Title: Approximate quantum encryption and entropic security

Date: Jun 04, 2007 11:40 AM

URL: http://pirsa.org/07060022

Abstract: An approximate quantum encryption scheme uses a private key to encrypt a quantum state while leaking only a very small (though non-zero) amount of information to the adversary. Previous work has shown that while we need 2n bits of key to encrypt n qubits exactly, we can get away with only n bits in the approximate case, provided that we know that the state to be encrypted is not entangled with something that the adversary already has in his possession. In this talk I will show a generalization of this result: approximate quantum encryption requires roughly n-t bits of key, where t is a lower bound on the conditional min-entropy of the state to be encrypted given the adversary's prior knowledge. Along the way, I will introduce a quantum version of entropic security and show how the approximate quantum encryption scheme fits within this framework. This is joint work with Simon-Pierre Desrosiers.

Approximate quantum encryption and entropic security

Frédéric Dupuis, Université de Montréal
joint work with Simon-Pierre Desrosiers

June 4, 2007

Outline

Quantum encryption

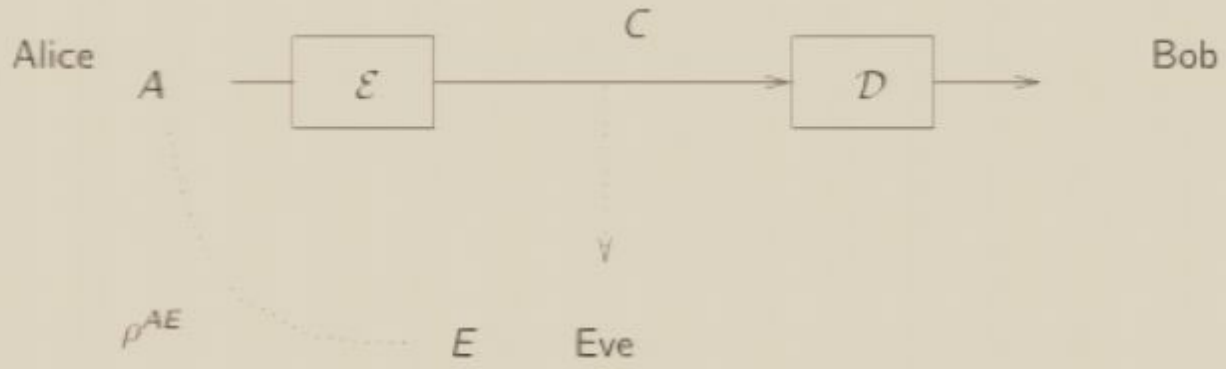

- Alice has a quantum state that she wants to send to Bob without Eve getting any information about it if she intercepts the transmission.

Quantum encryption

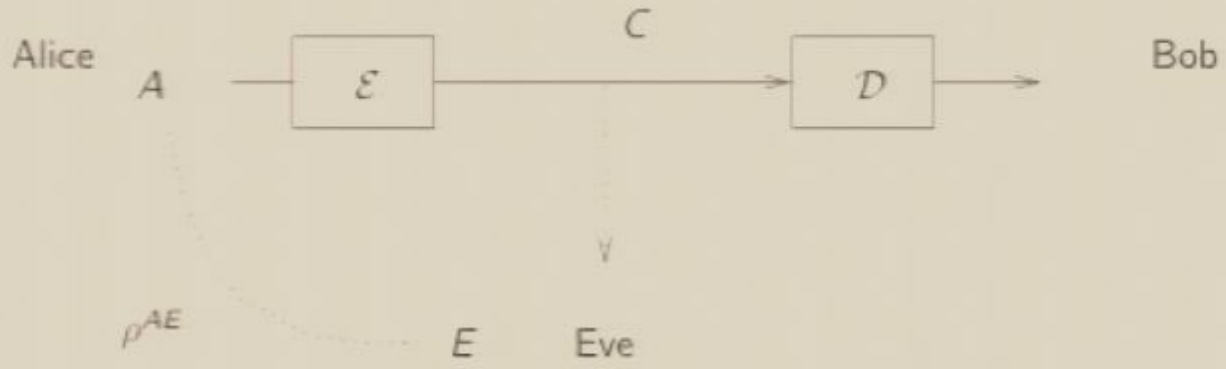

- Alice has a quantum state that she wants to send to Bob without Eve getting any information about it if she intercepts the transmission.
- Eve may have some partial quantum information about the message.

Quantum encryption

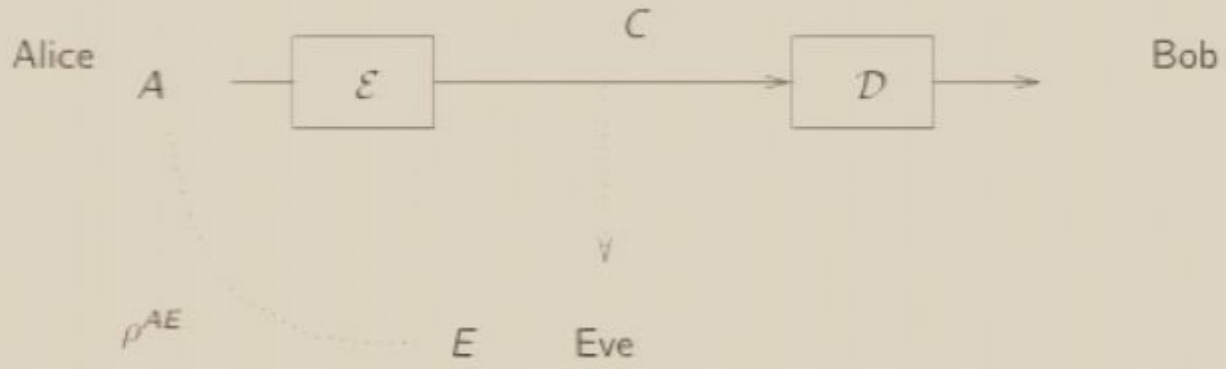

- Alice has a quantum state that she wants to send to Bob without Eve getting any information about it if she intercepts the transmission.
- Eve may have some partial quantum information about the message.

Quantum encryption

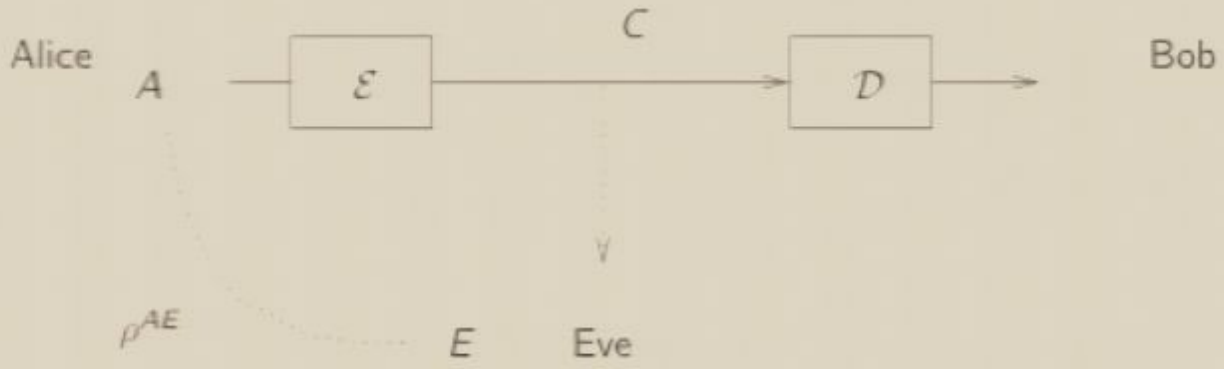

- Alice has a quantum state that she wants to send to Bob without Eve getting any information about it if she intercepts the transmission.
- Eve may have some partial quantum information about the message.

Quantum encryption

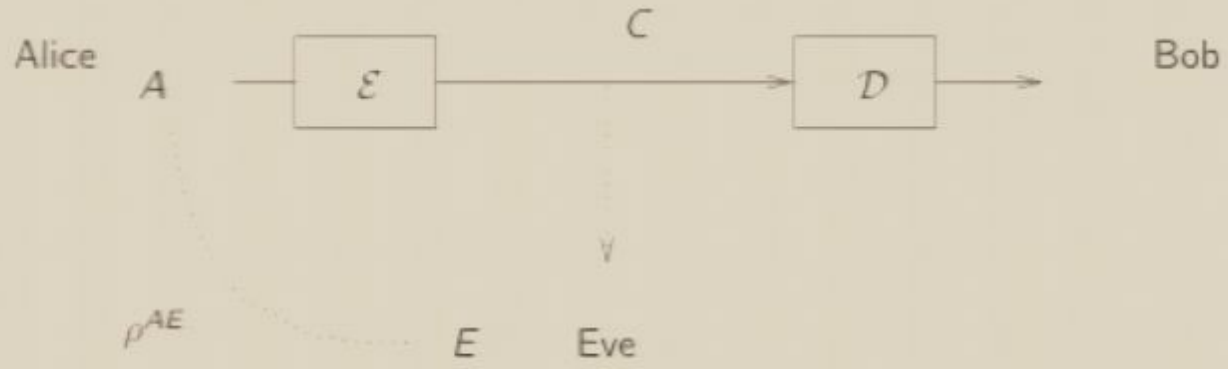

- Alice has a quantum state that she wants to send to Bob without Eve getting any information about it if she intercepts the transmission.
- Eve may have some partial quantum information about the message.
- Alice and Bob share a private classical key to encrypt the message.

Quantum encryption

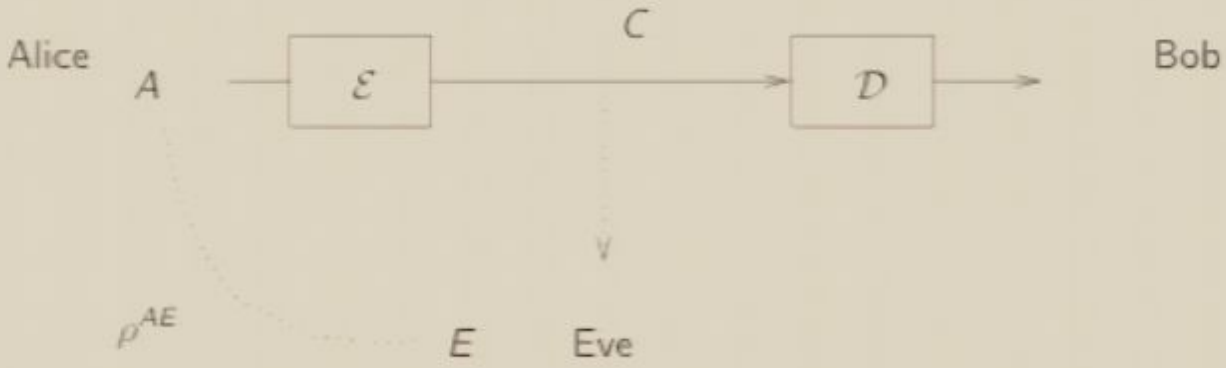

Security criterion:

$$\forall \rho^{AE} \qquad (\mathcal{E} \otimes \mathbb{I})(\rho^{AE}) = \frac{\mathbb{I}^A}{d_A} \otimes \rho^E$$

So, no matter what the input is, all Eve ever sees is her own prior information $\Rightarrow$ Eve gains no information!

Quantum encryption

Security criterion:

$$\forall \rho^{AE} \qquad (\mathcal{E} \otimes \mathbb{I})(\rho^{AE}) = \frac{\mathbb{I}^A}{d_A} \otimes \rho^E$$

So, no matter what the input is, all Eve ever sees is her own prior information $\Rightarrow$ Eve gains no information!

Quantum encryption

Quantum one-time pad

Here's a simple way to do it:

- Let the private key be the pair of n-bit strings $\langle k_X, k_Z \rangle$, chosen uniformly over all n bit strings.

Quantum encryption

Quantum one-time pad

Here's a simple way to do it:

- Let the private key be the pair of n-bit strings $\langle k_X, k_Z \rangle$, chosen uniformly over all n bit strings.

Quantum encryption

Quantum one-time pad

Here's a simple way to do it:

- Let the private key be the pair of n-bit strings $\langle k_X, k_Z \rangle$, chosen uniformly over all n bit strings.
- Then,

$$\mathcal{E}(\rho) := X^{k_x} Z^{k_z} \rho Z^{k_z} X^{k_x}$$

where $X^{k_x} := X^{k_{x,1}} \otimes X^{k_{x,2}} \otimes \cdots \otimes X^{k_{x,n}}$ and likewise for Z; and

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

Quantum encryption

Quantum one-time pad

Here's a simple way to do it:

- Let the private key be the pair of n-bit strings $\langle k_X, k_Z \rangle$, chosen uniformly over all n bit strings.
- Then,

$$\mathcal{E}(\rho) := X^{k_x} Z^{k_z} \rho Z^{k_z} X^{k_x}$$

where $X^{k_x} := X^{k_{x,1}} \otimes X^{k_{x,2}} \otimes \cdots \otimes X^{k_{x,n}}$ and likewise for Z; and

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

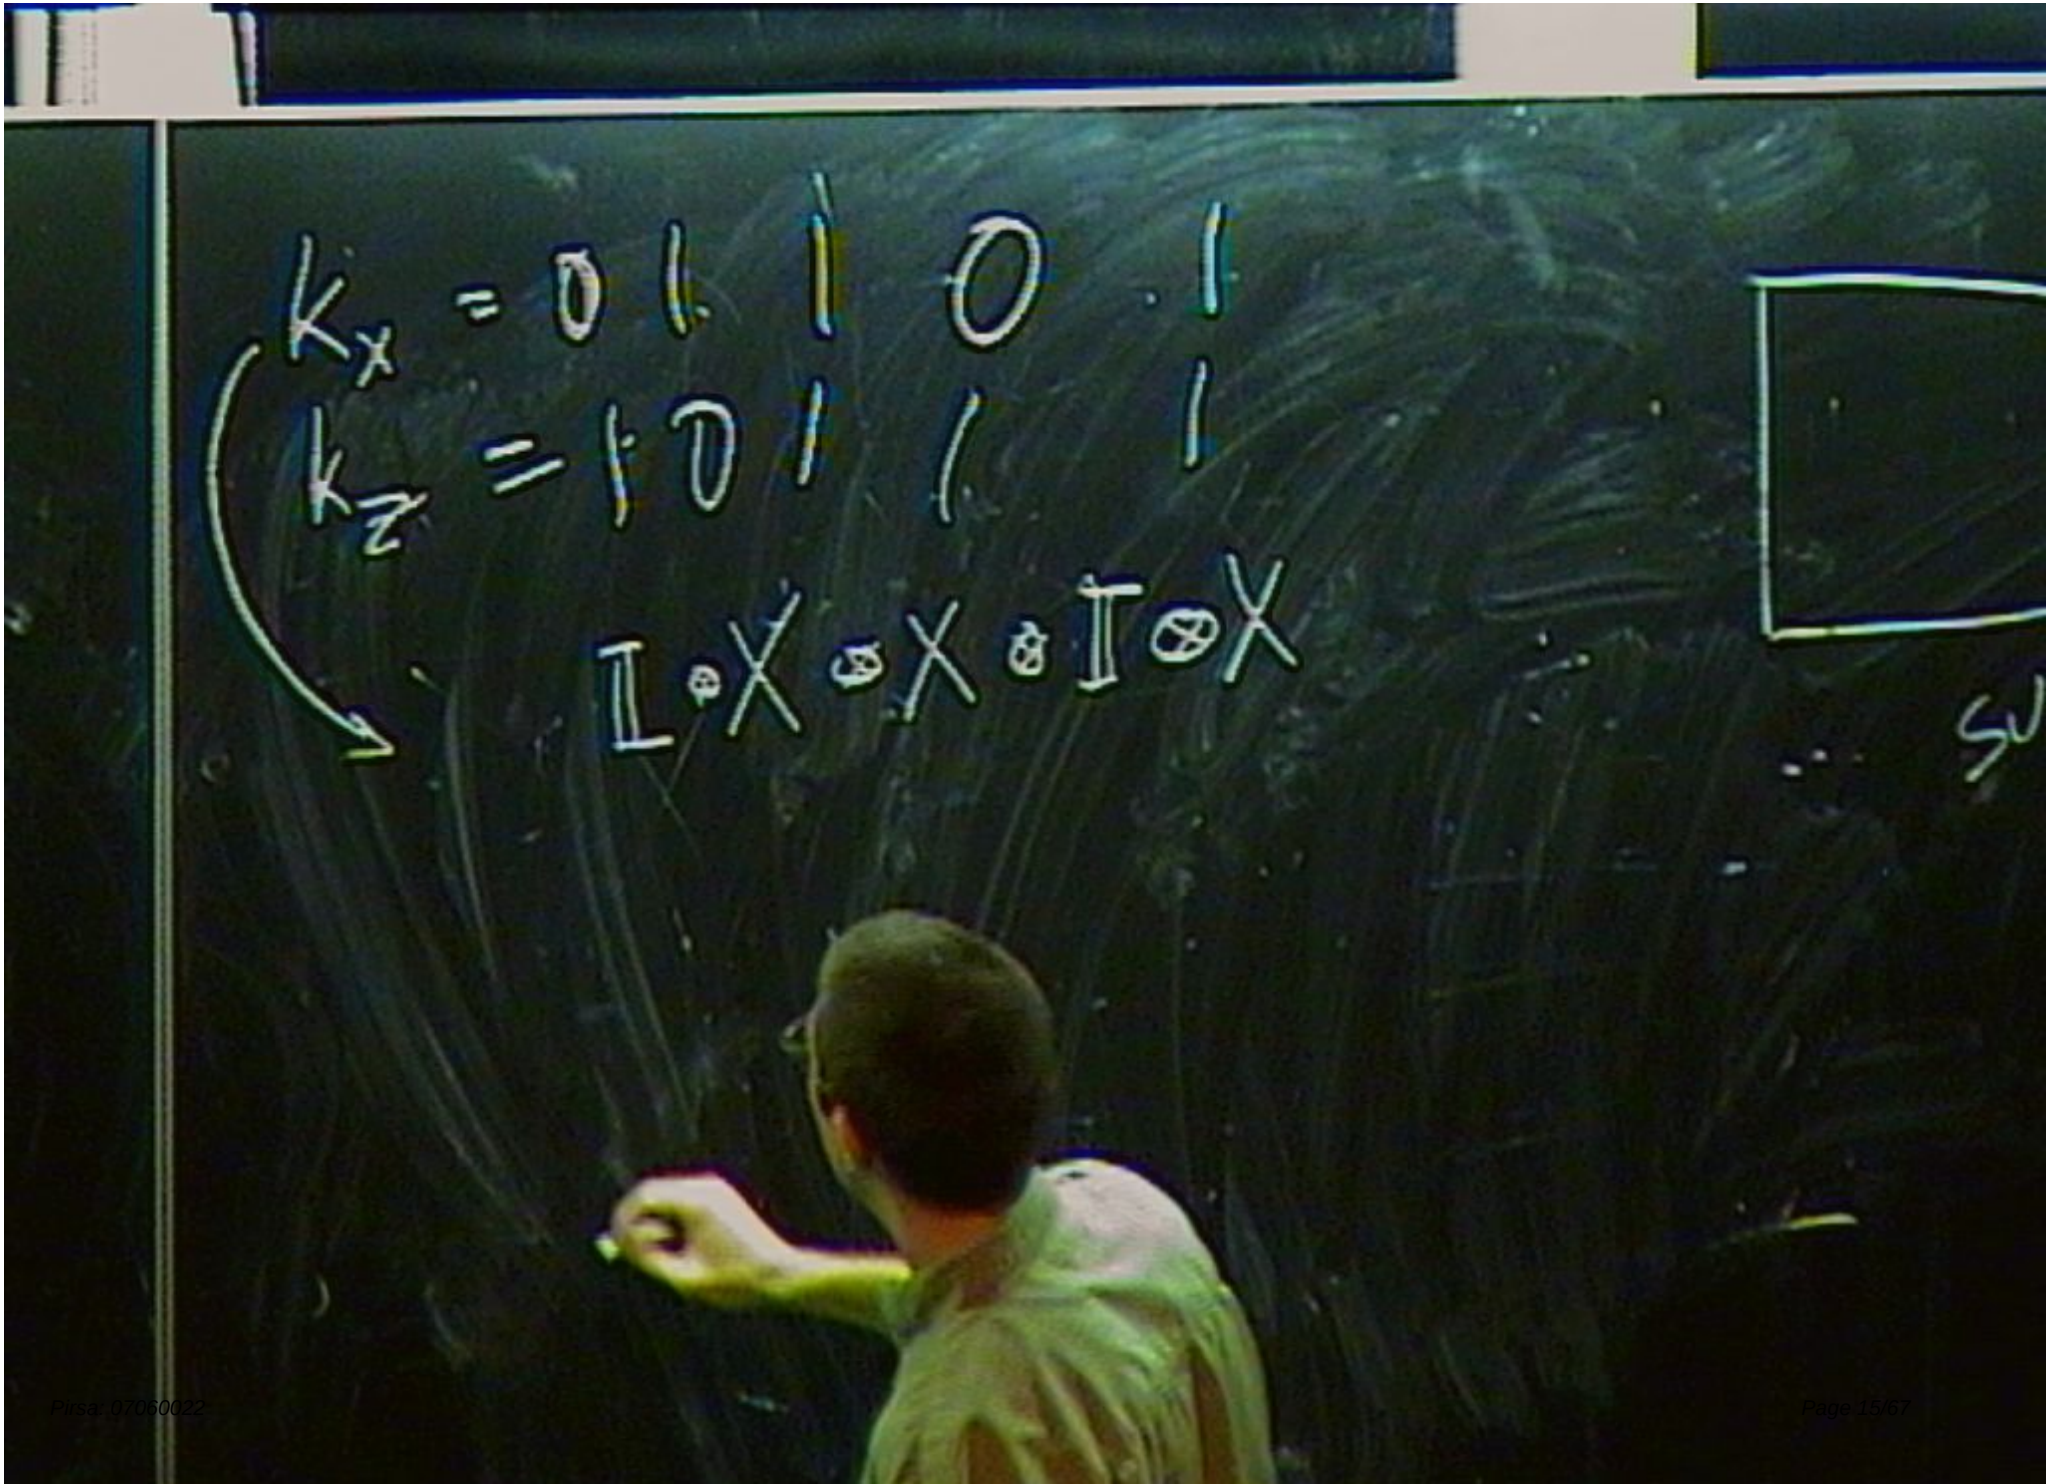
$k_x = 0\ 1\ 1\ 0\ 1$
$k_z = 1\ 0\ 1\ 1\ 1$
$I \otimes X \otimes X \otimes I \otimes X$

$k_x = 0\ 1\ 1\ 0\ 1$

$k_z = 1\ 0\ 1\ 1\ 1$

$I \otimes X \otimes X \otimes I \otimes X$

$Z \otimes I \otimes Z \otimes Z \otimes Z$

Quantum encryption

Quantum one-time pad

Here's a simple way to do it:

- Let the private key be the pair of n-bit strings $\langle k_X, k_Z \rangle$, chosen uniformly over all n bit strings.
- Then,

$$\mathcal{E}(\rho) := X^{k_x} Z^{k_z} \rho Z^{k_z} X^{k_x}$$

where $X^{k_x} := X^{k_{x,1}} \otimes X^{k_{x,2}} \otimes \cdots \otimes X^{k_{x,n}}$ and likewise for Z; and

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

Quantum encryption

Quantum one-time pad

Here's a simple way to do it:

- Let the private key be the pair of n-bit strings $\langle k_X, k_Z \rangle$, chosen uniformly over all n bit strings.
- Then,

$$\mathcal{E}(\rho) := X^{k_x} Z^{k_z} \rho Z^{k_z} X^{k_x}$$

where $X^{k_x} := X^{k_{x,1}} \otimes X^{k_{x,2}} \otimes \cdots \otimes X^{k_{x,n}}$ and likewise for Z; and

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$$

- It is easy to show that

$$(\mathcal{E} \otimes \mathbb{I})(\rho^{AE}) = \frac{\mathbb{I}}{d_A} \otimes \rho^E \qquad \forall \rho^{AE}$$

- Hence we can encrypt n qubits perfectly using $2n$ bits of key. Can we do better?

Approximate quantum encryption

- Answer: Not with this security definition. But what if we allow just a little bit of information leakage?
- Relaxed definition:

$$\forall \rho^{AE} \qquad \left\| (\mathcal{E} \otimes \mathbb{I})(\rho^{AE}) - \frac{\mathbb{I}}{d_A} \otimes \rho^E \right\|_1 \leqslant \varepsilon$$

  where $\|\rho - \sigma\|_1 := \operatorname{Tr} |\rho - \sigma|$.
- Still doesn't help: we need at least $2n - 1$ bits of key whenever $\varepsilon \leqslant \frac{1}{2}$.
- Need an additional assumption.
- In the literature so far: assume ρ^{AE} is not entangled.

Approximate quantum encryption

- Answer: Not with this security definition. But what if we allow just a little bit of information leakage?
- Relaxed definition:

$$\forall \rho^{AE} \qquad \left\| (\mathcal{E} \otimes \mathbb{I})(\rho^{AE}) - \frac{\mathbb{I}}{d_A} \otimes \rho^E \right\|_1 \leqslant \varepsilon$$

  where $\|\rho - \sigma\|_1 := \operatorname{Tr} |\rho - \sigma|$.
- Still doesn't help: we need at least $2n - 1$ bits of key whenever $\varepsilon \leqslant \frac{1}{2}$.
- Need an additional assumption.
- In the literature so far: assume ρ^{AE} is not entangled.

Approximate quantum encryption

Two methods for non-entangled states

- First proposal by Hayden, Leung, Shor, Winter using sets of random transformations brought it down to $n + \log n + 2\log(1/\varepsilon) + O(1)$.
- Two explicit constructions by Ambainis and Smith requiring $n + 2\log n + 2\log(1/\varepsilon)$ and $n + 2\log(1/\varepsilon)$ bits of key.
- So we have:
  - No entanglement: about n bits of key, same as classical
  - With entanglement: need $2n$ bits.
- Is there nothing in between? What if it's entangled only a little bit?

Approximate quantum encryption

Two methods for non-entangled states

- First proposal by Hayden, Leung, Shor, Winter using sets of random transformations brought it down to $n + \log n + 2\log(1/\varepsilon) + O(1)$.
- Two explicit constructions by Ambainis and Smith requiring $n + 2\log n + 2\log(1/\varepsilon)$ and $n + 2\log(1/\varepsilon)$ bits of key.
- So we have:
  - No entanglement: about n bits of key, same as classical
  - With entanglement: need $2n$ bits.
- Is there nothing in between? What if it's entangled only a little bit?

Classical encryption

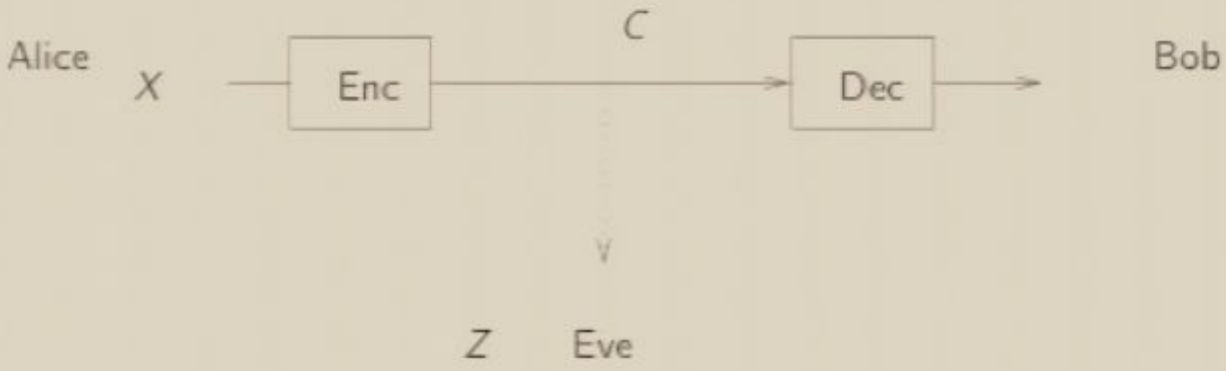

- Alice has a classical message that she wants to send to Bob without Eve getting any information about it if she intercepts the transmission.
- Eve may have some partial classical information about the message.
- Alice and Bob share a private classical key to encrypt the message.
- With no assumption on the input state, we need at least $n - 1$ bits of key. What assumption can we make to reduce this?

Classical encryption

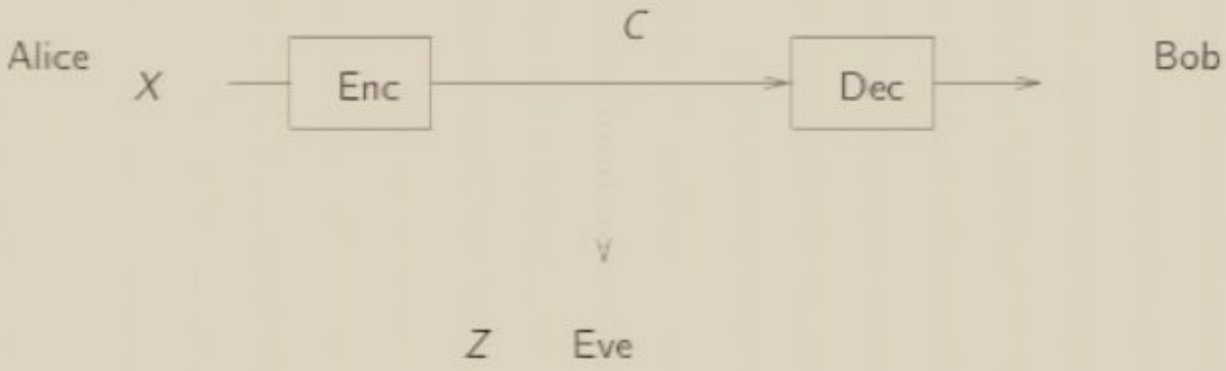

- Alice has a classical message that she wants to send to Bob without Eve getting any information about it if she intercepts the transmission.
- Eve may have some partial classical information about the message.
- Alice and Bob share a private classical key to encrypt the message.
- With no assumption on the input state, we need at least $n - 1$ bits of key. What assumption can we make to reduce this?

Entropic security

A security definition for classical information

- Suppose that we have a lower bound on Eve's prior knowledge of the message. Then we might be able to vary the key size based on this lower bound.
- A natural way to characterize this prior knowledge is the *conditional min-entropy* of the message X given the adversary's knowledge Z:

$$H_\infty(X|Z) := -\log\left[\max_{z\in\mathcal{Z},x\in\mathcal{X}} p(x|z)\right]$$

The min-entropy measures the adversary's best possible chance of guessing the message given her prior knowledge.

Entropic security

A security definition for classical information

- Suppose that we have a lower bound on Eve's prior knowledge of the message. Then we might be able to vary the key size based on this lower bound.
- A natural way to characterize this prior knowledge is the *conditional min-entropy* of the message X given the adversary's knowledge Z:

$$H_\infty(X|Z) := -\log\left[\max_{z\in\mathcal{Z},x\in\mathcal{X}} p(x|z)\right]$$

The min-entropy measures the adversary's best possible chance of guessing the message given her prior knowledge.

Entropic security and indistinguishability

Here is a security definition based on min-entropy:

Definition (Entropic indistinguishability, from Dodis and Smith)

A probabilistic encryption scheme E is (t, ε)-indistinguishable if for all message distributions such that $H_\infty(X|Z) \geqslant t$,

$$D(P_{E(X),Z}, P_U P_Z) \leqslant \varepsilon$$

Here, $D(P, Q) = \sum_{x \in \mathcal{X}} |P(x) - Q(x)|$.

Entropic security and indistinguishability

Here is a security definition based on min-entropy:

Definition (Entropic indistinguishability, from Dodis and Smith)

A probabilistic encryption scheme E is (t, ε)-indistinguishable if for all message distributions such that $H_\infty(X|Z) \geqslant t$,

$$D(P_{E(X),Z}, P_U P_Z) \leqslant \varepsilon$$

Here, $D(P, Q) = \sum_{x \in \mathcal{X}} |P(x) - Q(x)|$.

Entropic security and indistinguishability

Entropic indistinguishability can be shown to be equivalent to:

Definition (Entropic security, modified from Russell and Wang)

A probabilistic encryption scheme E is (t, ε)-entropically secure if for every adversary $\mathcal{A}$, there exists an adversary $\mathcal{A}'$ such that for all functions f, then

$$\left|\Pr[\mathcal{A}(E(X), Z) = f(X, Z)] - \Pr[\mathcal{A}'(Z) = f(X, Z)]\right| \leqslant \varepsilon$$

as long as $H_\infty(X|Z) \geqslant t$

up to small variations in the parameters t and ε.

In other words, as long as the conditional min-entropy of the input is high enough, then no adversary with the cyphertext can do significantly better than guessing without ever seeing the cyphertext.

Entropic security and indistinguishability

Entropic indistinguishability can be shown to be equivalent to:

Definition (Entropic security, modified from Russell and Wang)

A probabilistic encryption scheme E is (t, ε)-entropically secure if for every adversary $\mathcal{A}$, there exists an adversary $\mathcal{A}'$ such that for all functions f, then

$$\left|\Pr[\mathcal{A}(E(X), Z) = f(X, Z)] - \Pr[\mathcal{A}'(Z) = f(X, Z)]\right| \leqslant \varepsilon$$

as long as $H_\infty(X|Z) \geqslant t$

up to small variations in the parameters t and ε.

In other words, as long as the conditional min-entropy of the input is high enough, then no adversary with the cyphertext can do significantly better than guessing without ever seeing the cyphertext.

Entropic security and indistinguishability

Here is a security definition based on min-entropy:

Definition (Entropic indistinguishability, from Dodis and Smith)

A probabilistic encryption scheme E is (t, ε)-indistinguishable if for all message distributions such that $H_\infty(X|Z) \geqslant t$,

$$D(P_{E(X),Z}, P_U P_Z) \leqslant \varepsilon$$

Here, $D(P, Q) = \sum_{x \in \mathcal{X}} |P(x) - Q(x)|$.

Entropic security and indistinguishability

Entropic indistinguishability can be shown to be equivalent to:

Definition (Entropic security, modified from Russell and Wang)

A probabilistic encryption scheme E is (t, ε)-entropically secure if for every adversary $\mathcal{A}$, there exists an adversary $\mathcal{A}'$ such that for all functions f, then

$$\left|\Pr[\mathcal{A}(E(X), Z) = f(X, Z)] - \Pr[\mathcal{A}'(Z) = f(X, Z)]\right| \leqslant \varepsilon$$

as long as $H_\infty(X|Z) \geqslant t$

up to small variations in the parameters t and ε.

In other words, as long as the conditional min-entropy of the input is high enough, then no adversary with the cyphertext can do significantly better than guessing without ever seeing the cyphertext.

Entropic security and indistinguishability

Entropic indistinguishability can be shown to be equivalent to:

Definition (Entropic security, modified from Russell and Wang)

A probabilistic encryption scheme E is (t, ε)-entropically secure if for every adversary $\mathcal{A}$, there exists an adversary $\mathcal{A}'$ such that for all functions f, then

$$\left|\Pr[\mathcal{A}(E(X), Z) = f(X, Z)] - \Pr[\mathcal{A}'(Z) = f(X, Z)]\right| \leqslant \varepsilon$$

as long as $H_\infty(X|Z) \geqslant t$

up to small variations in the parameters t and ε.

In other words, as long as the conditional min-entropy of the input is high enough, then no adversary with the cyphertext can do significantly better than guessing without ever seeing the cyphertext.

Entropic security and indistinguishability

- How much key do we need to achieve this?
- Dodis and Smith present a scheme with $n - t + 2\log(1/\varepsilon) + 2$ bits of key.
- Can we get a quantum version of this?

Entropic security and indistinguishability

- How much key do we need to achieve this?
- Dodis and Smith present a scheme with $n - t + 2\log(1/\varepsilon) + 2$ bits of key.
- Can we get a quantum version of this?

Our results

- We generalize the notions of entropic security and indistinguishability to the quantum world.
- We show how to make (t, ε)-indistinguishable quantum encryption schemes requiring approximately $n - t$ bits of key.
- We prove that the two Ambainis-Smith schemes fit these security definitions unmodified.
- We give a simple lower bound of $n - t - 1$ bits of key.

Our results

- We generalize the notions of entropic security and indistinguishability to the quantum world.
- We show how to make (t, ε)-indistinguishable quantum encryption schemes requiring approximately $n - t$ bits of key.
- We prove that the two Ambainis-Smith schemes fit these security definitions unmodified.
- We give a simple lower bound of $n - t - 1$ bits of key.

Quantum conditional min-entropy

- What would be a good quantum generalization of conditional min-entropy?
- Renner gave one in his PhD thesis:

$$H_\infty(\rho^{AE}|\rho^E) = -\log\lambda$$

where λ is the minimum real number such that the Hermitian operator $\lambda\mathbb{I}^A \otimes \rho^E - \rho^{AE}$ is positive semi-definite.

Quantum conditional min-entropy

- What would be a good quantum generalization of conditional min-entropy?
- Renner gave one in his PhD thesis:

$$H_\infty(\rho^{AE}|\rho^E) = -\log\lambda$$

where λ is the minimum real number such that the Hermitian operator $\lambda\mathbb{I}^A \otimes \rho^E - \rho^{AE}$ is positive semi-definite.

Quantum conditional min-entropy

An equivalent formulation of quantum conditional min-entropy:

$$H_\infty(\rho^{AE}|\rho^E) = -\log\left[\max_{|\psi\rangle} \frac{\langle\psi|\rho^{AE}|\psi\rangle}{\langle\psi|\mathbb{I}\otimes\rho^E|\psi\rangle}\right]$$

analogous to the classical case:

$$H_\infty(X|Z) = -\log\left[\max_{x,z} \frac{p(x,z)}{p(z)}\right]$$

Quantum conditional min-entropy

An equivalent formulation of quantum conditional min-entropy:

$$H_\infty(\rho^{AE}|\rho^E) = -\log\left[\max_{|\psi\rangle} \frac{\langle\psi|\rho^{AE}|\psi\rangle}{\langle\psi|\mathbb{I}\otimes\rho^E|\psi\rangle}\right]$$

analogous to the classical case:

$$H_\infty(X|Z) = -\log\left[\max_{x,z} \frac{p(x,z)}{p(z)}\right]$$

Quantum conditional min-entropy

A few properties of quantum conditional min-entropy:

- For an n-qubit state:

$$-n \leqslant H_\infty(\rho^{AE}|\rho^E) \leqslant n$$

  The lower bound is saturated by a maximally entangled state between A and E, and the upper bound, by the maximally-mixed state.
- If ρ^{AE} is separable, then $H_\infty(\rho^{AE}|\rho^E) \geqslant 0$.

Quantum conditional min-entropy

A few properties of quantum conditional min-entropy:

- For an n-qubit state:

$$-n \leqslant H_\infty(\rho^{AE}|\rho^E) \leqslant n$$

  The lower bound is saturated by a maximally entangled state between A and E, and the upper bound, by the maximally-mixed state.
- If ρ^{AE} is separable, then $H_\infty(\rho^{AE}|\rho^E) \geqslant 0$.

Quantum entropic indistinguishability

We can use this to define a new notion of security:

Definition (Quantum entropic indistinguishability)

A quantum encryption system $\mathcal{E}$ is (t, ϵ)-indistinguishable if for all states ρ^{AE} such that $H_\infty(\rho^{AE}|\rho^E) \geqslant t$ we have that:

$$\left\| (\mathcal{E} \otimes \mathbb{I})(\rho^{AE}) - \frac{\mathbb{I}}{d_A} \otimes \rho^E \right\|_1 \leqslant \varepsilon$$

Quantum entropic indistinguishability

We can use this to define a new notion of security:

Definition (Quantum entropic indistinguishability)

A quantum encryption system $\mathcal{E}$ is (t, ϵ)-indistinguishable if for all states ρ^{AE} such that $H_\infty(\rho^{AE}|\rho^E) \geqslant t$ we have that:

$$\left\| (\mathcal{E} \otimes \mathbb{I})(\rho^{AE}) - \frac{\mathbb{I}}{d_A} \otimes \rho^E \right\|_1 \leqslant \varepsilon$$

Quantum entropic security

How do we generalize the concept of "function on the input"?

Definition (Quantum entropic security)

An quantum encryption system $\mathcal{E}$ is (t, ε)-entropically secure if for all states ρ^{AE} such that $H_{\min}(\rho^{AE}|\rho^{E}) \geqslant t$, all interpretations $\{(p_j, \sigma_j^{AE})\}$ and all adversaries A, there exists an A' such that for all functions f, we have

$$\left|\Pr[A((\mathcal{E} \otimes \mathbb{I})(\sigma_i^{AE})) = f(i)] - \Pr[A'(\sigma_i^{E}) = f(i)]\right| \leqslant \varepsilon.$$

Quantum entropic security

How do we generalize the concept of "function on the input"?

Definition (Quantum entropic security)

An quantum encryption system $\mathcal{E}$ is (t, ε)-entropically secure if for all states ρ^{AE} such that $H_{\min}(\rho^{AE}|\rho^{E}) \geqslant t$, all interpretations $\{(p_j, \sigma_j^{AE})\}$ and all adversaries A, there exists an A' such that for all functions f, we have

$$\left|\Pr[A((\mathcal{E} \otimes \mathbb{I})(\sigma_i^{AE})) = f(i)] - \Pr[A'(\sigma_i^{E}) = f(i)]\right| \leqslant \varepsilon.$$

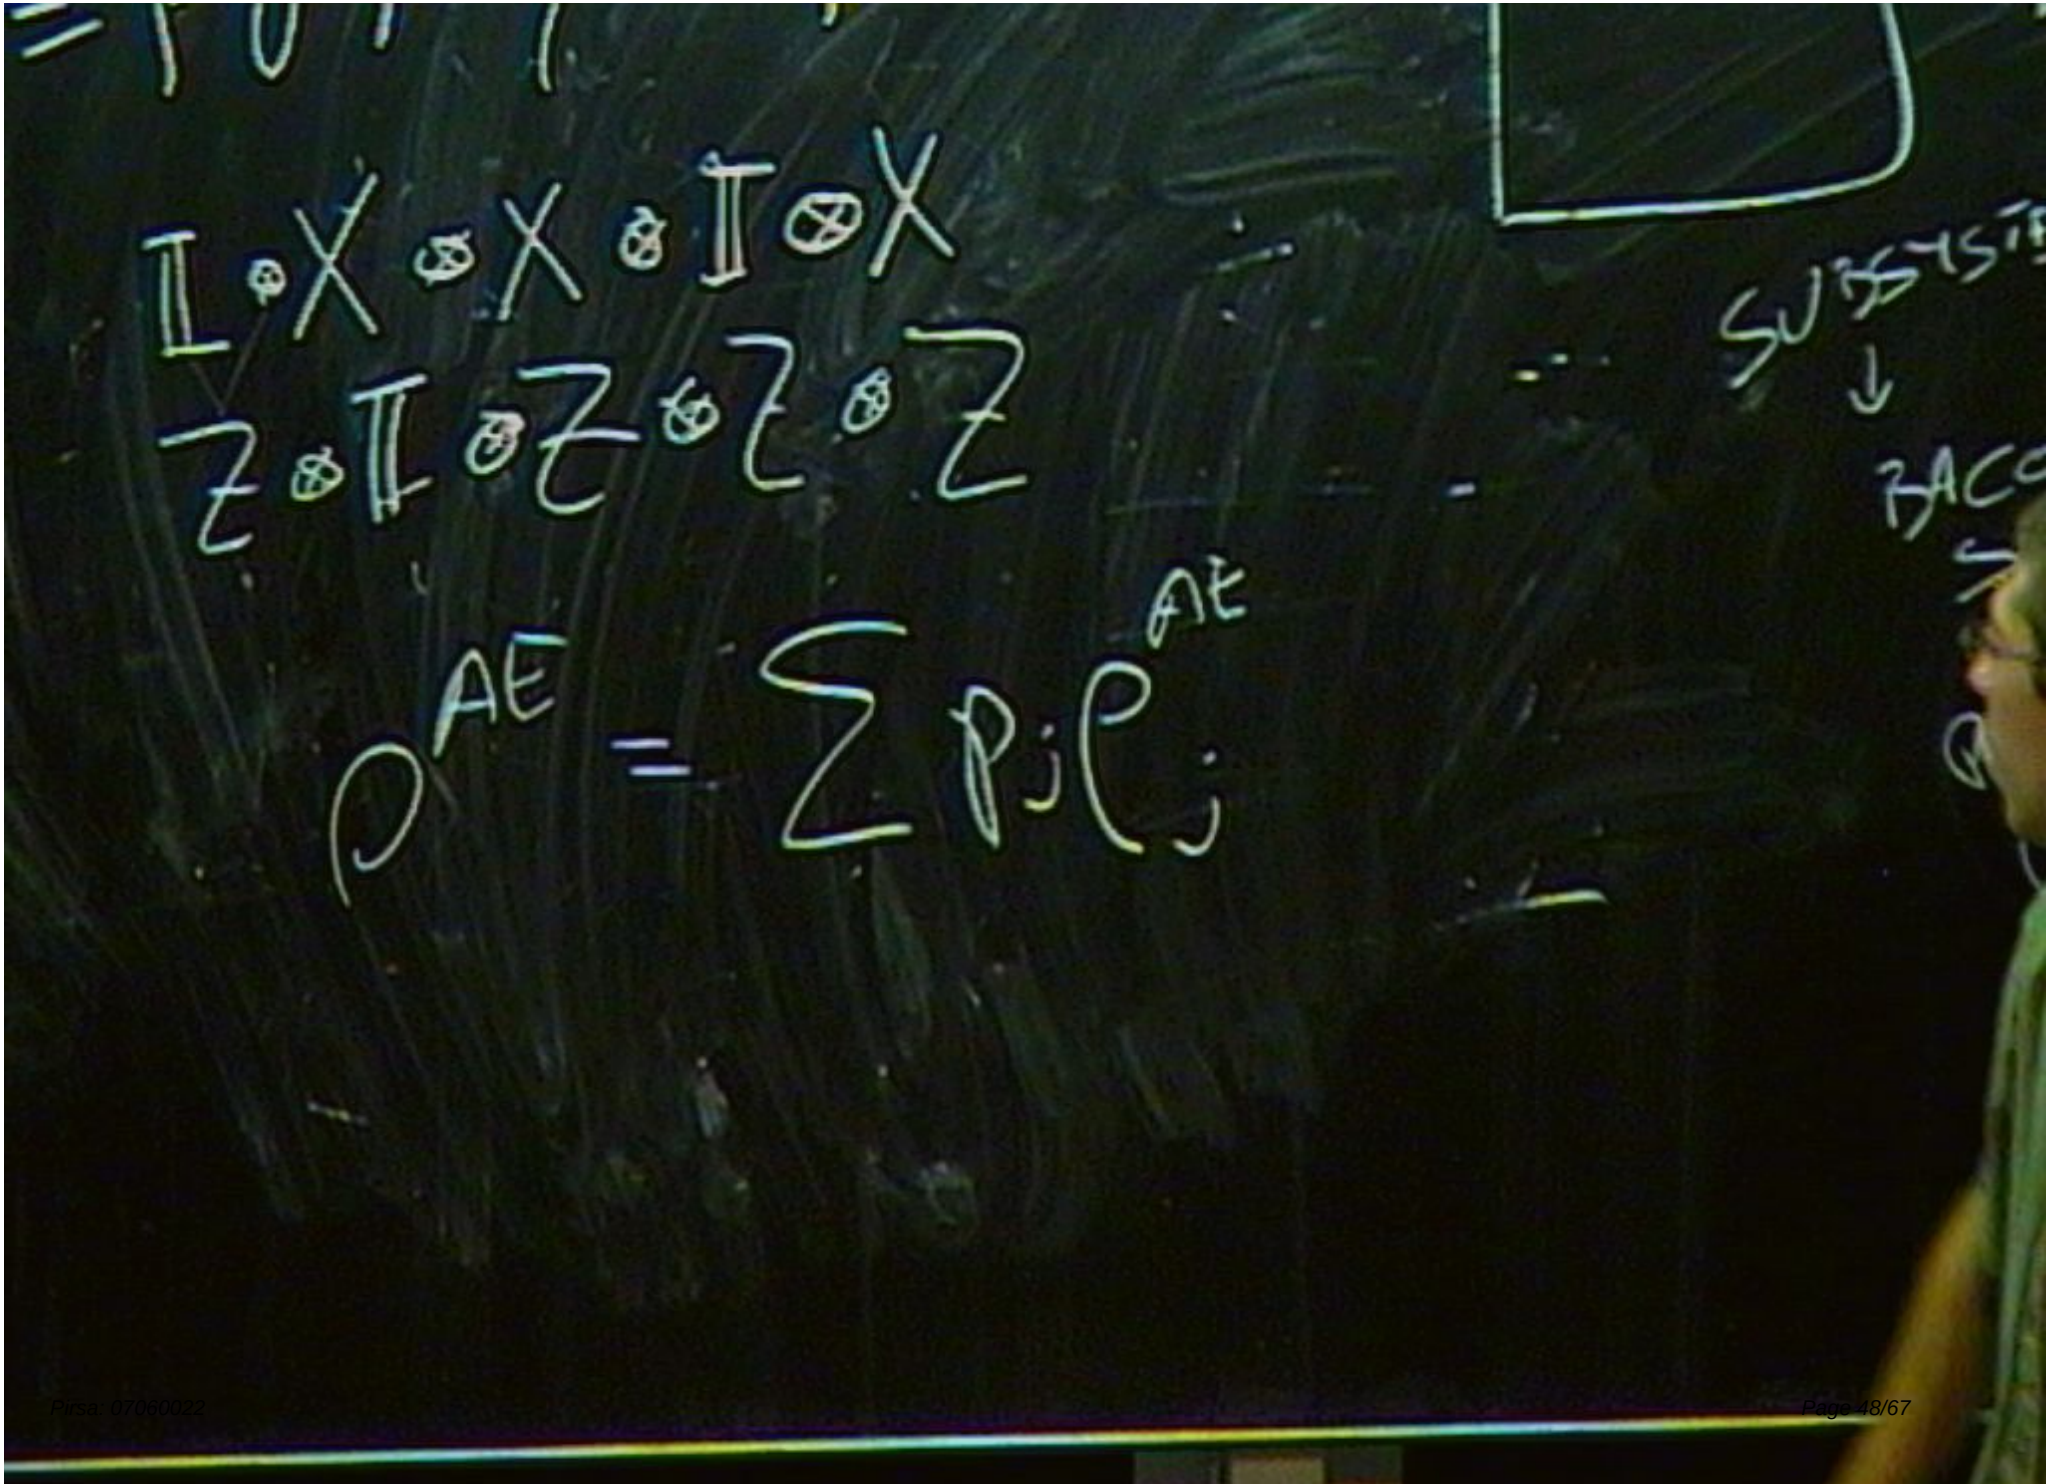
I⊗X⊗X⊗I⊗X
Z⊗I⊗Z⊗Z⊗Z
$\rho^{AE} = \sum p_j \rho_j^{AE}$

Quantum entropic security

How do we generalize the concept of "function on the input"?

Definition (Quantum entropic security)

An quantum encryption system $\mathcal{E}$ is (t, ε)-entropically secure if for all states ρ^{AE} such that $H_{\min}(\rho^{AE}|\rho^{E}) \geqslant t$, all interpretations $\{(p_j, \sigma_j^{AE})\}$ and all adversaries A, there exists an A' such that for all functions f, we have

$$\left|\Pr[A((\mathcal{E} \otimes \mathbb{I})(\sigma_i^{AE})) = f(i)] - \Pr[A'(\sigma_i^{E}) = f(i)]\right| \leqslant \varepsilon.$$

Quantum entropic security and indistinguishability

We can show that

- $(t-1, \varepsilon/2)$-indistinguishability implies (t, ε)-entropic security.
- (t, ε)-entropic security implies $(t-1, 6\varepsilon)$-indistinguishability as long as $t \leqslant n-1$.

Quantum entropic security and indistinguishability

We can show that

- $(t-1, \varepsilon/2)$-indistinguishability implies (t, ε)-entropic security.
- (t, ε)-entropic security implies $(t-1, 6\varepsilon)$-indistinguishability as long as $t \leqslant n-1$.

The first Ambainis-Smith scheme

- Perfect encryption: $\mathcal{E}(\rho) = \sum_{(k_x,k_z)} X^{k_x} Z^{k_z} \rho Z^{k_z} X^{k_x}$ with k_x and k_z picked at random.
- How about picking k_x and k_z from a smaller set of bitstrings?

The first Ambainis-Smith scheme

- Perfect encryption: $\mathcal{E}(\rho) = \sum_{(k_x,k_z)} X^{k_x} Z^{k_z} \rho Z^{k_z} X^{k_x}$ with k_x and k_z picked at random.
- How about picking k_x and k_z from a smaller set of bitstrings?

The first Ambainis-Smith scheme

Definition (δ-biased set)

A set $S \subseteq \{0,1\}^n$ is said to be δ-biased iff for every $s' \in \{0,1\}^n, s' \neq 0^n$, we have that $\left|\mathbb{E}_{s \leftarrow S}\left[(-1)^{s \odot s'}\right]\right| \leqslant \delta$.

Picking strings from a δ-biased set is almost like choosing strings at random when it comes to taking parities.

The construction we need [Alon, Goldreich, Håstad, Peralta] yields sets of size n^2/δ^2.

The first Ambainis-Smith scheme

Definition (δ-biased set)

A set $S \subseteq \{0,1\}^n$ is said to be δ-biased iff for every $s' \in \{0,1\}^n, s' \neq 0^n$, we have that $\left|\mathbb{E}_{s \leftarrow S}\left[(-1)^{s \odot s'}\right]\right| \leqslant \delta$.

Picking strings from a δ-biased set is almost like choosing strings at random when it comes to taking parities.

The construction we need [Alon, Goldreich, Håstad, Peralta] yields sets of size n^2/δ^2.

The first Ambainis-Smith scheme

$$\mathcal{E}(\rho) = \sum_{(k_x,k_z)} X^{k_x} Z^{k_z} \rho Z^{k_z} X^{k_x}$$

If we pick (k_x, k_z) from a δ-biased set of $2n$-bit strings, and that $H_\infty(\rho^{AE}|\rho^E) \geqslant t$, we can show that

$$\left\| (\mathcal{E} \otimes \mathbb{I})(\rho^{AE}) - \frac{\mathbb{I}}{2^n} \otimes \rho^E \right\|_1 \leqslant \delta \sqrt{2^{n-t}}$$

So if we pick $\delta = \varepsilon / \sqrt{2^{n-t}}$, we're fine.

How many bits of key do we need for that?
$\log n^2/\delta^2 = \log[2^{(n-t)} n^2/\varepsilon^2] = n - t + 2\log n + 2\log(1/\varepsilon)$.

The first Ambainis-Smith scheme

Definition (δ-biased set)

A set $S \subseteq \{0,1\}^n$ is said to be δ-biased iff for every $s' \in \{0,1\}^n, s' \neq 0^n$, we have that $\left|\mathbb{E}_{s \leftarrow S}\left[(-1)^{s \odot s'}\right]\right| \leqslant \delta$.

Picking strings from a δ-biased set is almost like choosing strings at random when it comes to taking parities.

The construction we need [Alon, Goldreich, Håstad, Peralta] yields sets of size n^2/δ^2.

The first Ambainis-Smith scheme

$$\mathcal{E}(\rho) = \sum_{(k_x,k_z)} X^{k_x} Z^{k_z} \rho Z^{k_z} X^{k_x}$$

If we pick (k_x, k_z) from a δ-biased set of $2n$-bit strings, and that $H_\infty(\rho^{AE}|\rho^E) \geqslant t$, we can show that

$$\left\| (\mathcal{E} \otimes \mathbb{I})(\rho^{AE}) - \frac{\mathbb{I}}{2^n} \otimes \rho^E \right\|_1 \leqslant \delta \sqrt{2^{n-t}}$$

So if we pick $\delta = \varepsilon / \sqrt{2^{n-t}}$, we're fine.

How many bits of key do we need for that?
$\log n^2/\delta^2 = \log[2^{(n-t)} n^2/\varepsilon^2] = n - t + 2\log n + 2\log(1/\varepsilon)$.

The first Ambainis-Smith scheme

$$\mathcal{E}(\rho) = \sum_{(k_x,k_z)} X^{k_x} Z^{k_z} \rho Z^{k_z} X^{k_x}$$

If we pick (k_x, k_z) from a δ-biased set of $2n$-bit strings, and that $H_\infty(\rho^{AE}|\rho^E) \geqslant t$, we can show that

$$\left\| (\mathcal{E} \otimes \mathbb{I})(\rho^{AE}) - \frac{\mathbb{I}}{2^n} \otimes \rho^E \right\|_1 \leqslant \delta \sqrt{2^{n-t}}$$

So if we pick $\delta = \varepsilon / \sqrt{2^{n-t}}$, we're fine.

How many bits of key do we need for that?
$\log n^2/\delta^2 = \log[2^{(n-t)} n^2/\varepsilon^2] = n - t + 2\log n + 2\log(1/\varepsilon)$.

The Ambainis-Smith scheme

Ambainis and Smith also give a second scheme which requires only $n - t + 2\log(1/\varepsilon)$ bits of key, but doubles the size of the ciphertext.

The Ambainis-Smith scheme

Ambainis and Smith also give a second scheme which requires only $n - t + 2\log(1/\varepsilon)$ bits of key, but doubles the size of the ciphertext.

Connection with previous results

Notice that we can retrieve previous results from this:

- If we assume no entanglement between Alice and Eve, we implicitly have $t \geqslant 0$, and hence we need $\approx n$ bits of key.
- If we have no bound whatsoever on Eve's knowledge, our best bound is $t \geqslant -n$ and we need around $2n$ bits of key.

Connection with previous results

Notice that we can retrieve previous results from this:

- If we assume no entanglement between Alice and Eve, we implicitly have $t \geqslant 0$, and hence we need $\approx n$ bits of key.
- If we have no bound whatsoever on Eve's knowledge, our best bound is $t \geqslant -n$ and we need around $2n$ bits of key.

A simple lower bound on the key length

We can show that $n - t$ is essentially the optimal number of key bits. Let's assume Alice wants to encrypt a state which consists of:

- $(n - t)/2$ halves of EPR pairs, with the other halves in Eve's hands
- $(n + t)/2$ maximally mixed qubits.

It is easy to show that that $H_\infty(\rho^{AE}|\rho^E) = t$: the EPR pairs contribute $-(n - t)/2$ and the rest contributes $(n + t)/2$ to the conditional min-entropy.

To encrypt this, we have to at least be able to encrypt the halves of EPR pairs; but we can show that that takes at least $2(n - t)/2 - 1 = n - t - 1$ bits of key.

Conclusion and open problems

Recap of what we have done:

- Gave a meaningful quantum version of entropic security equivalent to a simple definition involving the trace distance;
- Showed that these definitions lead to a more complete understanding of approximate quantum encryption;
- Presented a encryption scheme which achieves these security definitions;
- Showed that it is nearly optimal.

Conclusion and open problems

Some open problems:

- Explicit schemes using other norms (especially operator norm)?
- Quantum extractors?
  - An extractor takes a source with min-entropy at least t and a small random seed, and outputs a uniform distribution.

The End

Thank you!