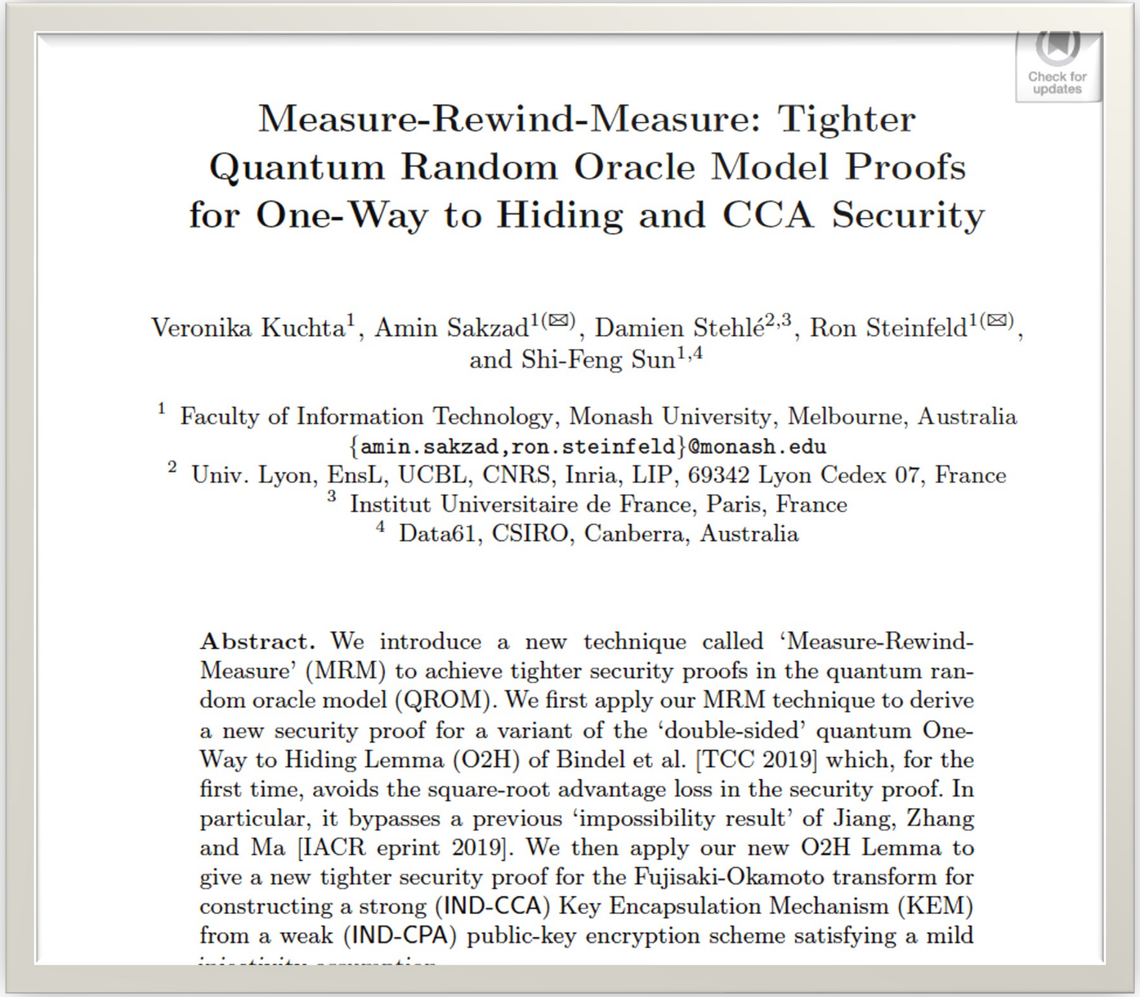

Check for updates

Measure-Rewind-Measure: Tighter Quantum Random Oracle Model Proofs for One-Way to Hiding and CCA Security

Veronika Kuchta[1], Amin Sakzad[1(✉)], Damien Stehlé[2,3], Ron Steinfeld[1(✉)], and Shi-Feng Sun[1,4]

[1] Faculty of Information Technology, Monash University, Melbourne, Australia
{amin.sakzad,ron.steinfeld}@monash.edu
[2] Univ. Lyon, EnsL, UCBL, CNRS, Inria, LIP, 69342 Lyon Cedex 07, France
[3] Institut Universitaire de France, Paris, France
[4] Data61, CSIRO, Canberra, Australia

Abstract. We introduce a new technique called 'Measure-Rewind-Measure' (MRM) to achieve tighter security proofs in the quantum random oracle model (QROM). We first apply our MRM technique to derive a new security proof for a variant of the 'double-sided' quantum One-Way to Hiding Lemma (O2H) of Bindel et al. [TCC 2019] which, for the first time, avoids the square-root advantage loss in the security proof. In particular, it bypasses a previous 'impossibility result' of Jiang, Zhang and Ma [IACR eprint 2019]. We then apply our new O2H Lemma to give a new tighter security proof for the Fujisaki-Okamoto transform for constructing a strong (IND-CCA) Key Encapsulation Mechanism (KEM) from a weak (IND-CPA) public-key encryption scheme satisfying a mild

The Quantum Random Oracle Model

A Crash Course

Let X and Y be finite sets.

A **quantum random oracle** is initiated by choosing a function $f : X \to Y$ uniformly at random.

It operates as shown below.

$$\frac{|x_1\rangle + |x_2\rangle}{\sqrt{2}} \longrightarrow$$

$$\longrightarrow \frac{|x_1, f(x_1)\rangle + |x_2, f(x_2)\rangle}{\sqrt{2}}$$

formal linear sums from X.

formal linear sums from $X \times Y$.

A Crash Course

If w… immediately,

the…

But… y operations).

A key point:

There are two basic operations in quantum information: **unitary operations**, and **measurements.**

Unitary operations are always reversible. Measurements typically are not.

$$\frac{|\ldots}{\sqrt{2}} \quad \ldots \quad \frac{\ldots_1))\rangle + |x_2, f(x_2))\rangle}{\sqrt{2}}$$

A Crash Course

If we merely **measure** the outcome of the oracle immediately, then it's basically just a classical random oracle.
But there are other things we can do (i.e., unitary operations).

$$\frac{|x_1\rangle + |x_2\rangle}{\sqrt{2}}$$

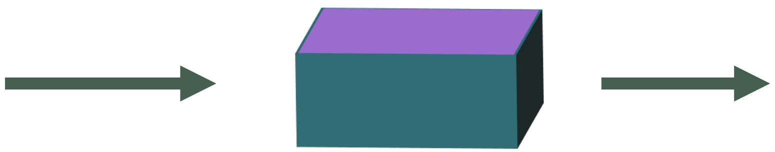

$$\frac{|x_1, f(x_1)\rangle + |x_2, f(x_2)\rangle}{\sqrt{2}}$$

The Fujisaki-Okamoto Transform

("This transform and its variants are used in all public-key encryption schemes and key establishment algorithms of the second round of the NIST PQC standardization process.")

Starting Point

We have a PK encryption protocol (KeyGen, Enc, Dec) which is IND-CPA secure.

KeyGen, Dec

Enc

Starting Point

We have a PK encryption protocol (KeyGen, Enc, Dec) which is IND-CPA secure.

(Meaning, Eve cannot distinguish between the encryptions of two chosen plaintexts.)

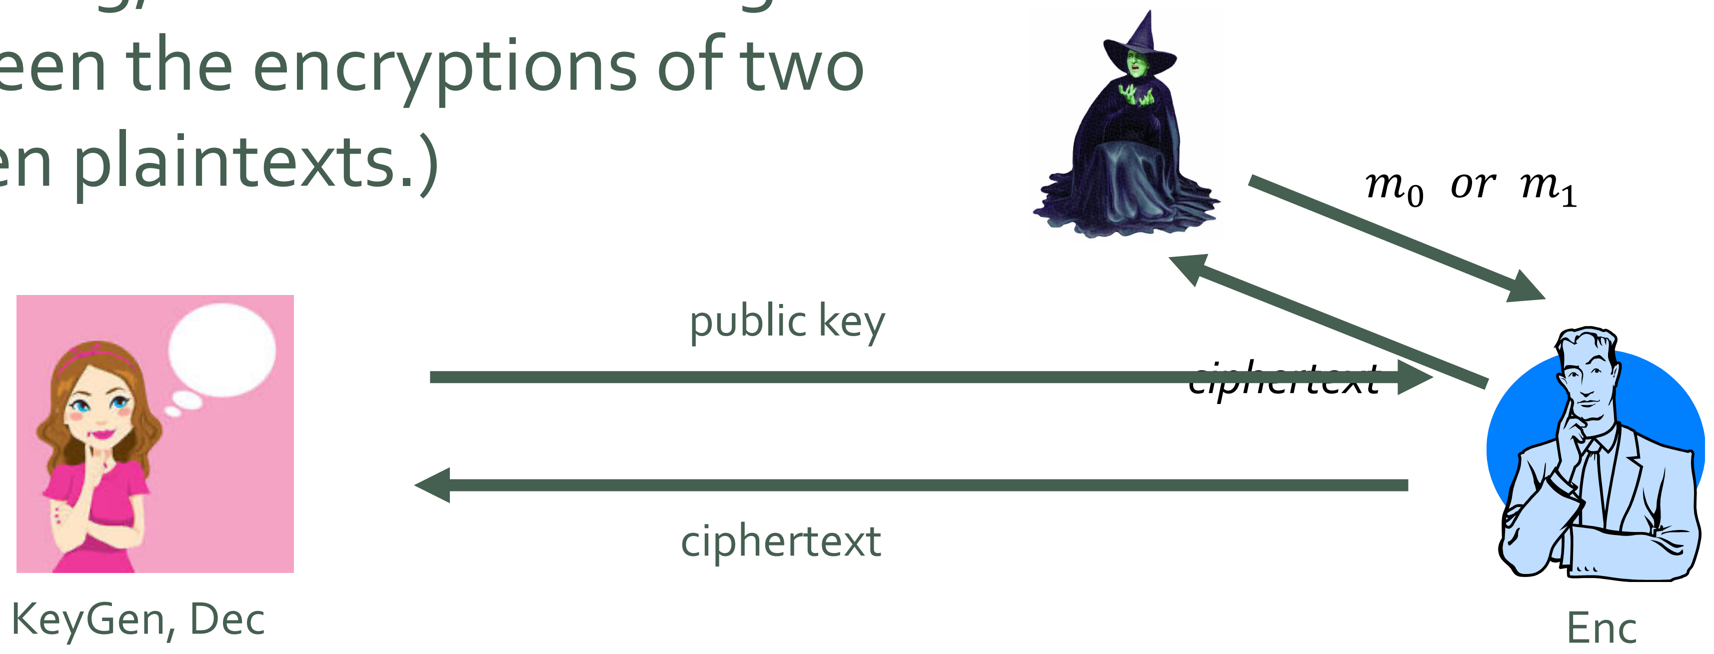

Starting Point

We want an IND-**CCA** secure KEM (KeyGen', Encaps, Decaps).
Idea: Use a hash function to strengthen security.

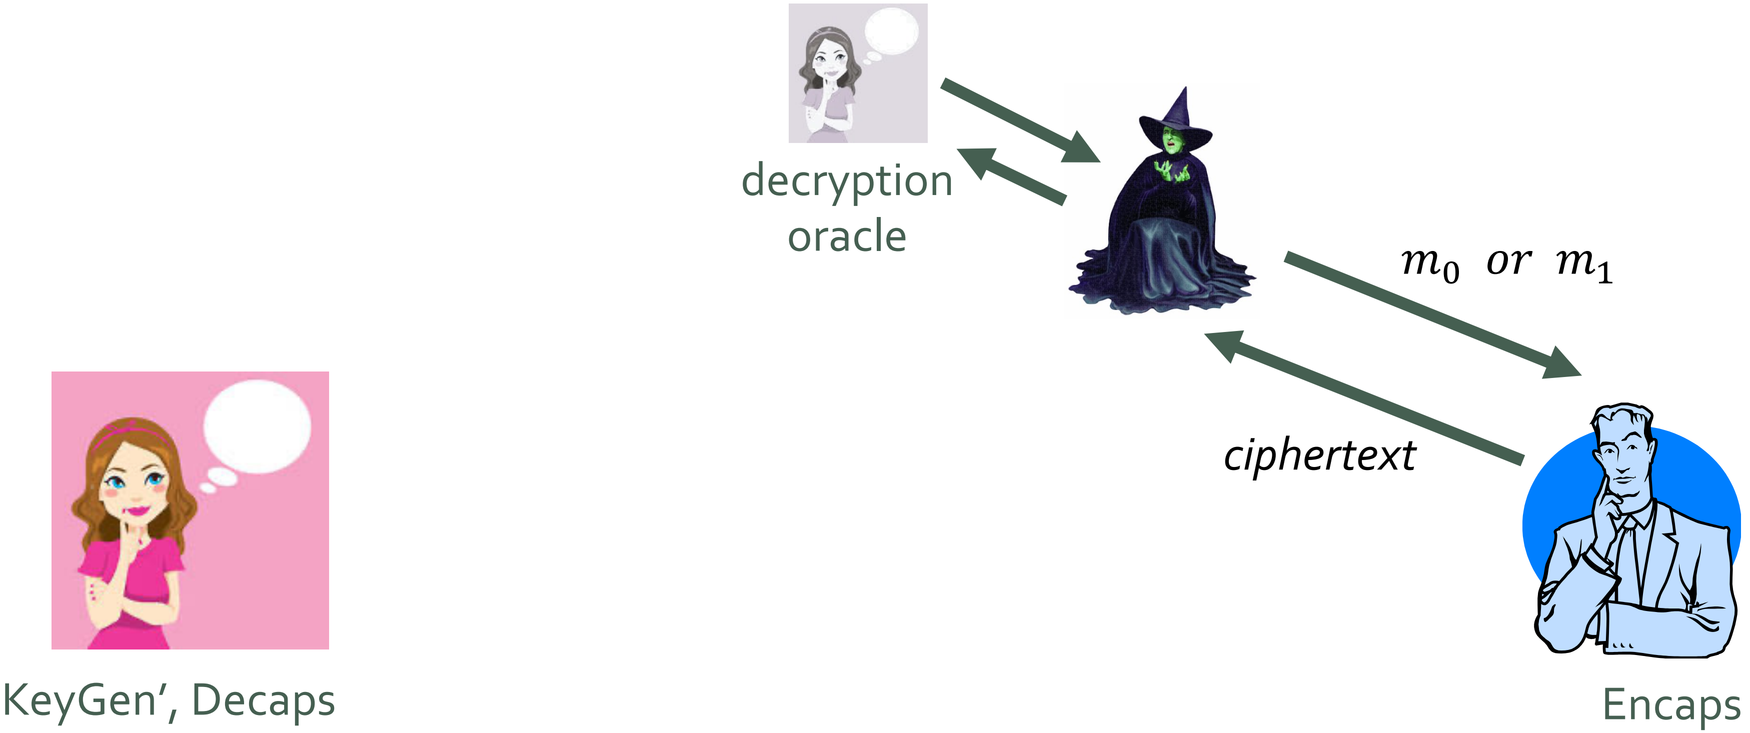

Building the KEM

k is “the key”

1. Bob generates a uniformly random m, sets c=Enc(m).
2. He sends c and computes k := H(c,m).
3. Alice sets m’=Dec(c), and computes k’ = H(c,m’).

Building the KEM

1. Bob generates a uniformly random m, sets c=Enc(m).
2. He sends c and computes k := H(c,m).
3. Alice sets m’=Dec(c), and computes k’ = H(c,m’) **and checks that c=Enc(m’).**

(Why the extra step?)

Building the KEM

1. Bob generates a uniformly random m, sets c=Enc(m).
2. He sends c and computes k := H(c,m).
3. Alice sets m'=Dec(c), and computes k' = H(c,m') and checks that c=Enc(m').

Problem: Enc might be a random algorithm. (Can't redo it.)

Fix: Derandomize it first. (Downgrades it to "OW-CPA".)

Building the KEM

"implicit rejection"

1. Bob generates a uniformly random m, sets c=Enc(m).
2. He sends c and computes k := H(c,m).
3. Alice sets m'=Dec(c), and computes k' = H(c,m') and checks that c=Enc(m').

Problem: What happens when Alice's step 3 fails?

Fix: Have her generate a fake response pseudorandomly.

Building the KEM

1. Bob generates a uniformly random m, sets c=Enc(m).
2. He sends c and computes k := H(c,m).
3. Alice sets m′=Dec(c), and computes k′ = H(c,m′) and checks that c=Enc(m′).

The Fujisaki-Okamoto is basically the above procedure, with additional "fixes" added in.

IND-CCA Security Proof

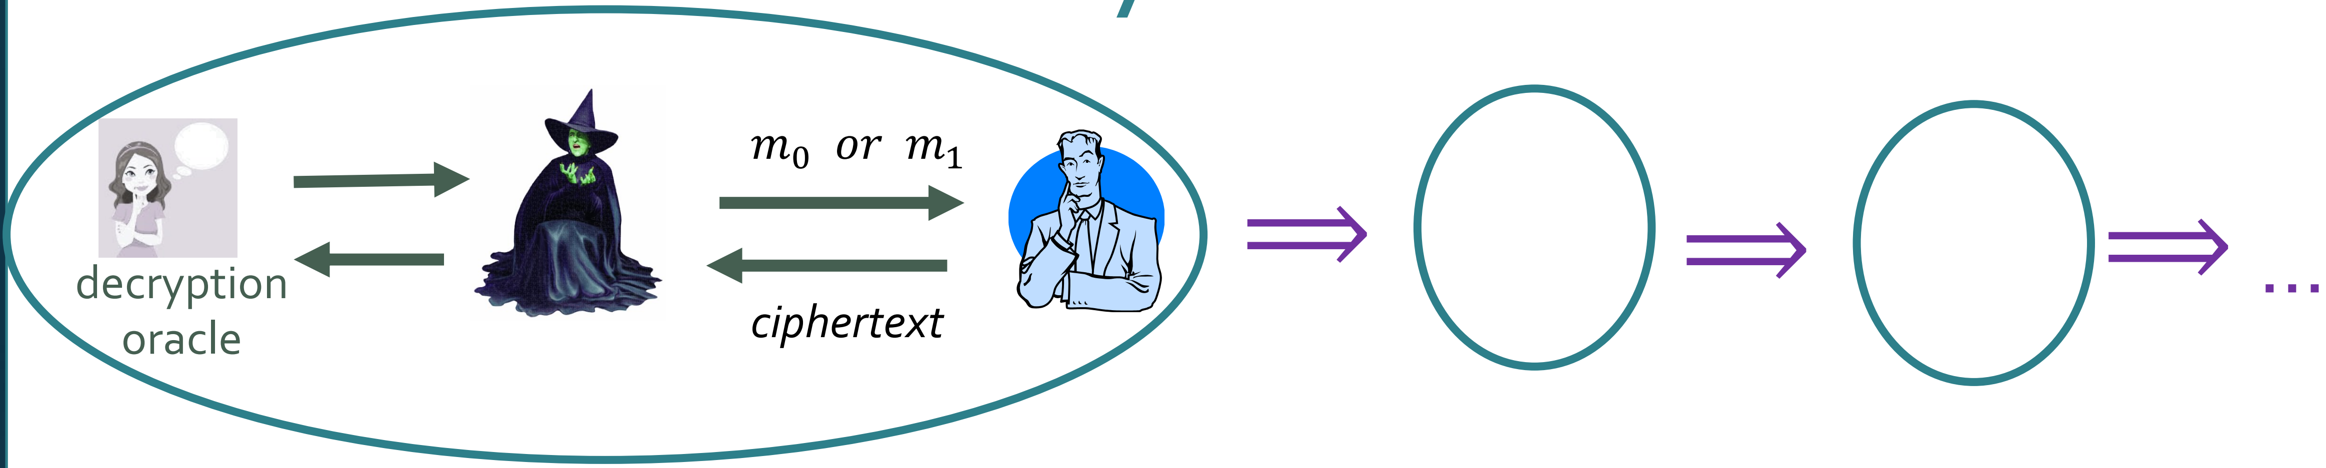

A CCA-hack of the KEM we've constructed ...

... implies a series of other types of hacks ...

... which implies a one-way hack of the original PKE scheme.

IND-CCA Security Proof

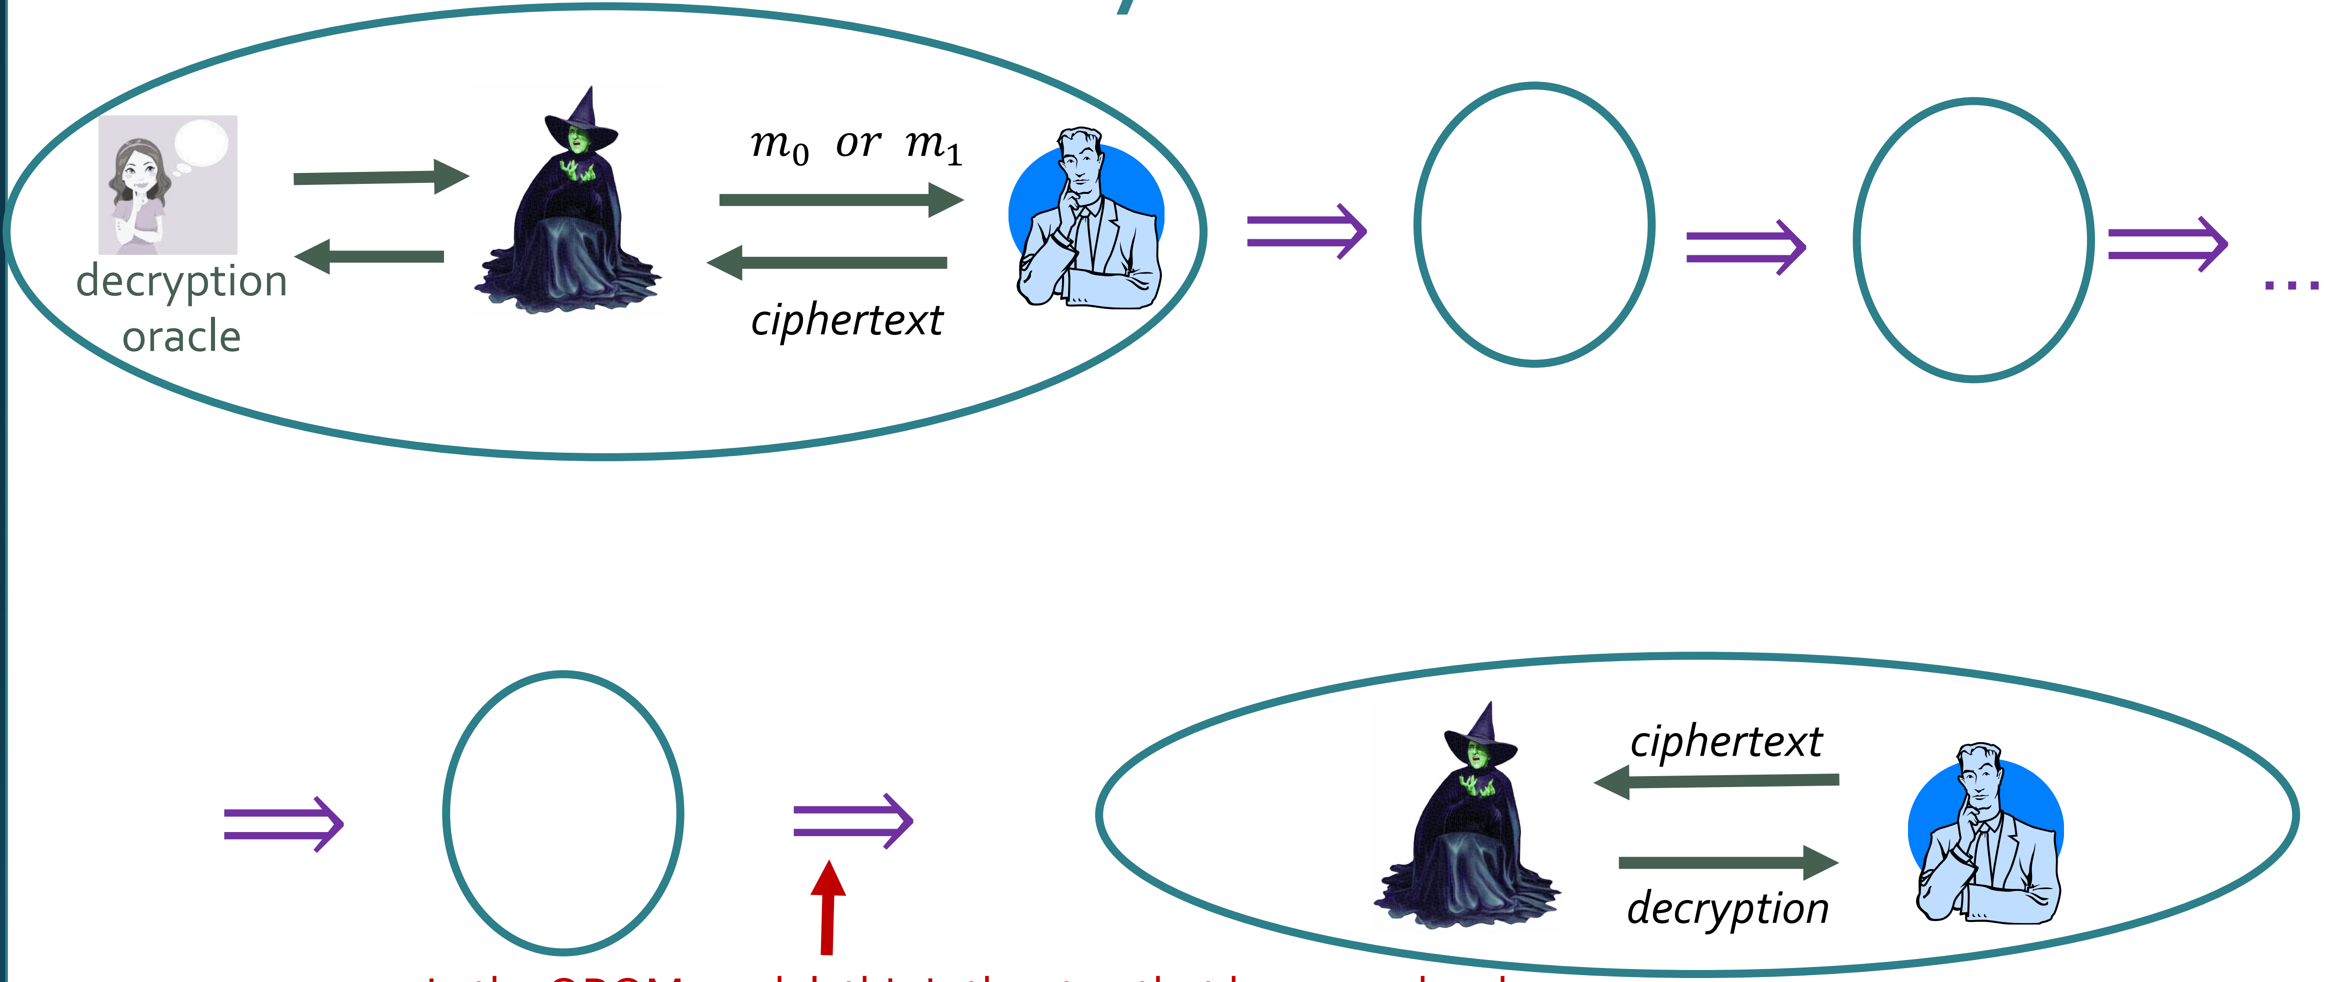

In the QROM model, this is the step that becomes hard.
("One-way to hiding lemma.")

One-Way to Hiding Lemmas

The Two-Oracles Problem

Let X, Y be finite sets.

Let $G, H: X \to Y$ be random functions such that $G = H$ everywhere outside of a subset $S \subseteq X$.

Problem: Eve wants to distinguish G from H, via oracle access.

Let's also assume that Eve has a "hint" z.
(z = random variable correlated with G,H,S).

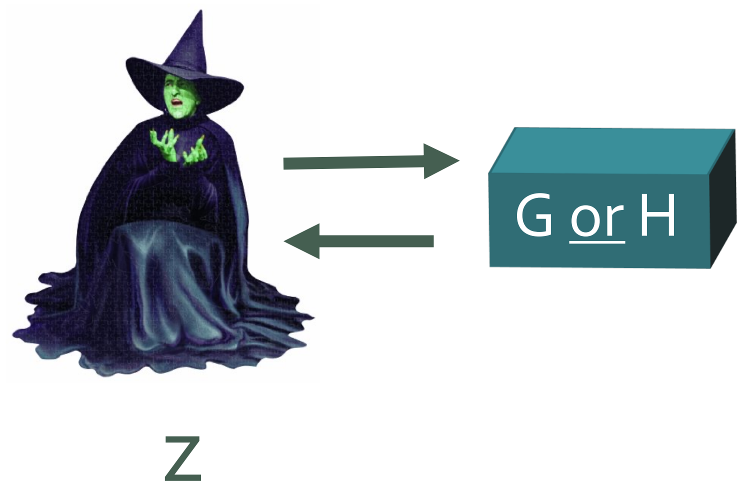

The Two-Oracles Problem

Intuition: This is like an IND experiment. Think of z as a public-key encryption of the set S.

z

The Two-Oracles Problem

It is not hard to show that **if Eve can distinguish G from H efficiently, then she can also guess an element of S efficiently.**

This is a classical **"one-way to hiding lemma,"** and it can be used to prove classical security for Fujisaki-Okamoto.

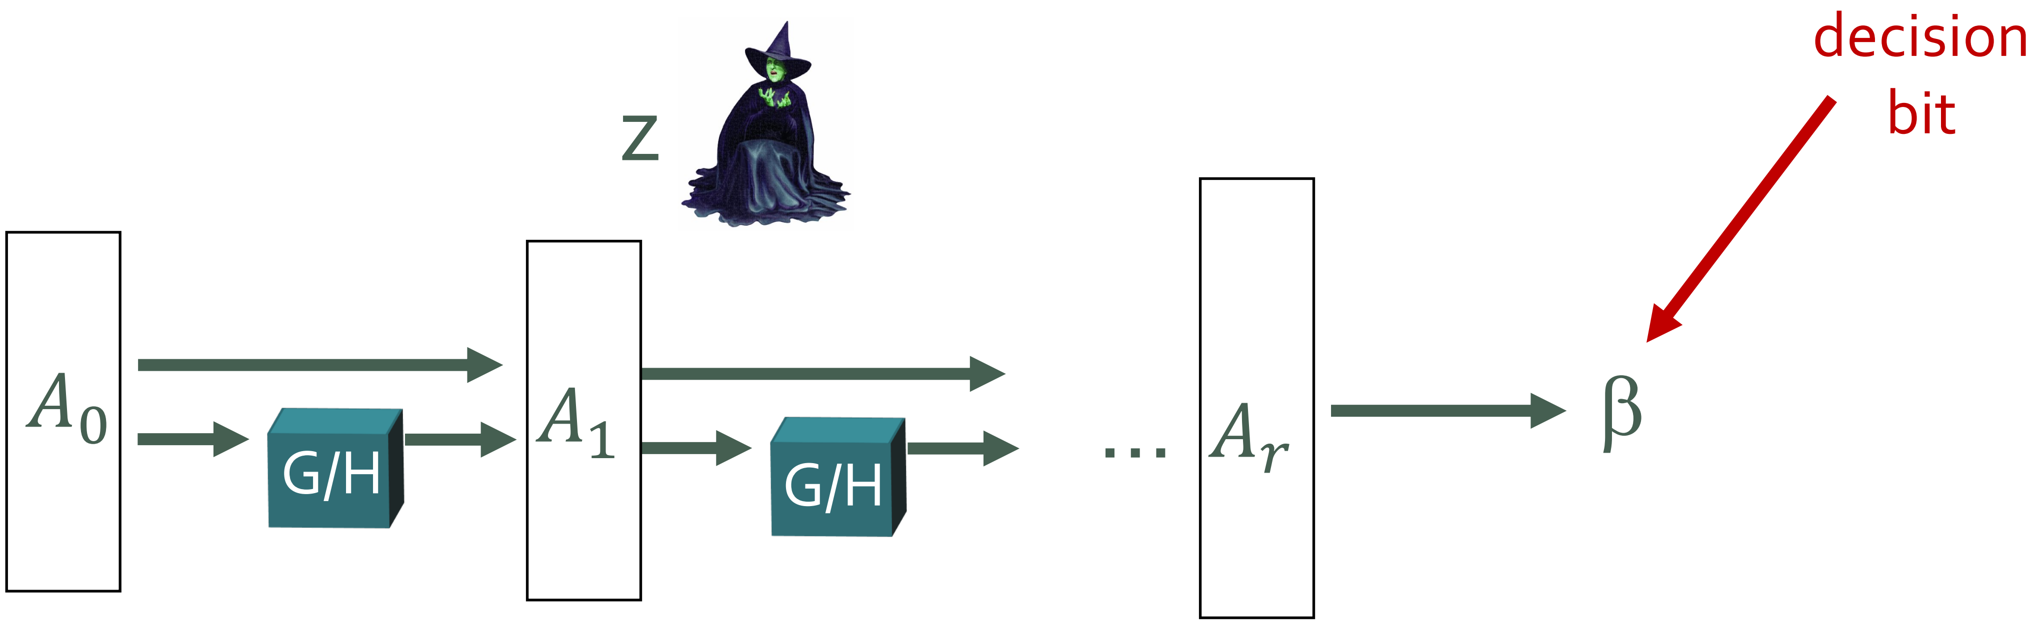

Quantum One-Way to Hiding

Can we prove the same if the unknown oracle is a quantum oracle?

Previous approach: Choose random $i \in \{1, \ldots, d-1\}$. Run distinguisher until just before the ith query, and then measure input register.

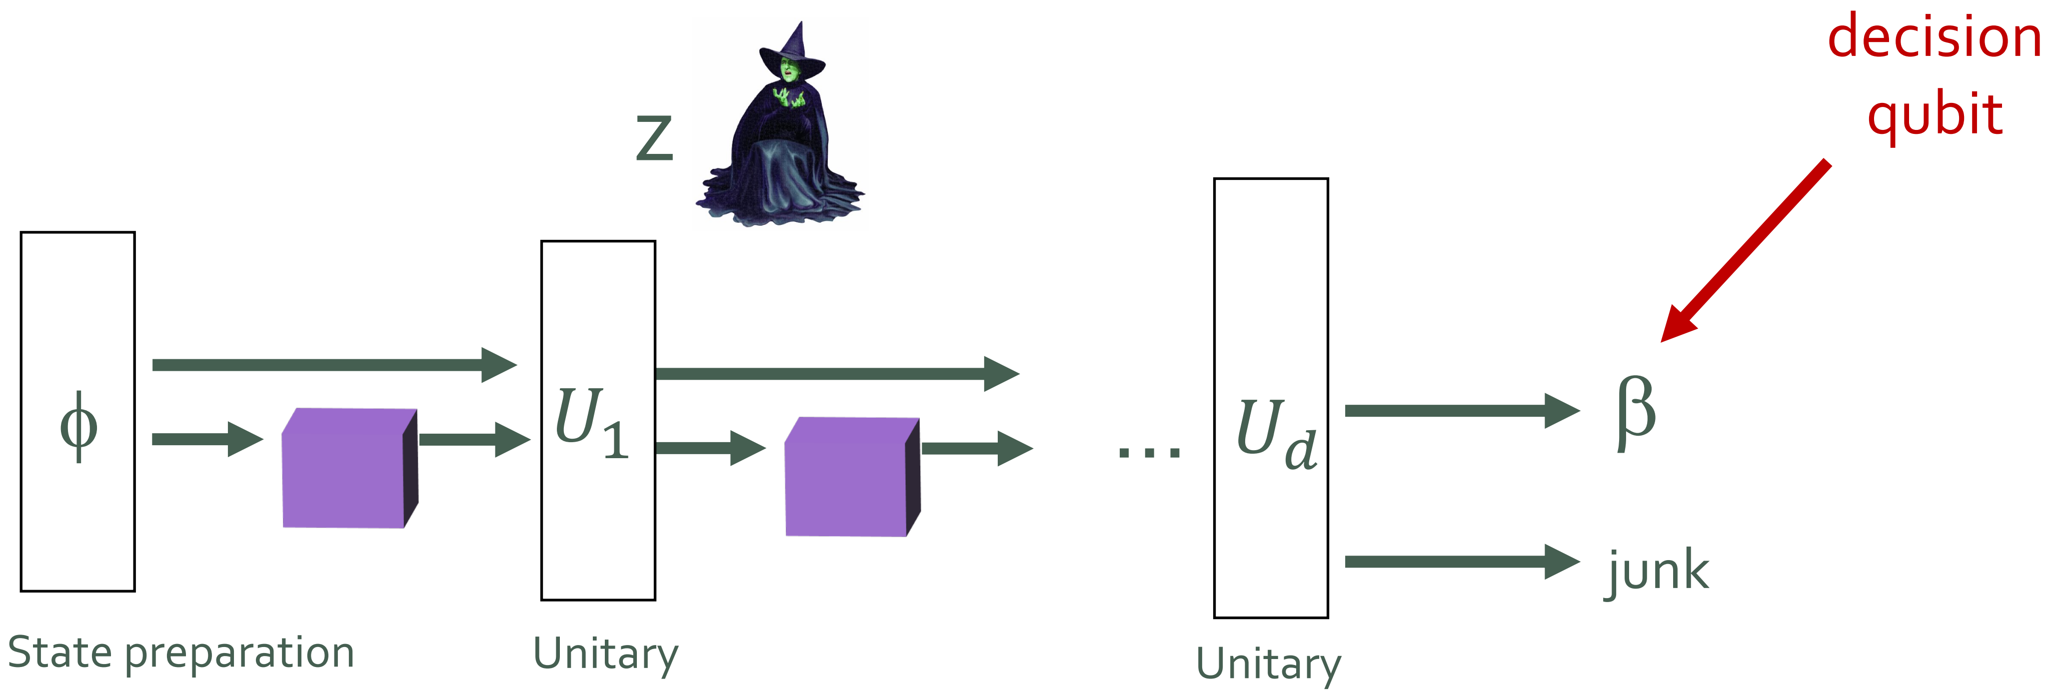

Quantum One-Way to Hiding

Can we prove the same if the unknown oracle is a quantum oracle?

Previous approach: Choose random $i \in \{1, \dots, d-1\}$. Run distinguisher until just before the *i*th query, and then measure input register.

This works, but it's got a square-root loss in effectiveness.

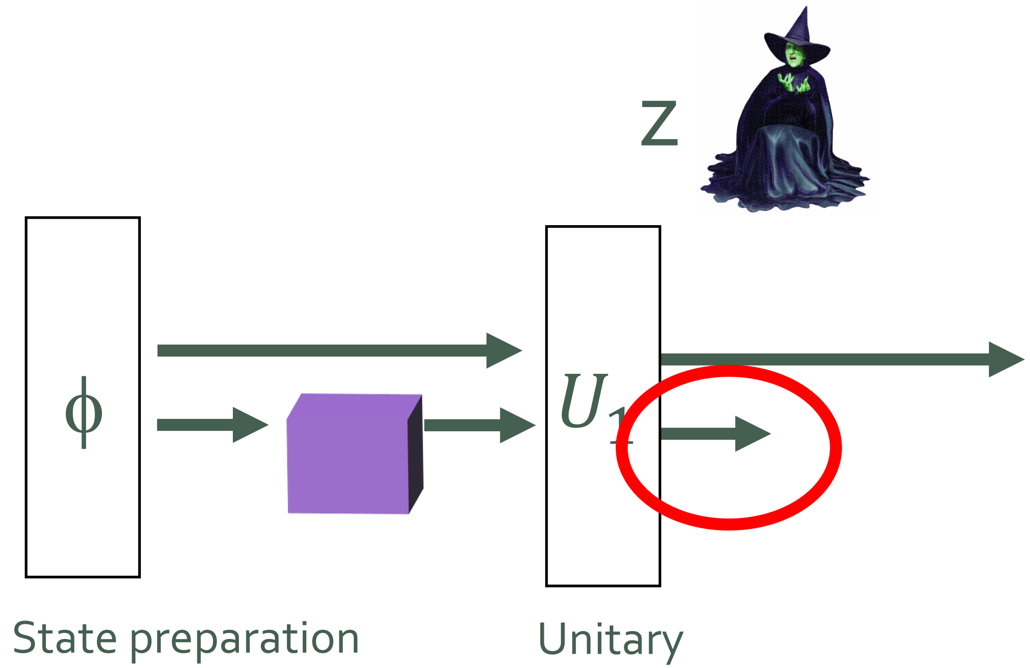

Quantum One-Way to Hiding

New approach [Kuchta '20]:

1. Run full algorithm and measure the decision qubit.
2. Rewind back to before ith round and measure the input register.

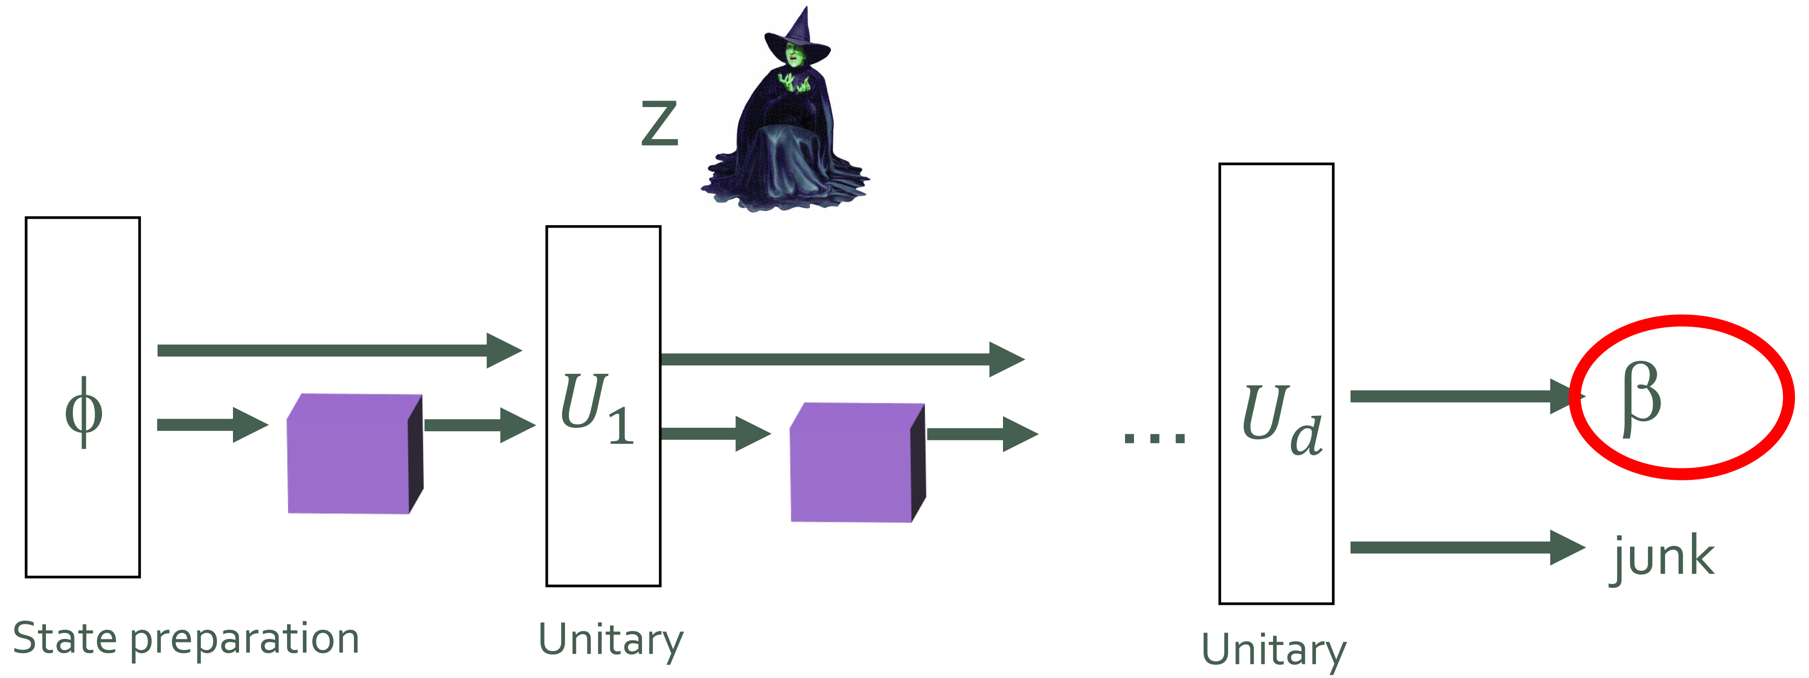

Quantum One-Way to Hiding

New approach [Kuchta '20]:

1. Run full algorithm and measure the decision qubit.

2. Rewind back to before ith round and measure the input register.

Step 1 magnifies the success probability. (No square-root loss.)

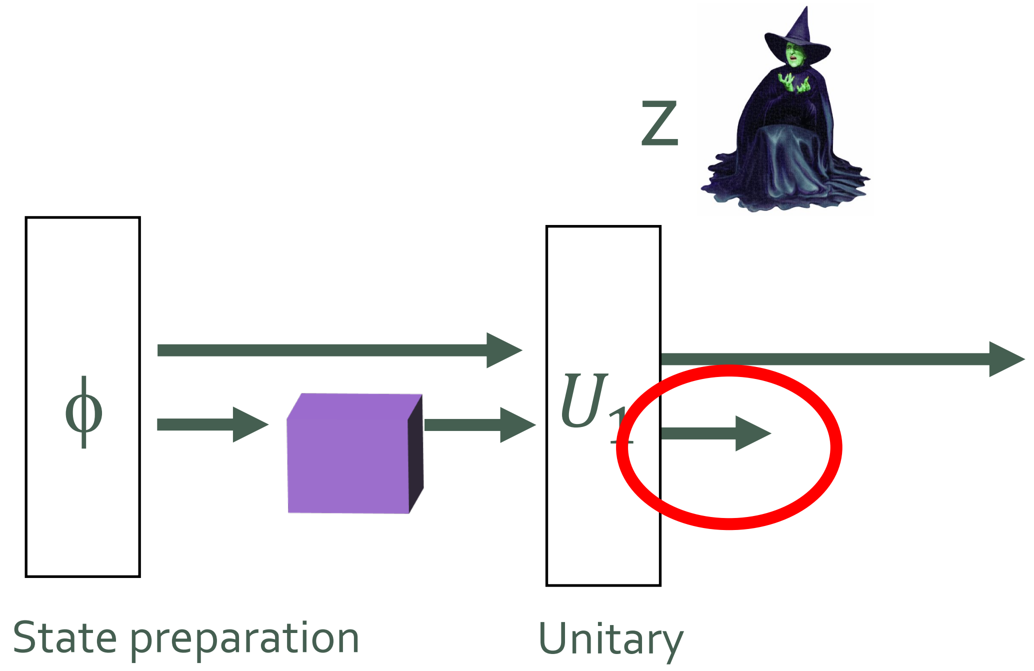

Outcome

Main Result

Goal: Show the tightest possible upper bound on the probability that a CCA-adversary can break a Fujisaki-Okamoto KEM.

	CCA bound	Security loss	Weak scheme
[10]	$q^{3/2} \cdot \varepsilon^{1/4}$	$3\lambda + 9 \log q$	IND-CPA
[11, 13, 15]	$d^{1/2} \cdot \varepsilon^{1/2}$	$\lambda + \log d$	IND-CPA
[5]	$d^{1/2} \cdot \varepsilon^{1/2}$	$\lambda + \log d$	IND-CPA injective
This work	$d^2 \cdot \varepsilon$	$4 \log d$	IND-CPA injective

(From source paper)

λ = target # of security bits
ε = probability that adversary can break original scheme
q = total # of hash function uses by adversary
d = sequential # of uses of hash function

Meaning of "IND-CPA injective"

Let E = (KeyGen, Enc, Dec) be a PKE scheme.

Recall that the 1st step of Fujisaki-Okamoto is to derandomize. If

$$Enc_{pk}(m) = F\,(pk, m, coins),$$

Then let

$$Enc_{pk}^{d}\,(m) = F\left(pk, m, H(m)\right).$$

The scheme E is ***η*-injective** if, with probability $\geq 1 - \eta$, the map Enc_{pk}^{d} is injective.

Question: How applicable is this to NIST PQC candidate schemes?