Semantic Security and Indistinguishability in the Quantum World

Tommaso Gagliardoni, Andreas Hülsing, Christian Schaffner
(slides by Tommaso, thanks a lot!!!)

University of Amsterdam and CWI

Tuesday, 20 October 2015
Aarhus, Denmark

Let's focus on symmetric-key encryption schemes

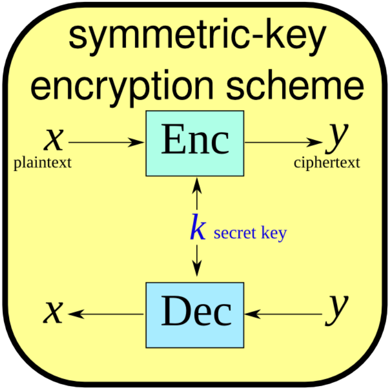

Adversaries

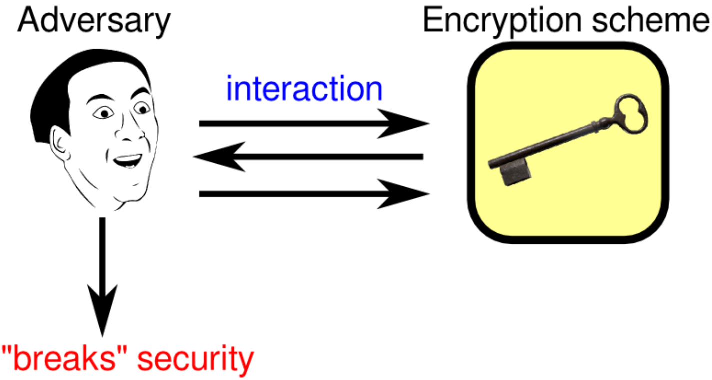

Adversary = PPT circuit family (classical security)

Adversaries

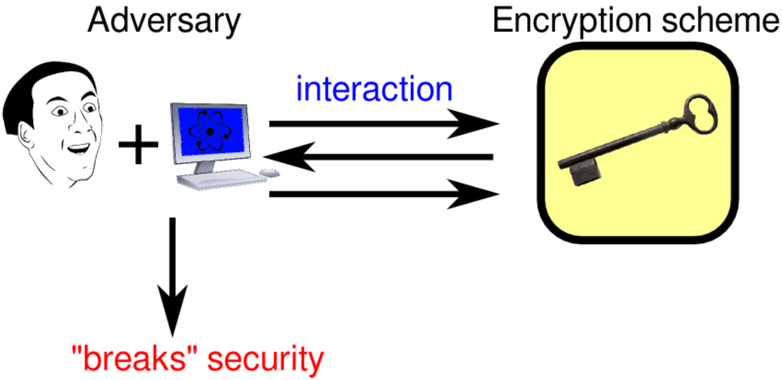

Adversaries

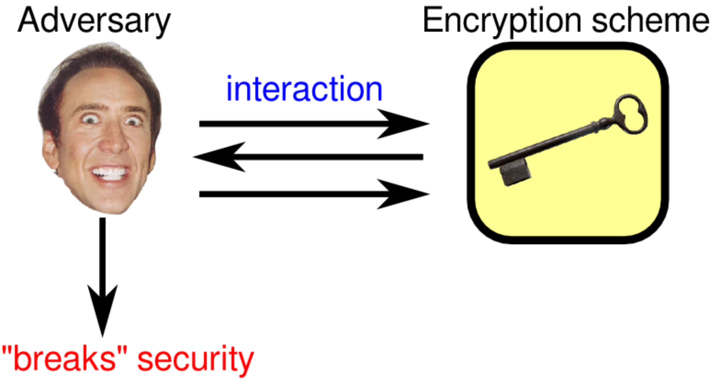

Adversary = QPPT circuit family (post-quantum security)

Not enough

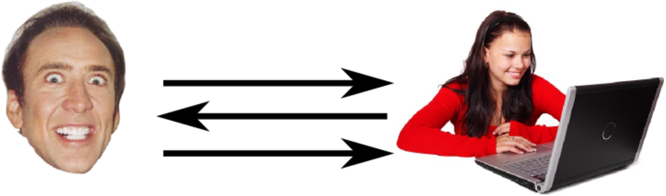

Not enough

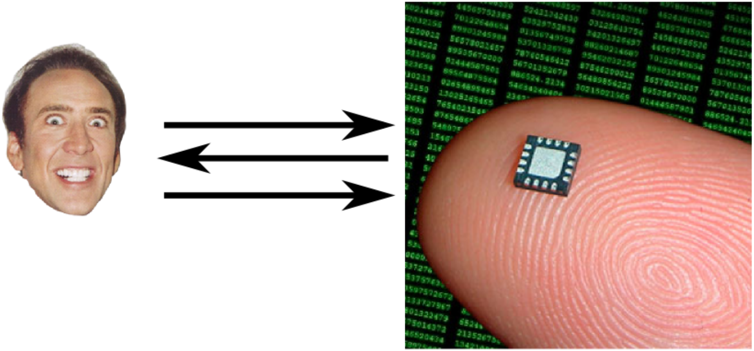

Not enough

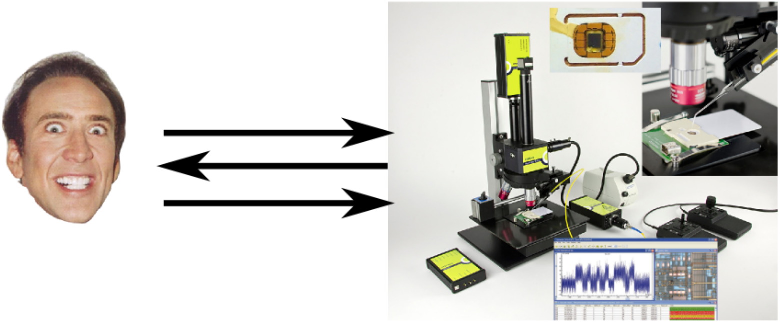

Not enough

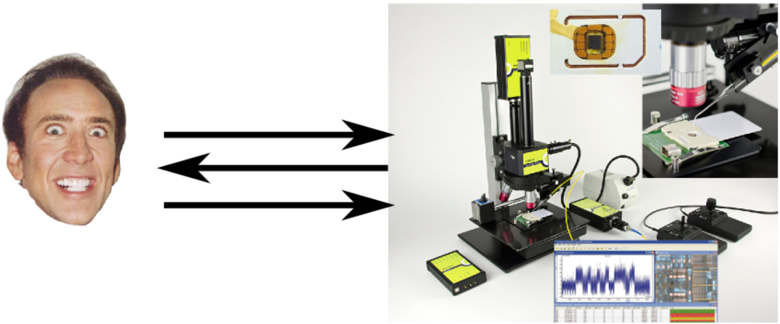

Quantum security beyond post-quantum: quantum interaction with classical schemes

Other examples

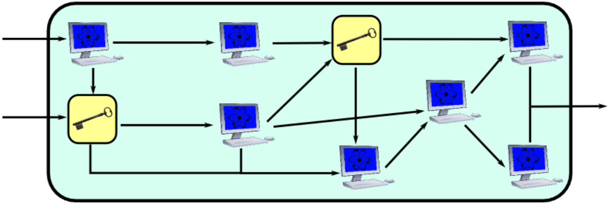

Other examples

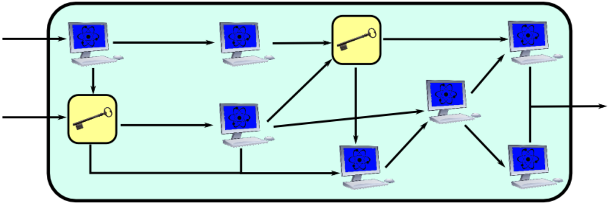

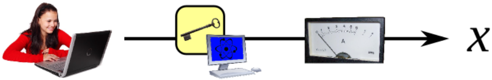

Other examples

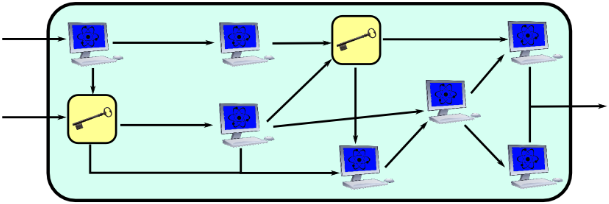

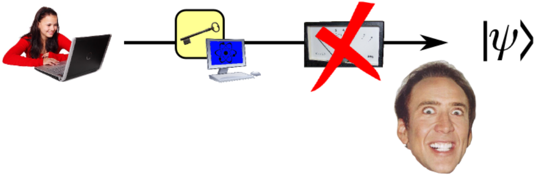

Previous work

[DFNS13] Ivan Damgård, Jesper Buus Nielsen, Jakob Løvstad Funder, Louis Salvail: *"Superposition Attacks on Cryptographic Protocols"*, ICITS 2013

[BZ13] Dan Boneh, Mark Zhandry: *"Secure Signatures and Chosen Ciphertext Security in a Quantum Computing World"*, CRYPTO 2013

Previous work

[DFNS13] Ivan Damgård, Jesper Buus Nielsen, Jakob Løvstad Funder, Louis Salvail: *"Superposition Attacks on Cryptographic Protocols"*, ICITS 2013

[BZ13] Dan Boneh, Mark Zhandry: *"Secure Signatures and Chosen Ciphertext Security in a Quantum Computing World"*, CRYPTO 2013

Model encryption as unitary operator defined by:

$$\sum_{x,y} |x, y\rangle \mapsto \sum_{x,y} |x, \mathrm{Enc}_k(x) \oplus y\rangle$$

(because we want to recover $x \mapsto \mathrm{Enc}_k(x)$ classically)

Results from [BZ13] & Our Contribution

- A 'natural' notion of security (fqIND-qCPA) is unachievable
- Compromise: 'almost classical' notion of security (IND-qCPA)
- IND-qCPA is achievable and stronger than IND-CPA

Results from [BZ13] & Our Contribution

- A ‘natural’ notion of security (fqIND-qCPA) is unachievable
- Compromise: ‘almost classical’ notion of security (IND-qCPA)
- IND-qCPA is achievable and stronger than IND-CPA

Situation:

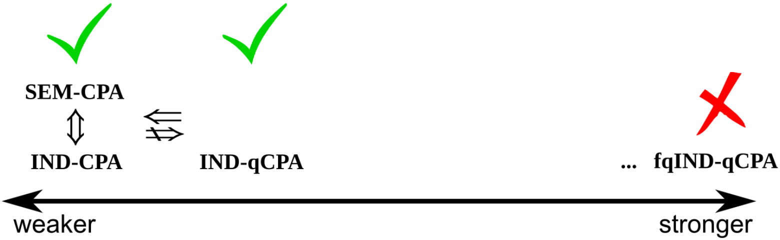

Results from [BZ13] & Our Contribution

- A 'natural' notion of security (fqIND-qCPA) is unachievable
- Compromise: 'almost classical' notion of security (IND-qCPA)
- IND-qCPA is achievable and stronger than IND-CPA

Situation:

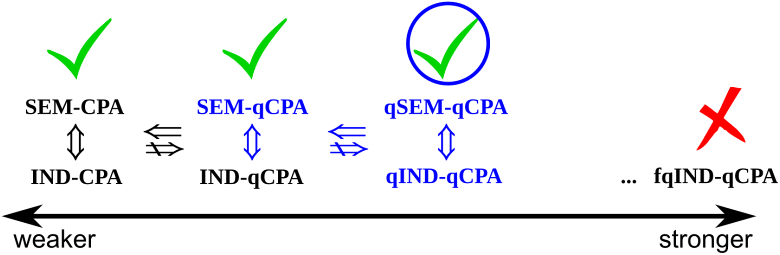

Our contribution!

Classical Indistinguishability (IND)

Game-based security: $\mathcal{A}$ plays an interactive game against a challenger $\mathcal{C}$.

Classical Indistinguishability (IND)

Game-based security: $\mathcal{A}$ plays an interactive game against a challenger $\mathcal{C}$.

IND game (challenge query): $\mathcal{A}$ sends $\mathcal{C}$ two plaintexts $x_0, x_1 \in \mathcal{M}$.

Classical Indistinguishability (IND)

Game-based security: $\mathcal{A}$ plays an interactive game against a challenger $\mathcal{C}$.

IND game (challenge query): $\mathcal{A}$ sends $\mathcal{C}$ two plaintexts $x_0, x_1 \in \mathcal{M}$.
$\mathcal{C}$ flips a random bit $b \xleftarrow{\$} \{0,1\}$,

Classical Indistinguishability (IND)

Game-based security: $\mathcal{A}$ plays an interactive game against a challenger $\mathcal{C}$.

IND game (challenge query): $\mathcal{A}$ sends $\mathcal{C}$ two plaintexts $x_0, x_1 \in \mathcal{M}$.

$\mathcal{C}$ flips a random bit $b \xleftarrow{\$} \{0,1\}$, computes $y \leftarrow \mathsf{Enc}_k(x_b)$,

Classical Indistinguishability (IND)

Game-based security: $\mathcal{A}$ plays an interactive game against a challenger $\mathcal{C}$.

IND game (challenge query): $\mathcal{A}$ sends $\mathcal{C}$ two plaintexts $x_0, x_1 \in \mathcal{M}$.
$\mathcal{C}$ flips a random bit $b \xleftarrow{\$} \{0,1\}$, computes $y \leftarrow \mathsf{Enc}_k(x_b)$, and finally sends ciphertext y to $\mathcal{A}$.

Classical Indistinguishability (IND)

Game-based security: $\mathcal{A}$ plays an interactive game against a challenger $\mathcal{C}$.

IND game (challenge query): $\mathcal{A}$ sends $\mathcal{C}$ two plaintexts $x_0, x_1 \in \mathcal{M}$. $\mathcal{C}$ flips a random bit $b \xleftarrow{\$} \{0,1\}$, computes $y \leftarrow \mathsf{Enc}_k(x_b)$, and finally sends ciphertext y to $\mathcal{A}$. $\mathcal{A}$'s goal is to guess b.

Classical Indistinguishability (IND)

Game-based security: $\mathcal{A}$ plays an interactive game against a challenger $\mathcal{C}$.

IND game (challenge query): $\mathcal{A}$ sends $\mathcal{C}$ two plaintexts $x_0, x_1 \in \mathcal{M}$. $\mathcal{C}$ flips a random bit $b \xleftarrow{\$} \{0,1\}$, computes $y \leftarrow \mathsf{Enc}_k(x_b)$, and finally sends ciphertext y to $\mathcal{A}$. $\mathcal{A}$'s goal is to guess b.

Classical Indistinguishability (IND)

For any efficient adversary $\mathcal{A}$ and any message x_0, x_1:

$$\left| \Pr[\mathcal{A}(y) = b] - \frac{1}{2} \right| \leq \mathsf{negl}\,(n)\,.$$

Classical Indistinguishability (IND)

Game-based security: $\mathcal{A}$ plays an interactive game against a challenger $\mathcal{C}$.

IND game (challenge query): $\mathcal{A}$ sends $\mathcal{C}$ two plaintexts $x_0, x_1 \in \mathcal{M}$. $\mathcal{C}$ flips a random bit $b \xleftarrow{\$} \{0,1\}$, computes $y \leftarrow \mathrm{Enc}_k(x_b)$, and finally sends ciphertext y to $\mathcal{A}$. $\mathcal{A}$'s goal is to guess b.

Classical Indistinguishability (IND)

For any efficient adversary $\mathcal{A}$ and any message x_0, x_1:

$$\left|\Pr[\mathcal{A}(y) = b] - \frac{1}{2}\right| \leq \mathrm{negl}\,(n)\,.$$

Theorem

IND $\iff$ SEM.

Quantum CPA (qCPA)

qCPA phase: $\mathcal{A}$ and $\mathcal{C}$ share a quantum channel:

Quantum CPA (qCPA)

qCPA phase: $\mathcal{A}$ and $\mathcal{C}$ share a quantum channel:

- $\mathcal{A}$ sends query: $\sum_x \alpha_{x,i} |x, 0\rangle$
- $\mathcal{C}$ replies with: $\sum_x \alpha_{x,i} |x, \mathsf{Enc}_k(x)\rangle$
- repeat for $i = 1, \ldots, q \leq \mathrm{poly}(n)$ times.

Quantum CPA (qCPA)

qCPA phase: $\mathcal{A}$ and $\mathcal{C}$ share a quantum channel:

- $\mathcal{A}$ sends query: $\sum_x \alpha_{x,i} |x, 0\rangle$
- $\mathcal{C}$ replies with: $\sum_x \alpha_{x,i} |x, \mathsf{Enc}_k(x)\rangle$
- repeat for $i = 1, \ldots, q \leq \text{poly}(n)$ times.

IND-qCPA

An encryption scheme is IND-qCPA secure if it is secure according to the (classical) IND notion, augmented by a qCPA learning phase.

Quantum CPA (qCPA)

qCPA phase: $\mathcal{A}$ and $\mathcal{C}$ share a quantum channel:

- $\mathcal{A}$ sends query: $\sum_x \alpha_{x,i} |x, 0\rangle$
- $\mathcal{C}$ replies with: $\sum_x \alpha_{x,i} |x, \mathsf{Enc}_k(x)\rangle$
- repeat for $i = 1, \ldots, q \leq \mathrm{poly}(n)$ times.

IND-qCPA

An encryption scheme is IND-qCPA secure if it is secure according to the (classical) IND notion, augmented by a qCPA learning phase.

Theorem [BZ13]

IND-qCPA is achievable and stronger than classical IND-CPA.

Quantum CPA (qCPA)

qCPA phase: $\mathcal{A}$ and $\mathcal{C}$ share a quantum channel:

- $\mathcal{A}$ sends query: $\sum_x \alpha_{x,i} |x, 0\rangle$
- $\mathcal{C}$ replies with: $\sum_x \alpha_{x,i} |x, \mathsf{Enc}_k(x)\rangle$
- repeat for $i = 1, \dots, q \leq \mathrm{poly}(n)$ times.

IND-qCPA

An encryption scheme is IND-qCPA secure if it is secure according to the (classical) IND notion, augmented by a qCPA learning phase.

Theorem [BZ13]

IND-qCPA is achievable and stronger than classical IND-CPA.

This makes sense for the public-key scenario, but in general it is clearly a 'compromise'... Why no better choice?

Fully Quantum Indistinguishability (fqIND)

fqIND phase: $\mathcal{A}$ and $\mathcal{C}$ share three quantum registers:

Fully Quantum Indistinguishability (fqIND)

fqIND phase: $\mathcal{A}$ and $\mathcal{C}$ share three quantum registers:

- $\mathcal{A}$ prepares state:

$$\sum_{x_0,x_1} \alpha_{x_0,x_1} |x_0, x_1, 0\rangle$$

Fully Quantum Indistinguishability (fqIND)

fqIND phase: $\mathcal{A}$ and $\mathcal{C}$ share three quantum registers:

- $\mathcal{A}$ prepares state:

$$\sum_{x_0,x_1} \alpha_{x_0,x_1} |x_0, x_1, 0\rangle$$

- $\mathcal{C}$ flips $b \xleftarrow{\$} \{0,1\}$ and transforms the last register to:

$$\sum_{x_0,x_1} \alpha_{x_0,x_1} |x_0, x_1, \mathsf{Enc}_k(x_b)\rangle$$

Fully Quantum Indistinguishability (fqIND)

fqIND phase: $\mathcal{A}$ and $\mathcal{C}$ share three quantum registers:

- $\mathcal{A}$ prepares state:

$$\sum_{x_0,x_1} \alpha_{x_0,x_1} |x_0, x_1, 0\rangle$$

- $\mathcal{C}$ flips $b \xleftarrow{\$} \{0,1\}$ and transforms the last register to:

$$\sum_{x_0,x_1} \alpha_{x_0,x_1} |x_0, x_1, \mathsf{Enc}_k(x_b)\rangle$$

- $\mathcal{A}$ must guess b.

Fully Quantum Indistinguishability (fqIND)

fqIND phase: $\mathcal{A}$ and $\mathcal{C}$ share three quantum registers:

- $\mathcal{A}$ prepares state:

$$\sum_{x_0,x_1} \alpha_{x_0,x_1} |x_0, x_1, 0\rangle$$

- $\mathcal{C}$ flips $b \xleftarrow{\$} \{0,1\}$ and transforms the last register to:

$$\sum_{x_0,x_1} \alpha_{x_0,x_1} |x_0, x_1, \mathsf{Enc}_k(x_b)\rangle$$

- $\mathcal{A}$ must guess b.

Theorem [BZ13]

fqIND is unachievable (too strong).

(attack exploits entanglement between ciphertext and plaintext)

BZ13 Attack (against fqIND schemes)

(example for 1-bit messages, with normalization amplitudes omitted)

$\mathcal{A}$ initializes register to: $H|0\rangle \otimes |0\rangle \otimes |0\rangle = \sum_x |x, 0, 0\rangle$
and then calls the encryption oracle with unknown bit b. Now:

- if $b = 0$, the state becomes: $\sum_x |x, 0, \mathrm{Enc}(x)\rangle$
  (notice the entanglement between 1st and 3rd register);
- if $b = 1$ instead, the state becomes:
  $\sum_x |x, 0, \mathrm{Enc}(0)\rangle = H|0\rangle \otimes |0\rangle \otimes |\mathrm{Enc}(0)\rangle$.

BZ13 Attack (against fqIND schemes)

(example for 1-bit messages, with normalization amplitudes omitted)

$\mathcal{A}$ initializes register to: $H|0\rangle \otimes |0\rangle \otimes |0\rangle = \sum_x |x, 0, 0\rangle$
and then calls the encryption oracle with unknown bit b. Now:

- if $b = 0$, the state becomes: $\sum_x |x, 0, \mathrm{Enc}(x)\rangle$
  (notice the entanglement between 1st and 3rd register);
- if $b = 1$ instead, the state becomes:
  $\sum_x |x, 0, \mathrm{Enc}(0)\rangle = H|0\rangle \otimes |0\rangle \otimes |\mathrm{Enc}(0)\rangle$.

Then $\mathcal{A}$ applies a Hadamard on the 1^{st} register and measures:

- if $b = 0$, the first register is completely mixed (irrespective of the Hadamard), and the measurement outcome is random;
- if $b = 1$ instead, the first register is: $H^2|0\rangle = |0\rangle$, and the outcome is 0.

The Road to qIND

(only focus on qIND- phase, but also assume a -qCPA phase)

The Road to qIND

(only focus on qIND- phase, but also assume a -qCPA phase)

For fqIND-qCPA many assumptions were implicitly made.

The Road to qIND

(only focus on qIND- phase, but also assume a -qCPA phase)

For fqIND-qCPA many assumptions were implicitly made. In our work, we explore every option: 'security tree' of definitions:

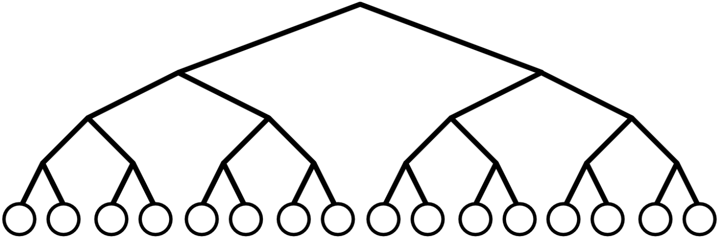

The Road to qIND

(only focus on qIND- phase, but also assume a -qCPA phase)

For fqIND-qCPA many assumptions were implicitly made. In our work, we explore every option: 'security tree' of definitions:

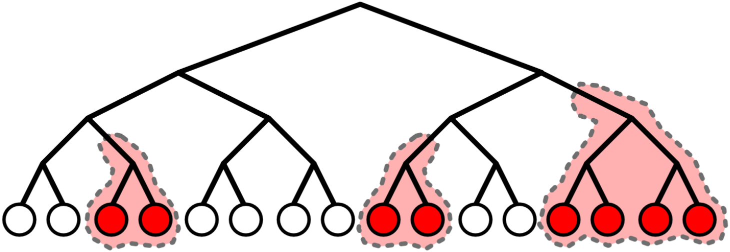

The Road to qIND

(only focus on qIND- phase, but also assume a -qCPA phase)

For fqIND-qCPA many assumptions were implicitly made. In our work, we explore every option: 'security tree' of definitions:

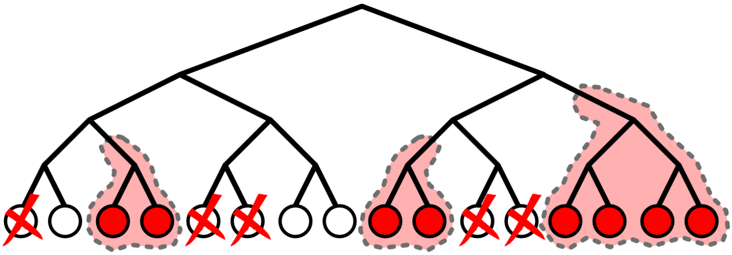

The Road to qIND

(only focus on qIND- phase, but also assume a -qCPA phase)

For fqIND-qCPA many assumptions were implicitly made. In our work, we explore every option: 'security tree' of definitions:

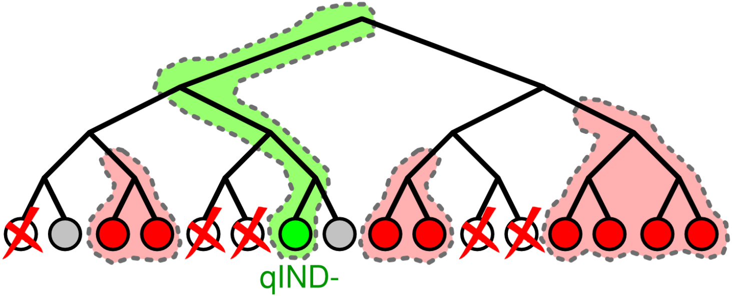

Model: $(\mathcal{O})$ vs. $(\mathcal{C})$

$$(\mathcal{O})$$

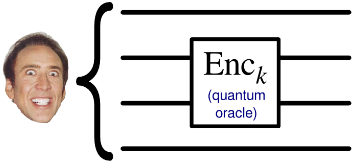

Model: $(\mathcal{O})$ vs. $(\mathcal{C})$

$(\mathcal{O})$

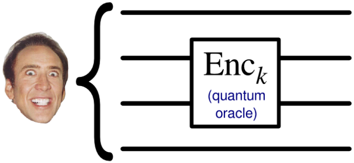

$(\mathcal{C})$

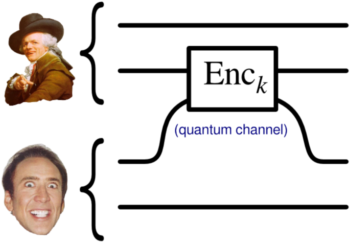

Model: (Q) vs. (c)

$$(Q)$$

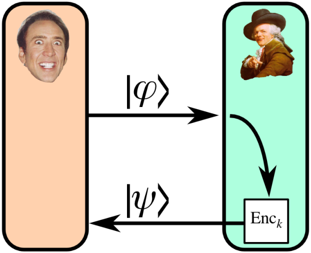

(Q)

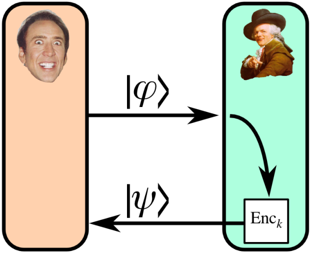
$|\varphi\rangle$
$|\psi\rangle$
Enc_k

(c)

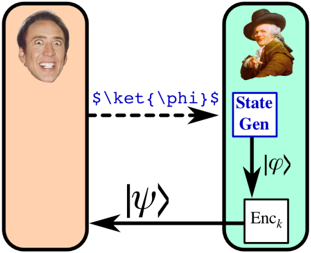
$\ket{\phi}$
State Gen
$|\varphi\rangle$
$|\psi\rangle$
Enc_k

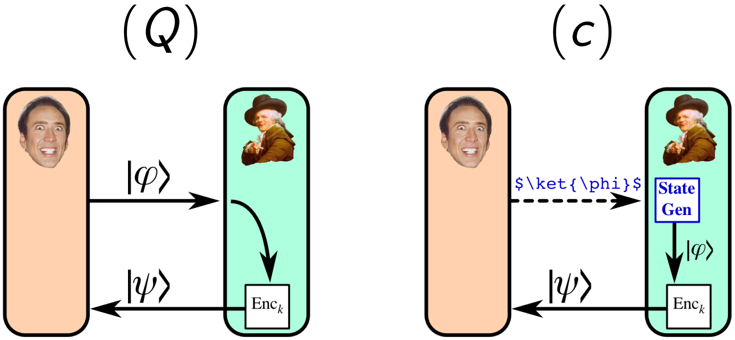

Classical description of a quantum state ρ: a classical bitstring describing the quantum circuit outputting ρ from $|0\ldots0\rangle$.

Model: (Q) vs. (c)

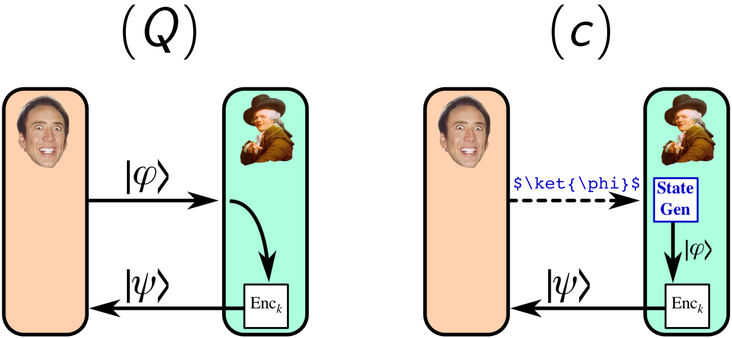

Classical description of a quantum state ρ: a classical bitstring describing the quantum circuit outputting ρ from $|0\dots0\rangle$.

Notice: if we restrict to BQP adversaries, the (c) model only differs from (Q) in the sense that the adversary is not allowed to entangle himself with the plaintext states.

Model: Type-(1) vs. Type-(2) Transformations

Type-(1)

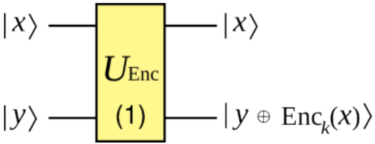

1

Model: Type-(1) vs. Type-(2) Transformations

Type-(1)

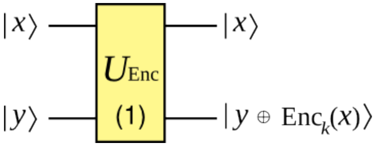

Type-(2)

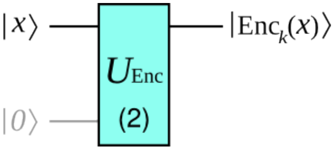

1

Model: Type-(1) vs. Type-(2) Transformations

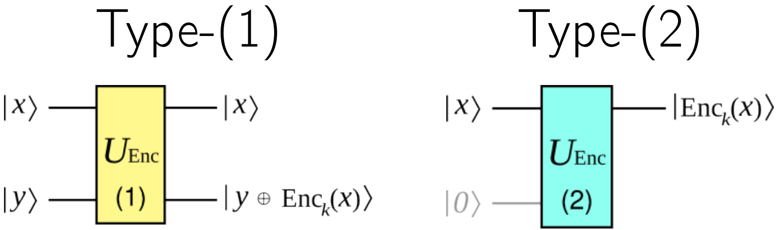

Type-(2) oracles are also called *minimal* oracles[1].

[1] E. Kashefi et al., 'A Comparison of Quantum Oracles', Phys. Rev. A 65

Model: Type-(1) vs. Type-(2) Transformations

Type-(1)

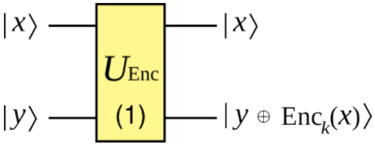

Type-(2)

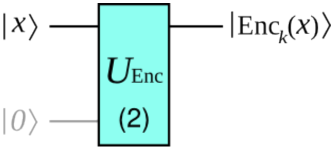

Type-(2) oracles are also called *minimal* oracles[1].

Notice: in our specific case, and limited to the qIND phase, the two types are both meaningful.

[1] E. Kashefi et al., 'A Comparison of Quantum Oracles', Phys. Rev. A 65

Quantum Indistinguishability (qIND)

qIND challenge query: $\mathcal{A}$ and $\mathcal{C}$ are two QPPT machines sharing a classical channel and a quantum channel.

Quantum Indistinguishability (qIND)

qIND challenge query: $\mathcal{A}$ and $\mathcal{C}$ are two QPPT machines sharing a classical channel and a quantum channel.

$\mathcal{A}$ sends $\mathcal{C}$ two classical, poly-sized descriptions of plaintext states ρ_0, ρ_1.

Quantum Indistinguishability (qIND)

qIND challenge query: $\mathcal{A}$ and $\mathcal{C}$ are two QPPT machines sharing a classical channel and a quantum channel.

$\mathcal{A}$ sends $\mathcal{C}$ two classical, poly-sized descriptions of plaintext states ρ_0, ρ_1.

$\mathcal{C}$ flips a random bit $b \xleftarrow{\$} \{0, 1\}$, creates ρ_b and computes:

$$\psi = U_{\mathsf{Enc}} \rho_b U_{\mathsf{Enc}}^{\dagger}$$

Quantum Indistinguishability (qIND)

qIND challenge query: $\mathcal{A}$ and $\mathcal{C}$ are two QPPT machines sharing a classical channel and a quantum channel.

$\mathcal{A}$ sends $\mathcal{C}$ two classical, poly-sized descriptions of plaintext states ρ_0, ρ_1.

$\mathcal{C}$ flips a random bit $b \xleftarrow{\$} \{0,1\}$, creates ρ_b and computes:

$$\psi = U_{\mathsf{Enc}} \rho_b U_{\mathsf{Enc}}^{\dagger}$$

and finally sends ciphertext state ψ to $\mathcal{A}$.

Quantum Indistinguishability (qIND)

qIND challenge query: $\mathcal{A}$ and $\mathcal{C}$ are two QPPT machines sharing a classical channel and a quantum channel.

$\mathcal{A}$ sends $\mathcal{C}$ two classical, poly-sized descriptions of plaintext states ρ_0, ρ_1.

$\mathcal{C}$ flips a random bit $b \xleftarrow{\$} \{0,1\}$, creates ρ_b and computes:

$$\psi = U_{\mathsf{Enc}} \rho_b U_{\mathsf{Enc}}^\dagger$$

and finally sends ciphertext state ψ to $\mathcal{A}$.

$\mathcal{A}$'s goal is to guess b.

qIND and qSEM

qIND challenge query: as the classical IND, but:

- $\mathcal{A}$ and $\mathcal{C}$ are two QPPT machines sharing a quantum channel;
- $\mathcal{A}$ can only choose classical descriptions of states;
- $\mathcal{C}$ performs type-(2) operations;
- the adversary has to distinguish the encryptions.

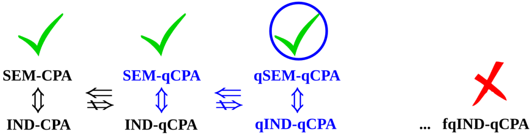

qIND and qSEM

qIND challenge query: as the classical IND, but:

- $\mathcal{A}$ and $\mathcal{C}$ are two QPPT machines sharing a quantum channel;
- $\mathcal{A}$ can only choose classical descriptions of states;
- $\mathcal{C}$ performs type-(2) operations;
- the adversary has to distinguish the encryptions.

qSEM challenge query: similar to classical SEM, but:

- template consisting of (descriptions of) quantum circuits;
- two copies of the plaintext are used to generate ciphertext and advice state (relies on classical descriptions);
- the goal is to produce a state *computationally indistinguishable* from the target state.

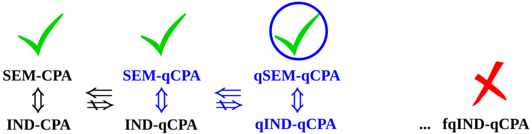

Separation Example

Theorem

IND-qCPA $\not\Rightarrow$ qIND-qCPA.

Separation Example

Theorem

IND-qCPA $\not\Rightarrow$ qIND-qCPA.

Consider [Gol04][2] : sample $r \xleftarrow{\$} \mathcal{R}$ and use a PRF $f : \mathcal{K} \times \mathcal{R} \to \mathcal{M}$. Then: $\mathsf{Enc}_k(x) := (x \oplus f_k(r), r)$.

[2] O. Goldreich: *'Foundations of Cryptography: Volume 2'*

Separation Example

Theorem

IND-qCPA $\not\Rightarrow$ qIND-qCPA.

Consider [Gol04][2] : sample $r \xleftarrow{\$} \mathcal{R}$ and use a PRF $f : \mathcal{K} \times \mathcal{R} \rightarrow \mathcal{M}$. Then: $\mathrm{Enc}_k(x) := (x \oplus f_k(r), r)$.

Theorem [BZ13]

The Goldreich scheme is IND-qCPA secure, provided the PRF is quantum-secure.

[2] O. Goldreich: *'Foundations of Cryptography: Volume 2'*

Separation Example

Theorem

IND-qCPA $\not\Rightarrow$ qIND-qCPA.

Consider [Gol04][2] : sample $r \xleftarrow{\$} \mathcal{R}$ and use a PRF $f : \mathcal{K} \times \mathcal{R} \to \mathcal{M}$. Then: $\mathsf{Enc}_k(x) := (x \oplus f_k(r), r)$.

Theorem [BZ13]

The Goldreich scheme is IND-qCPA secure, provided the PRF is quantum-secure.

Theorem

The Goldreich scheme is *not* qIND-qCPA secure.

[2] O. Goldreich: *'Foundations of Cryptography: Volume 2'*

Impossibility Result

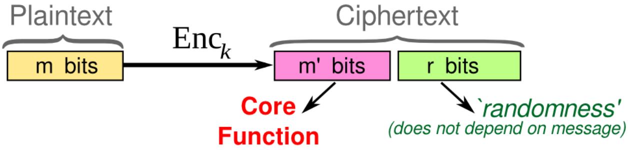

Impossibility Result

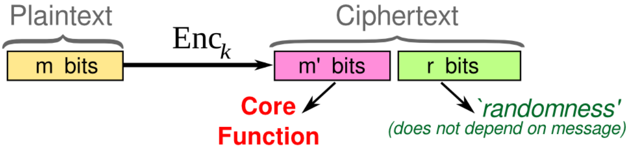

quasi–length-preserving (QLP): core function is bijective ($m = m'$).

Impossibility Result

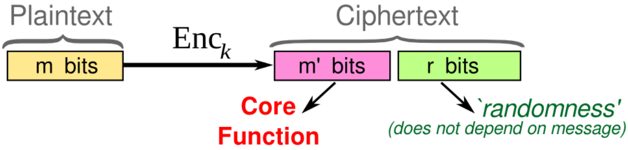

quasi–length-preserving (QLP): core function is bijective ($m = m'$).

- Goldreich's scheme
- OTP
- ECB block ciphers
- stream ciphers

Impossibility Result

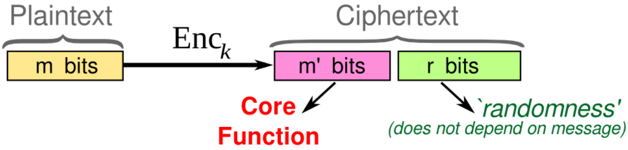

quasi–length-preserving (QLP): core function is bijective ($m = m'$).

- Goldreich's scheme
- OTP
- ECB block ciphers
- stream ciphers

Theorem

If a symmetric scheme is QLP, then it is *not* qIND-qCPA secure.

Impossibility Result

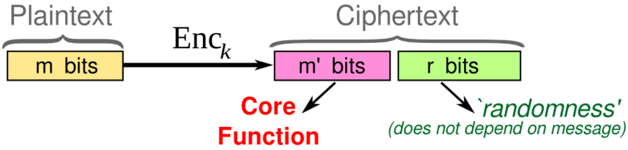

quasi–length-preserving (QLP): core function is bijective ($m = m'$).

- Goldreich's scheme
- OTP
- ECB block ciphers
- stream ciphers

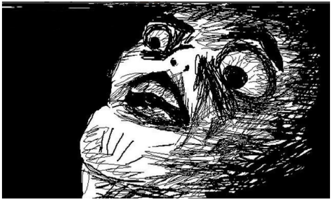

Theorem

If a symmetric scheme is QLP, then it is *not* qIND-qCPA secure.

The Attack

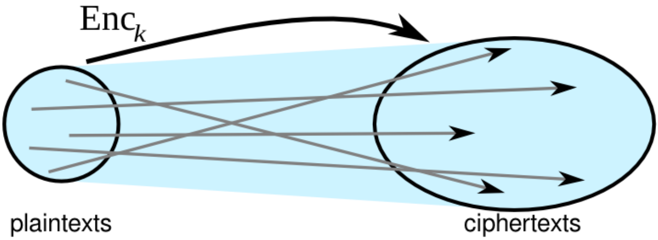

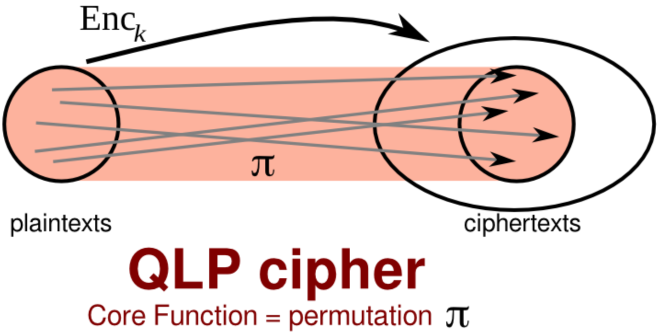
Enc$_k$
π
plaintexts
ciphertexts
QLP cipher
Core Function = permutation π

The Attack

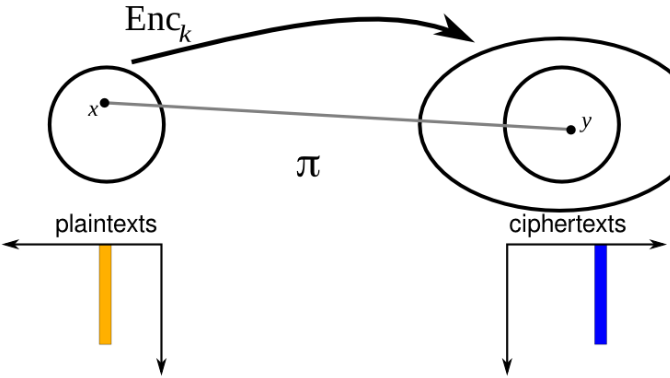

The Attack

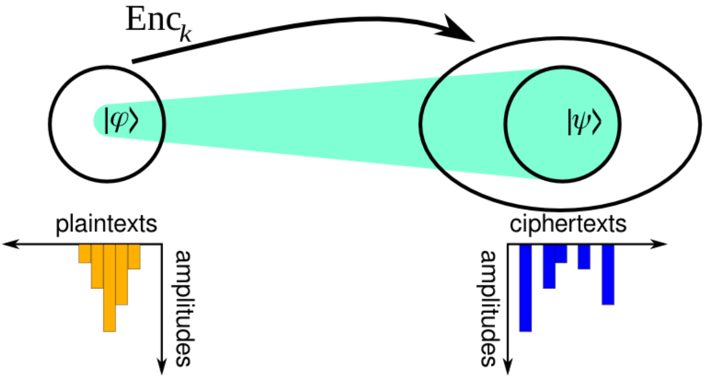

Enc_k
$|\varphi\rangle$
$|\psi\rangle$
plaintexts
amplitudes
ciphertexts
amplitudes

The Attack

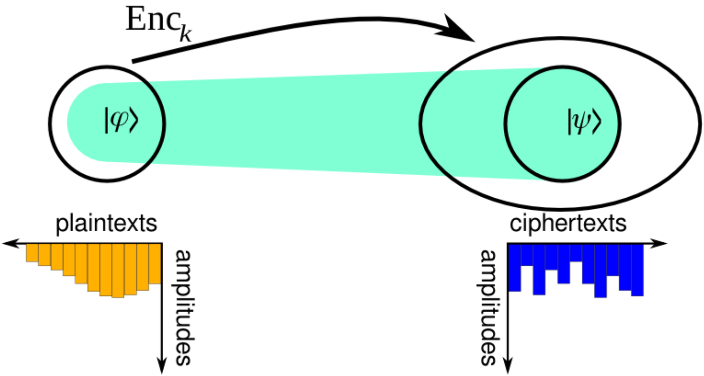

Enc_k
$|\varphi\rangle$
$|\psi\rangle$
plaintexts
amplitudes
ciphertexts
amplitudes

The Attack

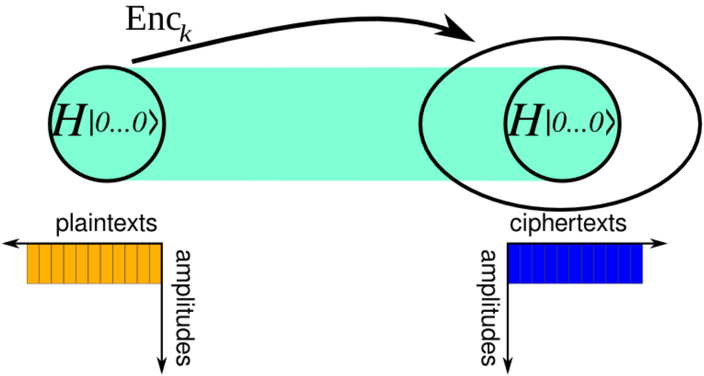

Enc_k
$H|0...0\rangle$
$H|0...0\rangle$
plaintexts
amplitudes
ciphertexts
amplitudes

The Attack

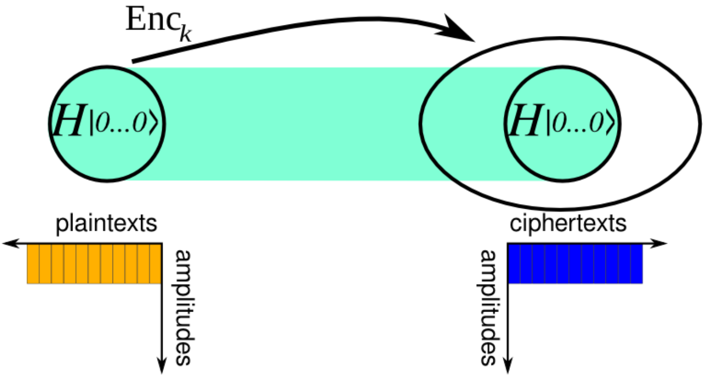

$$|+\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle \xrightarrow{\mathsf{Enc}_k} \frac{1}{\sqrt{2}}|\pi(0)\rangle + \frac{1}{\sqrt{2}}|\pi(1)\rangle = |+\rangle$$

Enc$_k(|+\rangle)$ is easy to distinguish from Enc$_k(|0\rangle)$,
e.g. by applying a Hadamard and measuring.

The Solution

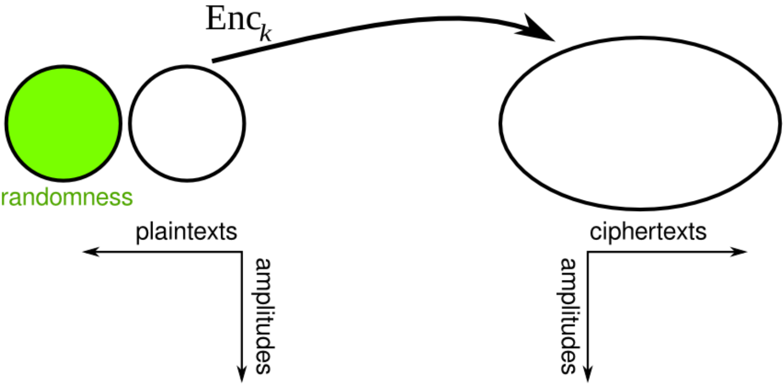

The Solution

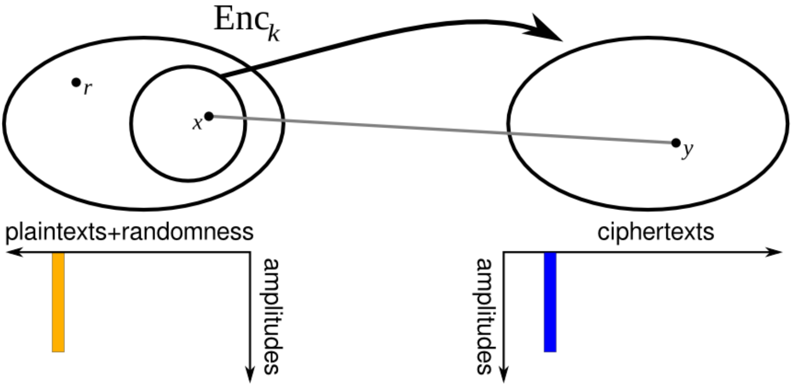

The Solution

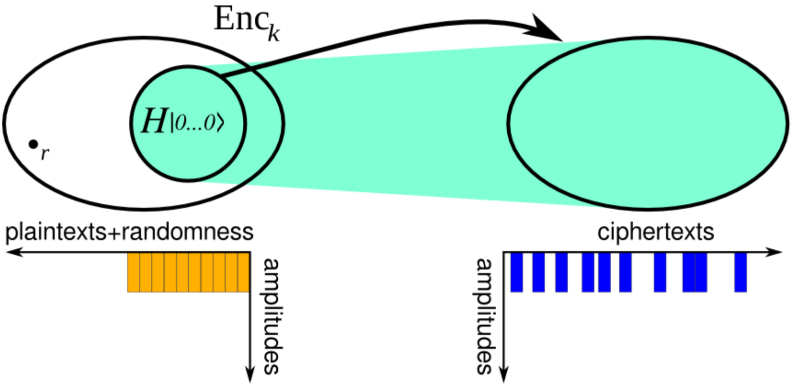

The Solution

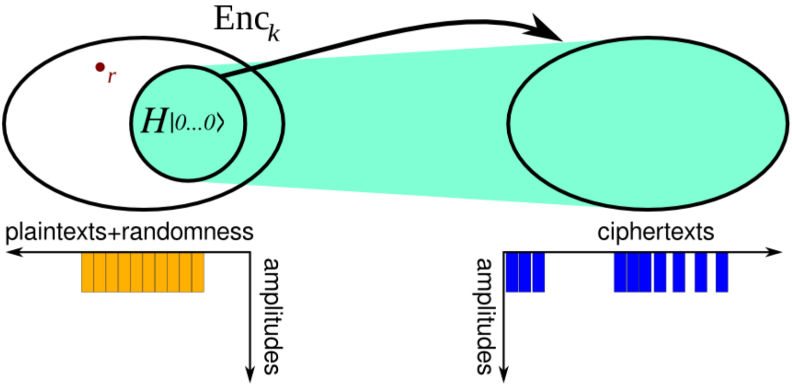

The Solution

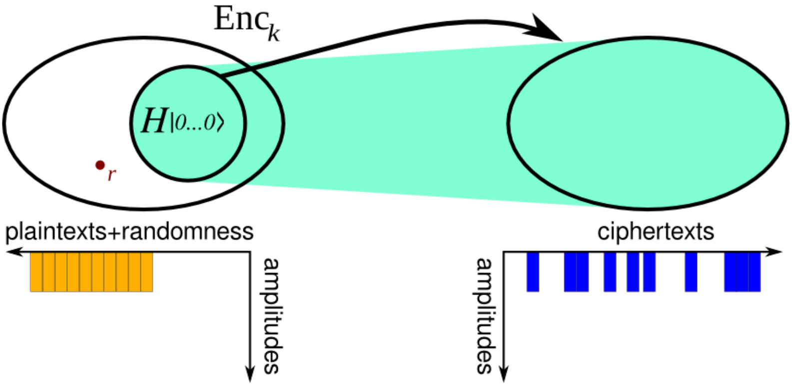

Secure Construction

Π family of quantum-secure pseudorandom permutations (QPRP).

Secure Construction

Π family of quantum-secure pseudorandom permutations (QPRP).

Construction

- Generate key: sample $(\pi, \pi^{-1}) \leftarrow \Pi$;
- Encrypt message x: pad with n bits of randomness r and set $y = \pi(r \| x)$;
- Decrypt y: truncate the first n bits of $\pi^{-1}(y)$.

Secure Construction

Π family of quantum-secure pseudorandom permutations (QPRP).

Construction

- Generate key: sample $(\pi, \pi^{-1}) \leftarrow \Pi$;
- Encrypt message x: pad with n bits of randomness r and set $y = \pi(r \| x)$;
- Decrypt y: truncate the first n bits of $\pi^{-1}(y)$.

Theorem

The above scheme is qIND-qCPA secure.

Secure Construction

Π family of quantum-secure pseudorandom permutations (QPRP).

Construction

- Generate key: sample $(\pi, \pi^{-1}) \leftarrow \Pi$;
- Encrypt message x: pad with n bits of randomness r and set $y = \pi(r \| x)$;
- Decrypt y: truncate the first n bits of $\pi^{-1}(y)$.

Theorem

The above scheme is qIND-qCPA secure.

(Idea of proof: show that for every two plaintext states $|\phi_0\rangle, |\phi_1\rangle$, the trace distance of the states ρ_0, ρ_1 obtained by considering their encryption under a mixture of every possible key is negligible)

Conclusions

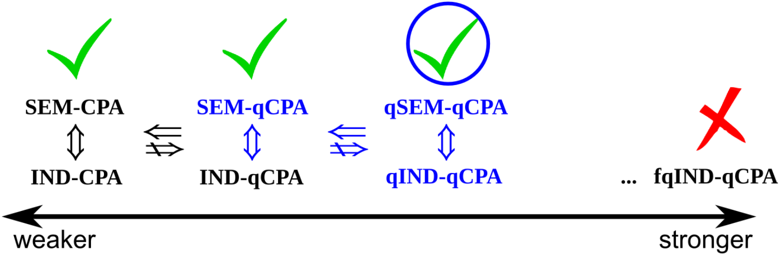

Conclusions

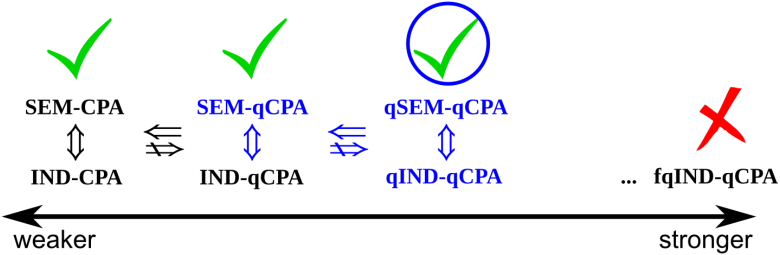

Future directions:

- public-key encryption;
- CCA security;
- qIND-qCPA security for longer messages, block-cipher mode of operations;
- 'fully quantum' IND and relation to our (Q2) notion;
- security of our construction also in the (Q2) model;
- patch IND-qCPA $\Rightarrow$ qIND-qCPA (using a HMAC).

End of this talk

Thanks for your attention!

c.schaffner@uva.nl

http://arxiv.org/abs/1504.05255

Nederlandse Organisatie voor Wetenschappelijk Onderzoek

Classical Semantic Security (SEM)

Simulation-based security: $\mathcal{A}$ is simulated by $\mathcal{S}$ in two different 'worlds' (real VS ideal).

Classical Semantic Security (SEM)

Simulation-based security: $\mathcal{A}$ is simulated by $\mathcal{S}$ in two different 'worlds' (real VS ideal).

SEM challenge query: $\mathcal{A}$ chooses a challenge template:

Classical Semantic Security (SEM)

Simulation-based security: $\mathcal{A}$ is simulated by $\mathcal{S}$ in two different 'worlds' (real VS ideal).

SEM challenge query: $\mathcal{A}$ chooses a challenge template:

- a message distribution X on plaintext space $\mathcal{M}$,

Classical Semantic Security (SEM)

Simulation-based security: $\mathcal{A}$ is simulated by $\mathcal{S}$ in two different 'worlds' (real VS ideal).

SEM challenge query: $\mathcal{A}$ chooses a challenge template:

- a message distribution X on plaintext space $\mathcal{M}$,
- an advice function $h : \mathcal{M} \to \mathbb{N}$,

Classical Semantic Security (SEM)

Simulation-based security: $\mathcal{A}$ is simulated by $\mathcal{S}$ in two different 'worlds' (real VS ideal).

SEM challenge query: $\mathcal{A}$ chooses a challenge template:

- a message distribution X on plaintext space $\mathcal{M}$,
- an advice function $h : \mathcal{M} \to \mathbb{N}$,
- a target function $f : \mathcal{M} \to \mathbb{N}$.

Classical Semantic Security (SEM)

Simulation-based security: $\mathcal{A}$ is simulated by $\mathcal{S}$ in two different 'worlds' (real VS ideal).

SEM challenge query: $\mathcal{A}$ chooses a challenge template:

- a message distribution X on plaintext space $\mathcal{M}$,
- an advice function $h : \mathcal{M} \to \mathbb{N}$,
- a target function $f : \mathcal{M} \to \mathbb{N}$.

x is sampled from X and $\mathcal{A}$ receives $(\mathsf{Enc}_k(x), h(x))$,

Classical Semantic Security (SEM)

Simulation-based security: $\mathcal{A}$ is simulated by $\mathcal{S}$ in two different 'worlds' (real VS ideal).

SEM challenge query: $\mathcal{A}$ chooses a challenge template:

- a message distribution X on plaintext space $\mathcal{M}$,
- an advice function $h : \mathcal{M} \to \mathbb{N}$,
- a target function $f : \mathcal{M} \to \mathbb{N}$.

x is sampled from X and $\mathcal{A}$ receives $(\mathsf{Enc}_k(x), h(x))$, but $\mathcal{S}$ only receives $h(x)$.

Classical Semantic Security (SEM)

Simulation-based security: $\mathcal{A}$ is simulated by $\mathcal{S}$ in two different 'worlds' (real VS ideal).

SEM challenge query: $\mathcal{A}$ chooses a challenge template:

- a message distribution X on plaintext space $\mathcal{M}$,
- an advice function $h : \mathcal{M} \to \mathbb{N}$,
- a target function $f : \mathcal{M} \to \mathbb{N}$.

x is sampled from X and $\mathcal{A}$ receives $(\mathrm{Enc}_k(x), h(x))$, but $\mathcal{S}$ only receives $h(x)$. The goal for both is to compute $f(x)$.

Classical Semantic Security (SEM)

Simulation-based security: $\mathcal{A}$ is simulated by $\mathcal{S}$ in two different 'worlds' (real VS ideal).

SEM challenge query: $\mathcal{A}$ chooses a challenge template:

- a message distribution X on plaintext space $\mathcal{M}$,
- an advice function $h : \mathcal{M} \to \mathbb{N}$,
- a target function $f : \mathcal{M} \to \mathbb{N}$.

x is sampled from X and $\mathcal{A}$ receives $(\mathsf{Enc}_k(x), h(x))$, but $\mathcal{S}$ only receives $h(x)$. The goal for both is to compute $f(x)$.

Classical Semantic Security (SEM)

For any efficient adversary $\mathcal{A}$ there exists an efficient simulator $\mathcal{S}$ such that the two 'worlds' are indistinguishable.

Classical Semantic Security (SEM)

Simulation-based security: $\mathcal{A}$ is simulated by $\mathcal{S}$ in two different 'worlds' (real VS ideal).

SEM challenge query: $\mathcal{A}$ chooses a challenge template:

- a message distribution X on plaintext space $\mathcal{M}$,
- an advice function $h : \mathcal{M} \to \mathbb{N}$,
- a target function $f : \mathcal{M} \to \mathbb{N}$.

x is sampled from X and $\mathcal{A}$ receives $(\mathsf{Enc}_k(x), h(x))$, but $\mathcal{S}$ only receives $h(x)$. The goal for both is to compute $f(x)$.

Classical Semantic Security (SEM)

For any efficient adversary $\mathcal{A}$ there exists an efficient simulator $\mathcal{S}$ such that the two 'worlds' are indistinguishable.

This definition is cumbersome.

Chosen Plaintext Attack (CPA)

CPA 'learning' phase: $\mathcal{A}$ chooses $\mathcal{C}$ up to $q = \mathrm{poly}(n)$ plaintexts $x_1, \ldots, x_q \in \mathcal{M}$ (possibly adaptively) and receives ciphertexts $\mathsf{Enc}_k(x_1), \ldots, \mathsf{Enc}_k(x_q)$.

Chosen Plaintext Attack (CPA)

CPA ‘learning’ phase: $\mathcal{A}$ chooses $\mathcal{C}$ up to $q = \text{poly}(n)$ plaintexts $x_1, \ldots, x_q \in \mathcal{M}$ (possibly adaptively) and receives ciphertexts $\text{Enc}_k(x_1), \ldots, \text{Enc}_k(x_q)$.

Can be done both before and/or after another challenge query.

Chosen Plaintext Attack (CPA)

CPA 'learning' phase: $\mathcal{A}$ chooses $\mathcal{C}$ up to $q = \text{poly}(n)$ plaintexts $x_1, \ldots, x_q \in \mathcal{M}$ (possibly adaptively) and receives ciphertexts $\text{Enc}_k(x_1), \ldots, \text{Enc}_k(x_q)$.

Can be done both before and/or after another challenge query.

Can be combined with other security notions:

Chosen Plaintext Attack (CPA)

CPA 'learning' phase: $\mathcal{A}$ chooses $\mathcal{C}$ up to $q = \text{poly}(n)$ plaintexts $x_1, \ldots, x_q \in \mathcal{M}$ (possibly adaptively) and receives ciphertexts $\text{Enc}_k(x_1), \ldots, \text{Enc}_k(x_q)$.

Can be done both before and/or after another challenge query.

Can be combined with other security notions:

CPA phase + SEM phase $\Rightarrow$ SEM-CPA security.

Chosen Plaintext Attack (CPA)

CPA 'learning' phase: $\mathcal{A}$ chooses $\mathcal{C}$ up to $q = \text{poly}(n)$ plaintexts $x_1, \ldots, x_q \in \mathcal{M}$ (possibly adaptively) and receives ciphertexts $\text{Enc}_k(x_1), \ldots, \text{Enc}_k(x_q)$.

Can be done both before and/or after another challenge query.

Can be combined with other security notions:

CPA phase + SEM phase $\Rightarrow$ SEM-CPA security.

CPA phase + IND phase $\Rightarrow$ IND-CPA security.

Chosen Plaintext Attack (CPA)

CPA 'learning' phase: $\mathcal{A}$ chooses $\mathcal{C}$ up to $q = \text{poly}(n)$ plaintexts $x_1, \ldots, x_q \in \mathcal{M}$ (possibly adaptively) and receives ciphertexts $\text{Enc}_k(x_1), \ldots, \text{Enc}_k(x_q)$.

Can be done both before and/or after another challenge query.

Can be combined with other security notions:

CPA phase + SEM phase $\Rightarrow$ SEM-CPA security.

CPA phase + IND phase $\Rightarrow$ IND-CPA security.

Theorem

IND-CPA $\iff$ SEM-CPA.

Chosen Plaintext Attack (CPA)

CPA 'learning' phase: $\mathcal{A}$ chooses $\mathcal{C}$ up to $q = \text{poly}(n)$ plaintexts $x_1, \ldots, x_q \in \mathcal{M}$ (possibly adaptively) and receives ciphertexts $\text{Enc}_k(x_1), \ldots, \text{Enc}_k(x_q)$.

Can be done both before and/or after another challenge query.

Can be combined with other security notions:

CPA phase + SEM phase $\Rightarrow$ SEM-CPA security.

CPA phase + IND phase $\Rightarrow$ IND-CPA security.

Theorem

IND-CPA $\iff$ SEM-CPA.

Note: deterministic schemes are insecure $\Rightarrow$ need for randomization.

BZ Attack

(example for 1-bit messages, with normalization amplitudes omitted)

$\mathcal{A}$ initializes register to: $H|0\rangle \otimes |0\rangle \otimes |0\rangle = \sum_x |x, 0, 0\rangle$ and then calls the encryption oracle with unknown bit b. Now:

- if $b = 0$, the state becomes: $\sum_x |x, 0, \mathsf{Enc}(x)\rangle$ (notice entanglement between 1^{st} and 3^{rd} registers);
- if $b = 1$ instead, the state becomes: $\sum_x |x, 0, \mathsf{Enc}(0)\rangle = H|0\rangle \otimes |0\rangle \otimes |\mathsf{Enc}(0)\rangle$.

Then $\mathcal{A}$ applies a Hadamard on the 1^{st} register and measures:

- if $b = 0$, the Hadamard maps the state to a complete mixture, and the measurement outcome is random;
- if $b = 1$ instead, the first register is: $H^2|0\rangle = |0\rangle$, and the outcome is 0.

Quantum Indistinguishability (qIND)

qIND challenge query: $\mathcal{A}$ and $\mathcal{C}$ are two QPPT machines sharing a classical channel and a quantum channel.

Quantum Indistinguishability (qIND)

qIND challenge query: $\mathcal{A}$ and $\mathcal{C}$ are two QPPT machines sharing a classical channel and a quantum channel.

$\mathcal{A}$ sends $\mathcal{C}$ two classical, poly-sized descriptions of plaintext states ρ_0, ρ_1.

Quantum Indistinguishability (qIND)

qIND challenge query: $\mathcal{A}$ and $\mathcal{C}$ are two QPPT machines sharing a classical channel and a quantum channel.

$\mathcal{A}$ sends $\mathcal{C}$ two classical, poly-sized descriptions of plaintext states ρ_0, ρ_1.

$\mathcal{C}$ flips a random bit $b \xleftarrow{\$} \{0,1\}$, and computes:

$$\psi = U_{\mathsf{Enc}} \rho_b U_{\mathsf{Enc}}^{\dagger}$$

Quantum Indistinguishability (qIND)

qIND challenge query: $\mathcal{A}$ and $\mathcal{C}$ are two QPPT machines sharing a classical channel and a quantum channel.

$\mathcal{A}$ sends $\mathcal{C}$ two classical, poly-sized descriptions of plaintext states ρ_0, ρ_1.

$\mathcal{C}$ flips a random bit $b \xleftarrow{\$} \{0,1\}$, and computes:

$$\psi = U_{\mathsf{Enc}} \rho_b U_{\mathsf{Enc}}^\dagger$$

and finally sends ciphertext state ψ to $\mathcal{A}$.

Quantum Indistinguishability (qIND)

qIND challenge query: $\mathcal{A}$ and $\mathcal{C}$ are two QPPT machines sharing a classical channel and a quantum channel.

$\mathcal{A}$ sends $\mathcal{C}$ two classical, poly-sized descriptions of plaintext states ρ_0, ρ_1.

$\mathcal{C}$ flips a random bit $b \xleftarrow{\$} \{0, 1\}$, and computes:

$$\psi = U_{\mathsf{Enc}} \rho_b U_{\mathsf{Enc}}^{\dagger}$$

and finally sends ciphertext state ψ to $\mathcal{A}$.

$\mathcal{A}$'s goal is to guess b.

Quantum Indistinguishability (qIND)

Quantum Indistinguishability (qIND)

For any QPPT adversary $\mathcal{A}$ and any ρ_0, ρ_1 with efficient classical representations:

$$\left| \Pr[\mathcal{A}(\psi) = b] - \frac{1}{2} \right| \leq \mathsf{negl}\left(n\right),$$

where $\psi = U_{\mathsf{Enc}} \rho_b U_{\mathsf{Enc}}^\dagger$, and $b \xleftarrow{\$} \{0, 1\}$.

Quantum Indistinguishability (qIND)

Quantum Indistinguishability (qIND)

For any QPPT adversary $\mathcal{A}$ and any ρ_0, ρ_1 with efficient classical representations:

$$\left|\Pr[\mathcal{A}(\psi) = b] - \frac{1}{2}\right| \leq \mathsf{negl}(n),$$

where $\psi = U_{\mathsf{Enc}}\rho_b U_{\mathsf{Enc}}^{\dagger}$, and $b \xleftarrow{\$} \{0,1\}$.

Quantum Indistinguishability under qCPA (qIND-qCPA)

An encryption scheme is IND-qCPA secure if it is secure according to the qIND notion, augmented by a qCPA learning phase.

Quantum Indistinguishability (qIND)

Quantum Indistinguishability (qIND)

For any QPPT adversary $\mathcal{A}$ and any ρ_0, ρ_1 with efficient classical representations:

$$\left|\Pr[\mathcal{A}(\psi) = b] - \frac{1}{2}\right| \leq \mathsf{negl}\left(n\right),$$

where $\psi = U_{\mathsf{Enc}} \rho_b U_{\mathsf{Enc}}^{\dagger}$, and $b \xleftarrow{\$} \{0, 1\}$.

Quantum Indistinguishability under qCPA (qIND-qCPA)

An encryption scheme is IND-qCPA secure if it is secure according to the qIND notion, augmented by a qCPA learning phase.

what about quantum semantic security?

Quantum Semantic Security

Classical Semantic Security under qCPA (SEM-qCPA)

An encryption scheme is SEM-qCPA secure if it is secure according to the SEM notion, augmented by a qCPA learning phase.

Quantum Semantic Security

Classical Semantic Security under qCPA (SEM-qCPA)

An encryption scheme is SEM-qCPA secure if it is secure according to the SEM notion, augmented by a qCPA learning phase.

Theorem

IND-qCPA $\iff$ SEM-qCPA.

Quantum Semantic Security

Classical Semantic Security under qCPA (SEM-qCPA)

An encryption scheme is SEM-qCPA secure if it is secure according to the SEM notion, augmented by a qCPA learning phase.

Theorem

IND-qCPA $\iff$ SEM-qCPA.

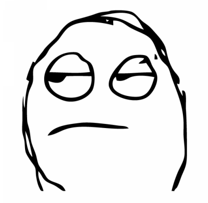

BOOOOORING...

Quantum Semantic Security

Classical Semantic Security under qCPA (SEM-qCPA)

An encryption scheme is SEM-qCPA secure if it is secure according to the SEM notion, augmented by a qCPA learning phase.

Theorem

IND-qCPA $\iff$ SEM-qCPA.

Proof Idea:

'$\Rightarrow$': provide $\mathcal{S}$ with $\mathcal{A}$'s code through h, impersonate $\mathcal{C}$ and use IND to argue same prob.

'$\Leftarrow$': assume distinguisher $\mathcal{A}$, choose constant h, then no $\mathcal{S}$ can infere anything w/o ciphertext.

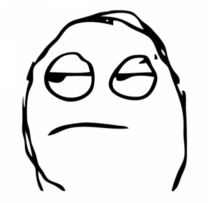

BOOOOORING...

Quantum Semantic Security

qSEM challenge query: $\mathcal{A}$ chooses a challenge template consisting of classical descriptions of:

Quantum Semantic Security

qSEM challenge query: $\mathcal{A}$ chooses a challenge template consisting of classical descriptions of:

- a quantum generator circuit $G : \mathbb{N} \to \mathcal{H}_M$,

Quantum Semantic Security

qSEM challenge query: $\mathcal{A}$ chooses a challenge template consisting of classical descriptions of:

- a quantum generator circuit $G : \mathbb{N} \to \mathcal{H}_M$,
- a quantum advice circuit $h : \mathcal{H}_\mathcal{M} \to \mathcal{H}_h$,

Quantum Semantic Security

qSEM challenge query: $\mathcal{A}$ chooses a challenge template consisting of classical descriptions of:

- a quantum generator circuit $G : \mathbb{N} \to \mathcal{H}_M$,
- a quantum advice circuit $h : \mathcal{H}_\mathcal{M} \to \mathcal{H}_h$,
- a quantum target circuit $f : \mathcal{H}_\mathcal{M} \to \mathcal{H}_f$.

Quantum Semantic Security

qSEM challenge query: $\mathcal{A}$ chooses a challenge template consisting of classical descriptions of:

- a quantum generator circuit $G : \mathbb{N} \to \mathcal{H}_M$,
- a quantum advice circuit $h : \mathcal{H}_\mathcal{M} \to \mathcal{H}_h$,
- a quantum target circuit $f : \mathcal{H}_\mathcal{M} \to \mathcal{H}_f$.

G is run twice on the same randomness, producing two copies of ρ (consider purification here);

Quantum Semantic Security

qSEM challenge query: $\mathcal{A}$ chooses a challenge template consisting of classical descriptions of:

- a quantum generator circuit $G : \mathbb{N} \to \mathcal{H}_M$,
- a quantum advice circuit $h : \mathcal{H}_\mathcal{M} \to \mathcal{H}_h$,
- a quantum target circuit $f : \mathcal{H}_\mathcal{M} \to \mathcal{H}_f$.

G is run twice on the same randomness, producing two copies of ρ (consider purification here);

- the first copy gets encrypted to $\psi = U_{\mathsf{Enc}} \rho U_{\mathsf{Enc}}^\dagger$,

Quantum Semantic Security

qSEM challenge query: $\mathcal{A}$ chooses a challenge template consisting of classical descriptions of:

- a quantum generator circuit $G : \mathbb{N} \to \mathcal{H}_M$,
- a quantum advice circuit $h : \mathcal{H}_\mathcal{M} \to \mathcal{H}_h$,
- a quantum target circuit $f : \mathcal{H}_\mathcal{M} \to \mathcal{H}_f$.

G is run twice on the same randomness, producing two copies of ρ (consider purification here);

- the first copy gets encrypted to $\psi = U_{\mathsf{Enc}} \rho U_{\mathsf{Enc}}^\dagger$,
- the second copy is used to compute $h(\rho)$.

Quantum Semantic Security

qSEM challenge query: $\mathcal{A}$ chooses a challenge template consisting of classical descriptions of:

- a quantum generator circuit $G : \mathbb{N} \to \mathcal{H}_M$,
- a quantum advice circuit $h : \mathcal{H}_\mathcal{M} \to \mathcal{H}_h$,
- a quantum target circuit $f : \mathcal{H}_\mathcal{M} \to \mathcal{H}_f$.

G is run twice on the same randomness, producing two copies of ρ (consider purification here);

- the first copy gets encrypted to $\psi = U_{\mathsf{Enc}} \rho U_{\mathsf{Enc}}^{\dagger}$,
- the second copy is used to compute $h(\rho)$.

$\mathcal{A}$ receives $(\psi, h(\rho))$

Quantum Semantic Security

qSEM challenge query: $\mathcal{A}$ chooses a challenge template consisting of classical descriptions of:

- a quantum generator circuit $G : \mathbb{N} \to \mathcal{H}_M$,
- a quantum advice circuit $h : \mathcal{H}_\mathcal{M} \to \mathcal{H}_h$,
- a quantum target circuit $f : \mathcal{H}_\mathcal{M} \to \mathcal{H}_f$.

G is run twice on the same randomness, producing two copies of ρ (consider purification here);

- the first copy gets encrypted to $\psi = U_{\mathsf{Enc}} \rho U_{\mathsf{Enc}}^\dagger$,
- the second copy is used to compute $h(\rho)$.

$\mathcal{A}$ receives $(\psi, h(\rho))$; but $\mathcal{S}$ only gets $h(\rho)$.

Quantum Semantic Security

qSEM challenge query: $\mathcal{A}$ chooses a challenge template consisting of classical descriptions of:

- a quantum generator circuit $G : \mathbb{N} \to \mathcal{H}_M$,
- a quantum advice circuit $h : \mathcal{H}_\mathcal{M} \to \mathcal{H}_h$,
- a quantum target circuit $f : \mathcal{H}_\mathcal{M} \to \mathcal{H}_f$.

G is run twice on the same randomness, producing two copies of ρ (consider purification here);

- the first copy gets encrypted to $\psi = U_{\mathsf{Enc}} \rho U_{\mathsf{Enc}}^\dagger$,
- the second copy is used to compute $h(\rho)$.

$\mathcal{A}$ receives $(\psi, h(\rho))$; but $\mathcal{S}$ only gets $h(\rho)$.

Goal is to compute a state φ computationally indistinguishable from $f(\rho)$.

Quantum Semantic Security

Quantum Semantic Security (qSEM)

For any efficient quantum adversary $\mathcal{A}$ there exists an efficient quantum simulator $\mathcal{S}$ such that their qSEM templates are identically distributed, and:

$$|\Pr[\mathcal{A}(\psi, h(\rho)) \text{ wins qSEM }] - \Pr[\mathcal{S}(h(\rho)) \text{ wins qSEM }]| \leq \mathsf{negl}(n)$$

Quantum Semantic Security

Quantum Semantic Security (qSEM)

For any efficient quantum adversary $\mathcal{A}$ there exists an efficient quantum simulator $\mathcal{S}$ such that their qSEM templates are identically distributed, and:

$$|\Pr[\mathcal{A}(\psi, h(\rho)) \text{ wins qSEM }] - \Pr[\mathcal{S}(h(\rho)) \text{ wins qSEM }]| \leq \mathsf{negl}(n)$$

Quantum Semantic Security under qCPA (qSEM-qCPA)

An encryption scheme is qSEM-qCPA secure if it is secure according to the qSEM notion, augmented by a qCPA learning phase.

Quantum Semantic Security

Quantum Semantic Security (qSEM)

For any efficient quantum adversary $\mathcal{A}$ there exists an efficient quantum simulator $\mathcal{S}$ such that their qSEM templates are identically distributed, and:

$$|\Pr[\mathcal{A}(\psi, h(\rho)) \text{ wins qSEM }] - \Pr[\mathcal{S}(h(\rho)) \text{ wins qSEM }]| \leq \mathrm{negl}\,(n)$$

Quantum Semantic Security under qCPA (qSEM-qCPA)

An encryption scheme is qSEM-qCPA secure if it is secure according to the qSEM notion, augmented by a qCPA learning phase.

Theorem

qIND-qCPA $\iff$ qSEM-qCPA.

qSEM $\Rightarrow$ qIND

By contradiction: let $\mathcal{A}$ be an efficient qIND distinguisher. We show that there exists an efficient $\mathcal{A}'$ for qSEM which does not admit simulator.

$\mathcal{A}'$ invokes $\mathcal{A}$, which starts a qIND challenge query consisting of two classical descriptions s_0, s_1 of states ρ_0, ρ_1.

$\mathcal{A}'$ records this template, then prepare his own qSEM challenge template consisting of:

- as generator G, the circuit outputting ρ_0 or ρ_1 uniformly;
- as advice h, a 'dumb' (constant output) circuit;
- as target f, the *identity* circuit $f(\rho) = \rho$.

$\mathcal{A}'$ receives $\mathcal{C}$'s response, forwards the ciphertext to $\mathcal{A}$, and observes output.

Since $\mathcal{A}$ recovers b with non-negligible probability, $\mathcal{A}'$ can then reconstruct the correct ρ_b (having recorded its description) and compute the target state $f(\rho_b)$.

Any simulator $\mathcal{S}$, on the other hand, only receives a constant state, and then cannot do better than guessing.

qSEM $\Leftarrow$ qIND

Let $\mathcal{A}$ be any QPT adversary against qSEM. Then its circuit has a short classical representation ξ.

Then here is a simulator $\mathcal{S}$ with the same success probability:

1. $\mathcal{S}$ receives ξ as nonuniform advice (this is allowed);
2. then $\mathcal{S}$ implements and run $\mathcal{A}$ through ξ;
3. when $\mathcal{A}$ produces a qSEM challenge template (G, h, f), $\mathcal{S}$ forwards it to $\mathcal{C}$;
4. when $\mathcal{C}$ replies with its advice state, $\mathcal{S}$ forwards it to $\mathcal{A}$, together with the encryption of a bogus state;
5. finally, $\mathcal{S}$ outputs whatever $\mathcal{A}$ does.

The presence of the bogus encryption state instead of the right one does not affect $\mathcal{A}$'s success probability. In fact, if this were the case, we could turn $\mathcal{S}$ into an efficient distinguisher against qIND.

The $(\mathcal{C})$ model

Objection:

The $(\mathcal{C})$ model is a problem if you need rewinding: how do you rewind the challenger?

The $(\mathcal{C})$ model

Objection:

The $(\mathcal{C})$ model is a problem if you need rewinding: how do you rewind the challenger?

Our response: rewinding the challenger would represent a scenario where the adversary has almost total control of the environment. In some cases, it would also allow unlimited superposition access to a *decryption oracle*.

The $(\mathcal{C})$ model

Objection:

The $(\mathcal{C})$ model is a problem if you need rewinding: how do you rewind the challenger?

Our response: rewinding the challenger would represent a scenario where the adversary has almost total control of the environment. In some cases, it would also allow unlimited superposition access to a *decryption oracle*.
In fact, if you could rewind the challenger, this would be equivalent to the $(\mathcal{O})$ model (which we prove to be unachievable in our 'security tree').

The $(\mathcal{C})$ model

Objection:

The $(\mathcal{C})$ model is a problem if you need rewinding: how do you rewind the challenger?

Our response: rewinding the challenger would represent a scenario where the adversary has almost total control of the environment. In some cases, it would also allow unlimited superposition access to a *decryption oracle*.

In fact, if you could rewind the challenger, this would be equivalent to the $(\mathcal{O})$ model (which we prove to be unachievable in our 'security tree').

Existing rewinding techniques (Watrous, Unruh) have *nothing* to do with this scenario. In fact, they rewind the adversary instead.

The (c) model

Objection:

Your (c) model is too restrictive. Consider the following example:

The (c) model

Objection:

Your (c) model is too restrictive. Consider the following example:

1. consider a collision-resistant hash function h;

The (c) model

Objection:

Your (c) model is too restrictive. Consider the following example:

1. consider a collision-resistant hash function h;
2. prepare the state $\sum_x |x, h(x)\rangle$;

The (c) model

Objection:

Your (c) model is too restrictive. Consider the following example:

1. consider a collision-resistant hash function h;
2. prepare the state $\sum_x |x, h(x)\rangle$;
3. trace out 2nd register, obtaining $\psi_y = \sum_{h(x)=y} |x\rangle \langle x|$ for random y.

The (c) model

Objection:

Your (c) model is too restrictive. Consider the following example:

1. consider a collision-resistant hash function h;
2. prepare the state $\sum_x |x, h(x)\rangle$;
3. trace out 2nd register, obtaining $\psi_y = \sum_{h(x)=y} |x\rangle \langle x|$ for random y.

Now, ψ_y was generated in poly-time, and is not entangled to anything else.

The (c) model

Objection:

Your (c) model is too restrictive. Consider the following example:

1. consider a collision-resistant hash function h;
2. prepare the state $\sum_x |x, h(x)\rangle$;
3. trace out 2nd register, obtaining $\psi_y = \sum_{h(x)=y} |x\rangle \langle x|$ for random y.

Now, ψ_y was generated in poly-time, and is not entangled to anything else. But it cannot have a classical description! Otherwise we could make two copies of it and find collisions for h.

The (c) model

Objection:

Your (c) model is too restrictive. Consider the following example:

1. consider a collision-resistant hash function h;
2. prepare the state $\sum_x |x, h(x)\rangle$;
3. trace out 2nd register, obtaining $\psi_y = \sum_{h(x)=y} |x\rangle\langle x|$ for random y.

Now, ψ_y was generated in poly-time, and is not entangled to anything else. But it cannot have a classical description! Otherwise we could make two copies of it and find collisions for h.

Our response: true, but ψ_y is not a meaningful state for the (Q) model, either! Any BQP adversary which can produce ψ_y can be purified to an adversary producing the mixture $\Psi = \sum_y \Pr(y)\psi_y$ - which *has* a classical description, and cannot be used to find collisions for h.

The Type-(2) model

Objection:

It is well known that type-(2) oracles are *more powerful* than type-(1). In fact, building an efficient circuit for a type-(2) oracle requires the secret key (or exponential loss).

The Type-(2) model

Objection:

It is well known that type-(2) oracles are *more powerful* than type-(1). In fact, building an efficient circuit for a type-(2) oracle requires the secret key (or exponential loss).

Our response: true, but recall that we are in the $(\mathcal{C})$ model, so this computation is performed by the challenger, who already knows the secret key!

The Type-(2) model

Objection:

It is well known that type-(2) oracles are *more powerful* than type-(1). In fact, building an efficient circuit for a type-(2) oracle requires the secret key (or exponential loss).

Our response: true, but recall that we are in the $(\mathcal{C})$ model, so this computation is performed by the challenger, who already knows the secret key! In fact, for the challenger it is equivalent:

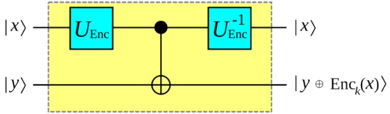

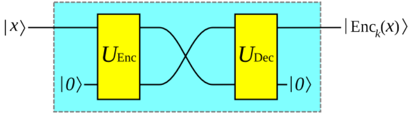

The Type-(2) model

Objection:

It is well known that type-(2) oracles are *more powerful* than type-(1). In fact, building an efficient circuit for a type-(2) oracle requires the secret key (or exponential loss).

Our response: true, but recall that we are in the ($\mathcal{C}$) model, so this computation is performed by the challenger, who already knows the secret key! In fact, for the challenger it is equivalent:

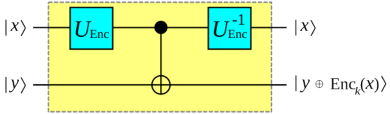

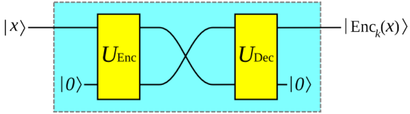

Moreover, if we use type-(1) operators we recover the (weaker) IND-qCPA notion by [BZ13] (modulo some caveats because of composition scenarios, see paper).