Empowered by Innovation NEC

Order-Preserving Encryption Secure Beyond One-Wayness

Isamu Teranishi （NEC）

Moti Yung (Google, Columbia University)

Tal Malkin (Columbia University)

Order Preserving Encryption (OPE)

Secret Key Encryption Scheme s.t.

- Plaintext and Ciphertext Spaces are intervals of the set of integers.
- It satisfies the order-preserving property:

$$m < m' \Leftrightarrow \mathrm{Enc}_K(m) < \mathrm{Enc}_K(m')$$

Application

- OPE can be used in encrypted outsourced database
- (**Range Query**) Because OPE enables one to find documents m satisfying

$$a<m<b$$

without decrypting ciphertexts.

- In fact, due to the order-pres. property, one can find such m by checking whether

$$\mathrm{Enc}_K(a) < \mathrm{Enc}_K(m) < \mathrm{Enc}_K(b)$$

holds or not.

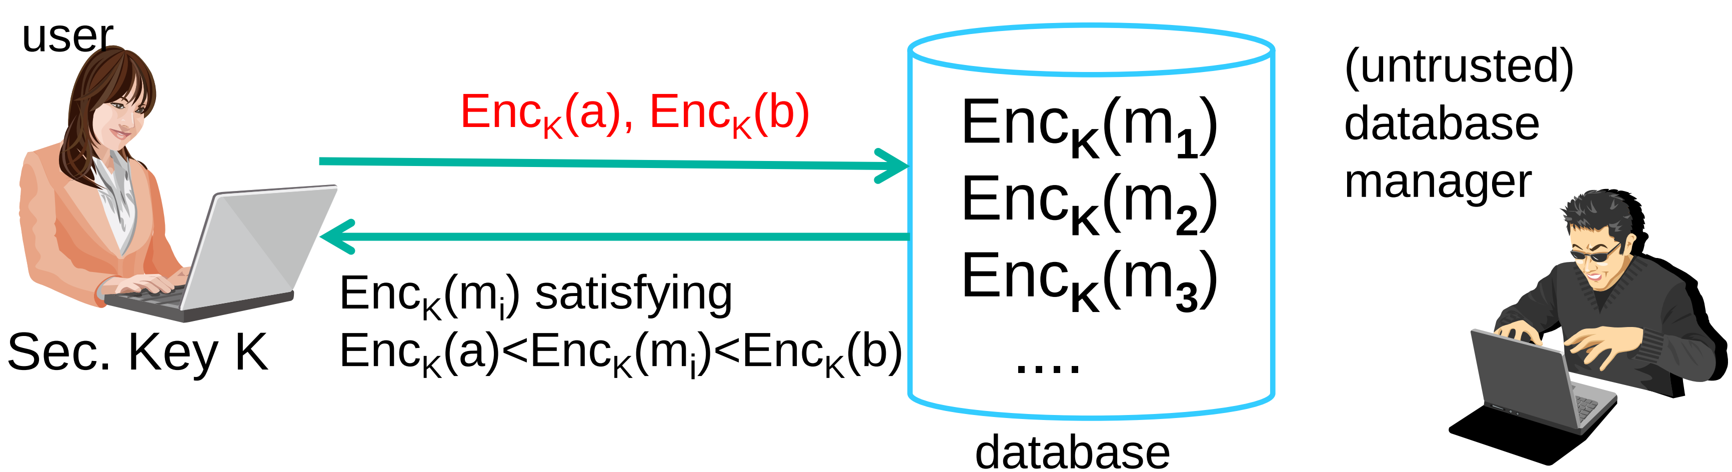

Subject and Results of This Paper

- However, security of OPE is far from being understood at this time.
  - In fact, a naturally defined indistinguishability notion (IND-O-CPA) cannot be achievable (under some natural condition) [1].

- In this paper we tackle the following fundamental problem for OPE:

what exactly must OPE leak?,

and what can it hide?

- And we show a positive results for it:
  - Define a weaker indistinguishability notion, (χ,T,q)-IND, for OPE than the known (unachievable) one while the known result[2]is about one-wayness
    - the notion is natural in the database setting mentioned before.
    - the notion can ensure that secrecy of lower bits of plaintext.
  - Propose a new OPE scheme satisfying our indistinguishability notion.

[1] Boldyreva, Chenette, Lee, O'Neill: Order-Preserving Symmetric Encryption. EUROCRYPT 2009: 224-241

Rest of This Talk

- Our Definition of Indistiguishability Notion
- Our Results
- Construction of Our scheme
- Security Proof

- Our Definition of Indistiguishability Notion
- Our Results
- Construction of Our scheme
- Security Proof

Review of (r,q+1)-WOW (Window-OneWay)

Our security notion is obtained by modifying the following known one-way based notion, (r,q+1)-WOW [2]

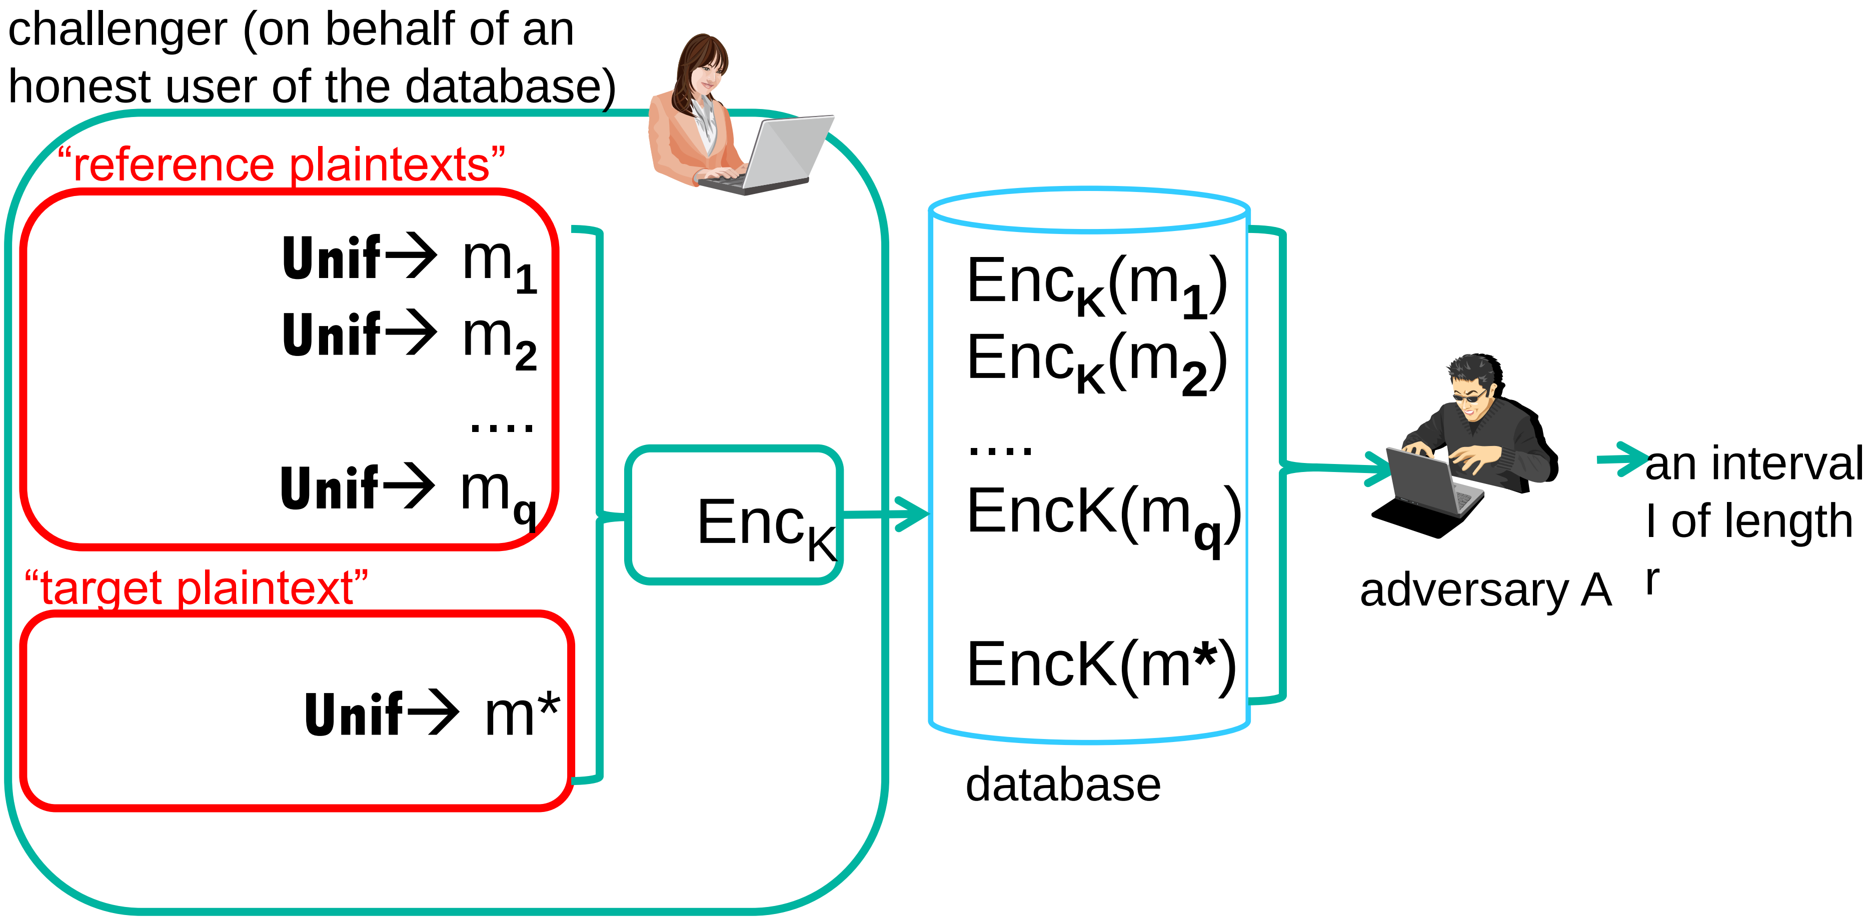

$$\forall A \text{ (polytime) } \Pr[m^* \in I] \leq \text{neg(Mess. Sp. Size)}$$

[2] Boldyreva, Chenette, O'Neill: Order-Preserving Encryption Revisited: Improved Security Analysis and Alternative Solutions. CRYPTO 2011: 578-595

Our notion $(\mathcal{X}, T, q)$-IND

Here $\mathcal{X} = (\mathcal{X}_1, \ldots, \mathcal{X}_q)$ be a tuple of (indep.) distributions on the Mess. Sp.

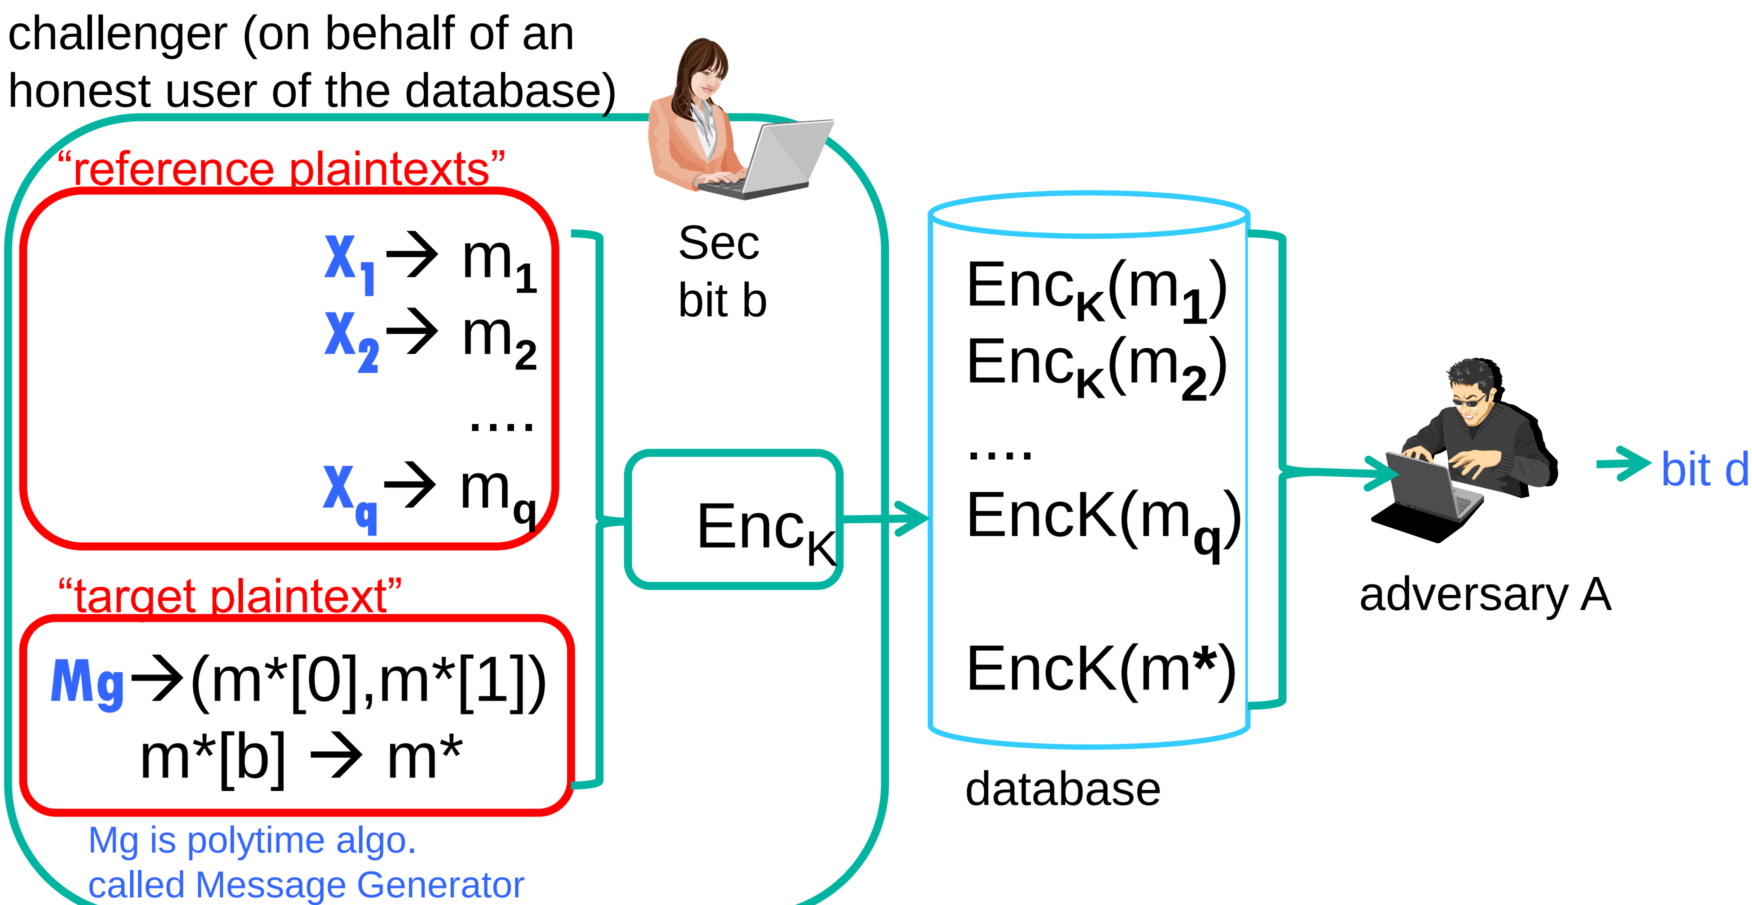

$\forall$Mg (polytime) whose output satisfies $|m^*[0] - m^*[1]| < T$

$\forall$ A (polytime) $|\Pr[d=b] - 1/2| \leq \mathrm{neg}(\text{Mess. Sp. Size})$

Why $|m^*[0]-m^*[1]| < T$?

- In our def., we require a message generator **Mg** to output $(m^*[0],m^*[1])$ satisfying

$$|m^*[0]-m^*[1]| < T$$

- This is because otherwise, an OPE is broken easily

  using the following idea [1]:

  - The order-pres. property

  $$m < m' \Rightarrow Enc_K(m) < Enc_K(m')$$

    means that Enc_K is monotone increasing.

  - Hence, if we allow an adversary to select $(m^*[0],m^*[1])$ such that

  $$m^*[1] - m^*[0]$$

    is large, the difference

  $$Enc_K(m^*[1])-Enc_K(m^*[0])$$

    has to become noticeably large.

  - Therefore, the adversary can distinguish $Enc_K(m^*[0])$ and $Enc_K(m^*[1])$ easily

Property of ($\boldsymbol{\chi}$,T,q)-IND

Our ($\boldsymbol{\chi}$,T,q)-IND implies that the least significant log T bits of a plaintext are hidden from the adversary in our database setting.

Proof (rough idea)

- Consider the following two messages:

  m*[0] : any message

  m*[1] : lower log T bits are selected randomly and the other bits are the same as those of m*[0]

- Then, it holds that

$$|m^*[0]-m^*[1]| < T,$$

  which is our condition for ($\boldsymbol{\chi}$,R,q)-IND.

- Hence, Enc_K(m*[0]) is indis. from Enc_K(m*[1]).
- Recall that the lower log T bits of m*[1] is random.
- This means that an adversary given Enc_K(m*[0]) cannot know the lower log T bits of m*[0].

- Our Definition of Indistiguishability Notion
- Our Results
- Construction of Our scheme
- Security Proof

Our Result (Informal)

- Very roughly, we construct an OPE scheme such that

Main Thm.(informal) if min-entropies of $\mathbf{X}_1,\ldots,\mathbf{X}_q$ are large, our scheme is $(\mathbf{X},T,q)$-IND for a large T. (Here $\mathbf{X}=(\mathbf{X}_1,\ldots,\mathbf{X}_q)$.)

- To formalize the above statement, we give some def.
  - The **min-entropy** of random variable $\mathbf{X}_i$ on a Mess. Sp. is

$$H_\infty(\mathbf{X}_i) := \min \{ -\log \Pr[\mathbf{X}_i = m] \mid m \in \text{Mess. Sp.}\}$$

  - It is known that the min-entropy of $\mathbf{X}_i$ has to less than that of **Unif** on Mess. Sp:

$$H_\infty(\mathbf{X}_i) \leq H_\infty(\mathbf{Unif}) \quad (= \log \#(\text{Mess. Sp.}))$$

  - So we define "***normalized***" **min-entropy** of $\mathbf{X}$ as follows:

$$H^*_\infty(\mathbf{X}_i) := H_\infty(\mathbf{X}_i) / H_\infty(\mathbf{Unif}) \leq 1$$

  - for a *tuple* $\mathbf{X}=(\mathbf{X}_1,\ldots,\mathbf{X}_q)$ of random variables, we also define

Our Result (Formal)

We construct an OPE scheme E[α, β] satisfying the following property:

Main Thm (Formal):
For a tuple of (indep) rand. variable $\mathbf{X} = (\mathbf{X}_1, \ldots, \mathbf{X}_q)$ satisfying

$$H^*_\infty(\mathbf{X}) > \beta,$$

our scheme E[α, β] satisfies

$$(\mathbf{X}, M^\alpha, q)\text{-IND}$$

for any $0 < \alpha < \beta$.

Here M is Mess. Sp.Size.

Our scheme is based on a PRF and the above result holds under security of PRF.

Corollaries

- Recall that our ($\mathbf{X}$,M^{α},q)-IND can hide lower bits of a plaintext
- Hence, the following corollaries hold (under the same assumption as above).

Corollary: Our scheme $E[\alpha, \beta]$ can hide fraction α of lower bits of plaintexts for any $\alpha < \beta$ satisfying $\beta < H^*_{\infty}(\mathbf{X})$.

In particular, if $\mathbf{X}$ is a tuple of the **Unif** distributions, it follows that

Corollary: Our scheme $E[\alpha, 1]$ can hide any fraction of lower bits of plaintexts.

(r,q+1)-WOW of Our Scheme.

- We can show the following fact as well:

Theorem: ($\mathbf{Unif}^q$,T,q)-IND implies (r,q+1)-WOW for suitable r.

- In particular, we can conclude the following corollary:

Corollary: Our scheme satisfies (M^s,q+1)-WOW for any

$0<s<1$

- In the case of the known scheme [1], it is shown that
  - the known scheme is (1,q+1)-WOW
  - but it is *not* (M^s,q+1)-WOW for $s > 1/2$.

Hence, our scheme achieve (r,q+1)-WOW for better parameter r than the known scheme [1].

- Our Definition of Indistiguishability Notion
- Our Results
- Construction of Our scheme
- Security Proof

Construction (1/4)

We construct our scheme in the following two steps:

- First, we construct a scheme
  - which satisfies our ($\boldsymbol{\chi}$,M^{α},q)-IND without assuming any computational assumption.
  - But the enc. and dec. of this scheme requires super-polytime

→ Today we talk about this scheme

- Second, we improve the above scheme
  - Here we use the "lazy sampling" technique [2],
  - So we use a PRF
  - and the security of this scheme is based on PRF.
  - The scheme achieves poly-time enc. and dec. costs.

→ See our paper for this scheme

Construction (2/4)

- For an encryption function Enc_K, we let

$$R := \mathrm{Enc}_K(0)$$

$$D[i] := \mathrm{Enc}_K(i) - \mathrm{Enc}_K(i-1)$$

- Then we can write $\mathrm{Enc}_K(m)$ as follows:

$$\mathrm{Enc}_K(m) = R + \sum_{i=1}^{m} D[i].$$

- Therefore, a design of Enc_K can be reduced to the selections of R and D[i].

Construction (3/4)

How to select D[i]:

we set D[i] $\leftarrow$ small value with high probability,
but set it to a "large random value" with low probability.

- Specifically,
  - Let p be a "small" fixed value.
  - Take a coin r[i] which becomes 1 with high prob 1-p.
  - if (r[i] = 1)
    - D[i] $\leftarrow$ small value (say, 1).
  - Otherwise,
    - D[i] $\xleftarrow{\$}$ {1,...,L},

      where L = large value (say, $2^{\mathrm{poly(SecParam)}}$)

- We take a value R in a similar manner

Construction (4/4)

- Then we set

  Key K $\leftarrow$ (R,D[1],...,D[M]), (Here Mess.Sp={0,...,M})

  $$\mathrm{Enc}_K(m) \leftarrow R + \Sigma_{i=1}^{m} D[i].$$

- But the problems are that,

  when the Mess. Sp. size M is super-polynomial of SecParam,

  - the above key K is ***not*** polysize
  - the above Enc_K is ***not*** polytime

- So, finally, we improve the above scheme using “lazy sampling” technique [1].
  - We omit the explanation of this final part. See our paper.

- Our Definition of Indistiguishability Notion
- Our Results
- Construction of Our scheme
- Security Proof

($\boldsymbol{\chi}$,M^{α},q)-IND of Our Scheme

Proof)

Consider the Mess. Sp. ={1…M}

- Due to the def. of ($\boldsymbol{\chi}$,M^{α},q)-IND , messages m*[0] and m*[1] of the challenge have to be within the distance T=M^{α}.
- Since α<1, the distance T=M^{α} is small compare to M (when M→ ∞)
- Recall that we consider the case where components of $\boldsymbol{\chi}$ has high min-

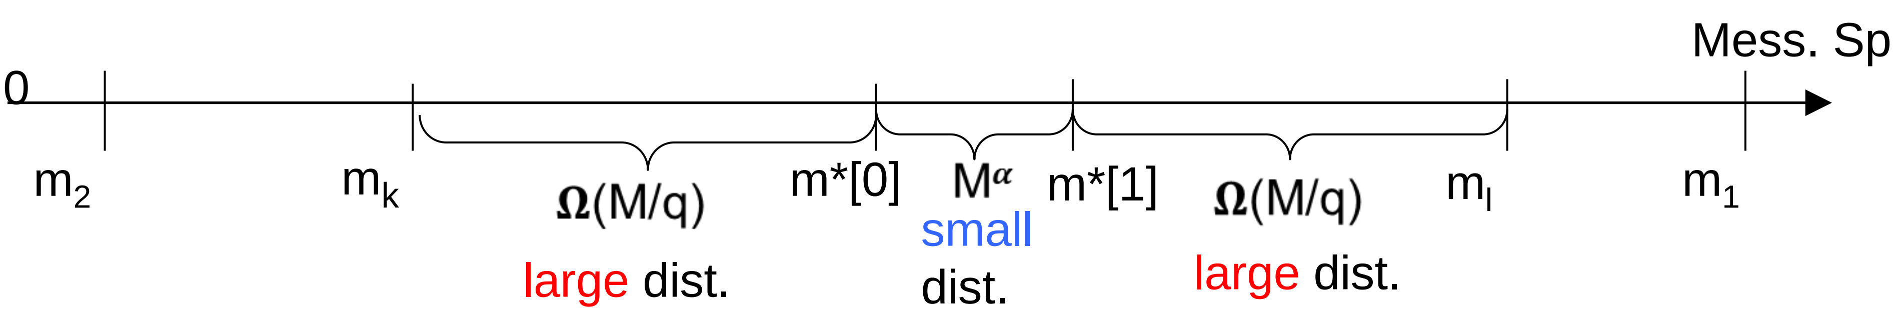

$(\mathbf{X}, M^{\alpha}, q)$-IND of Our Scheme

Recall that we take D[i] as follows:

- with high probability D[i] ← small value.
- with small probability D[i] becomes large random value.

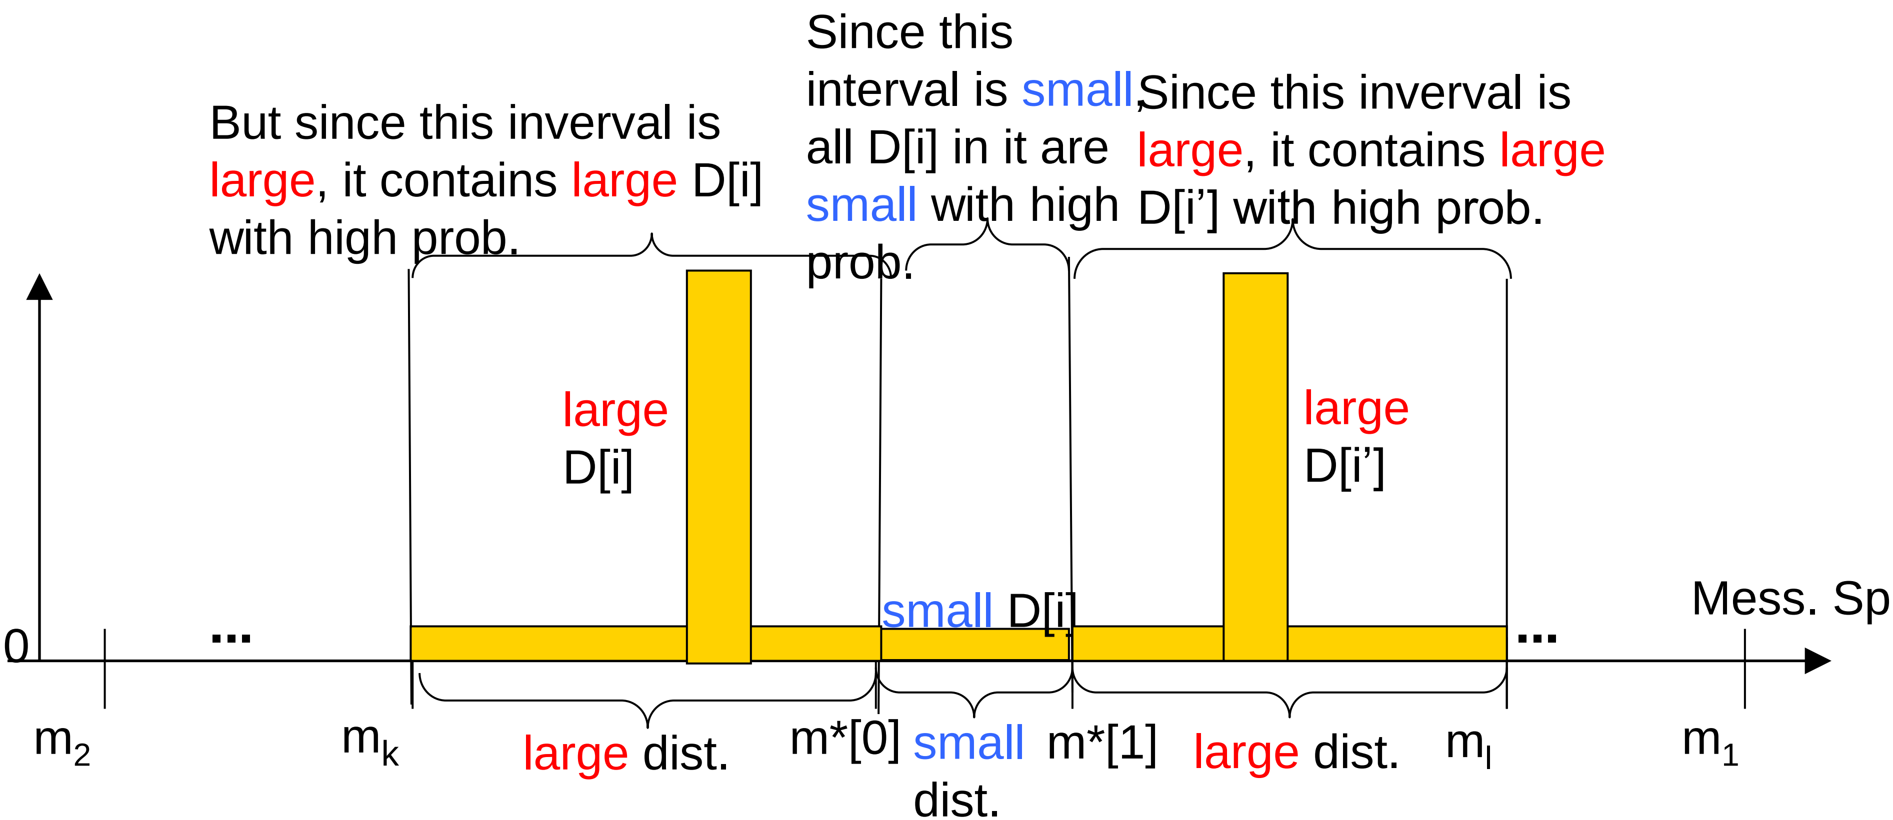

($\mathbf{X}$,M^{α},q)-IND of Our Scheme

- Consider an adversary who want to know b from

$$\mathrm{Enc}_K(m^*[b]) - \mathrm{Enc}_K(m_k) \qquad (\text{for } m_k < m^*[0])$$

$$= \Sigma_{j=m_k}^{m^*[b]} D[j] \qquad (\text{by definition.})$$

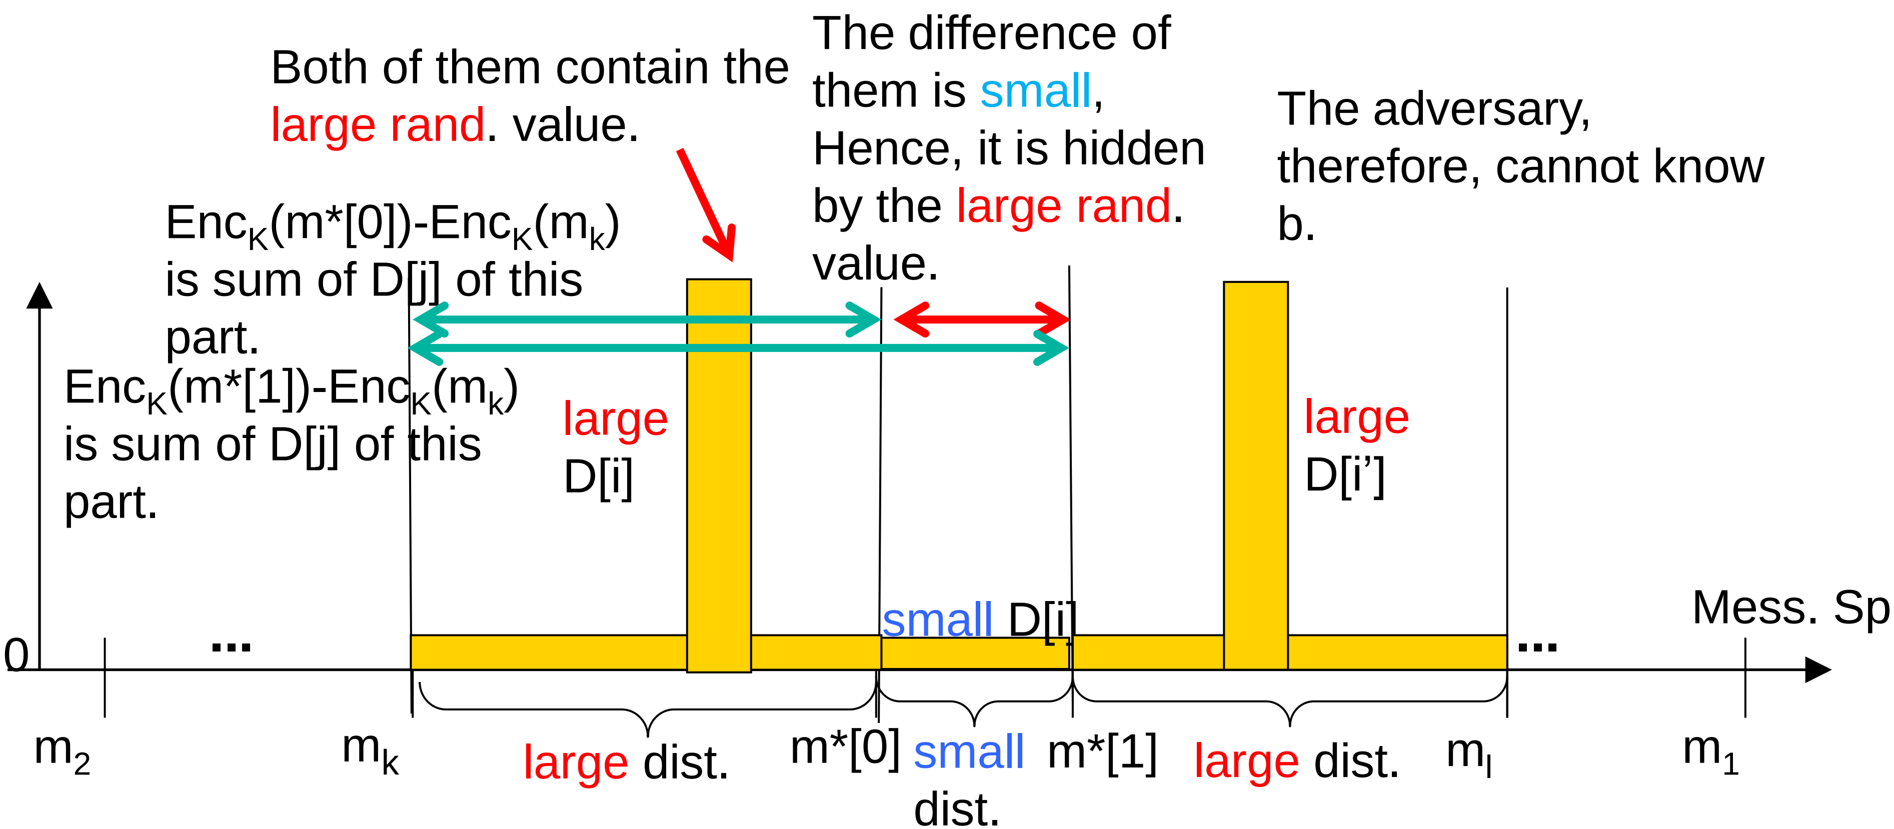

$(\mathbf{X}, M^{\alpha}, q)$-IND of Our Scheme

- Similarly, even if an adversary tries to know b from

  $Enc_K(m_l)$ - $Enc_K(m^*[b])$ (for $m_l > m^*[1]$),

  he cannot know it due to a similar reason.

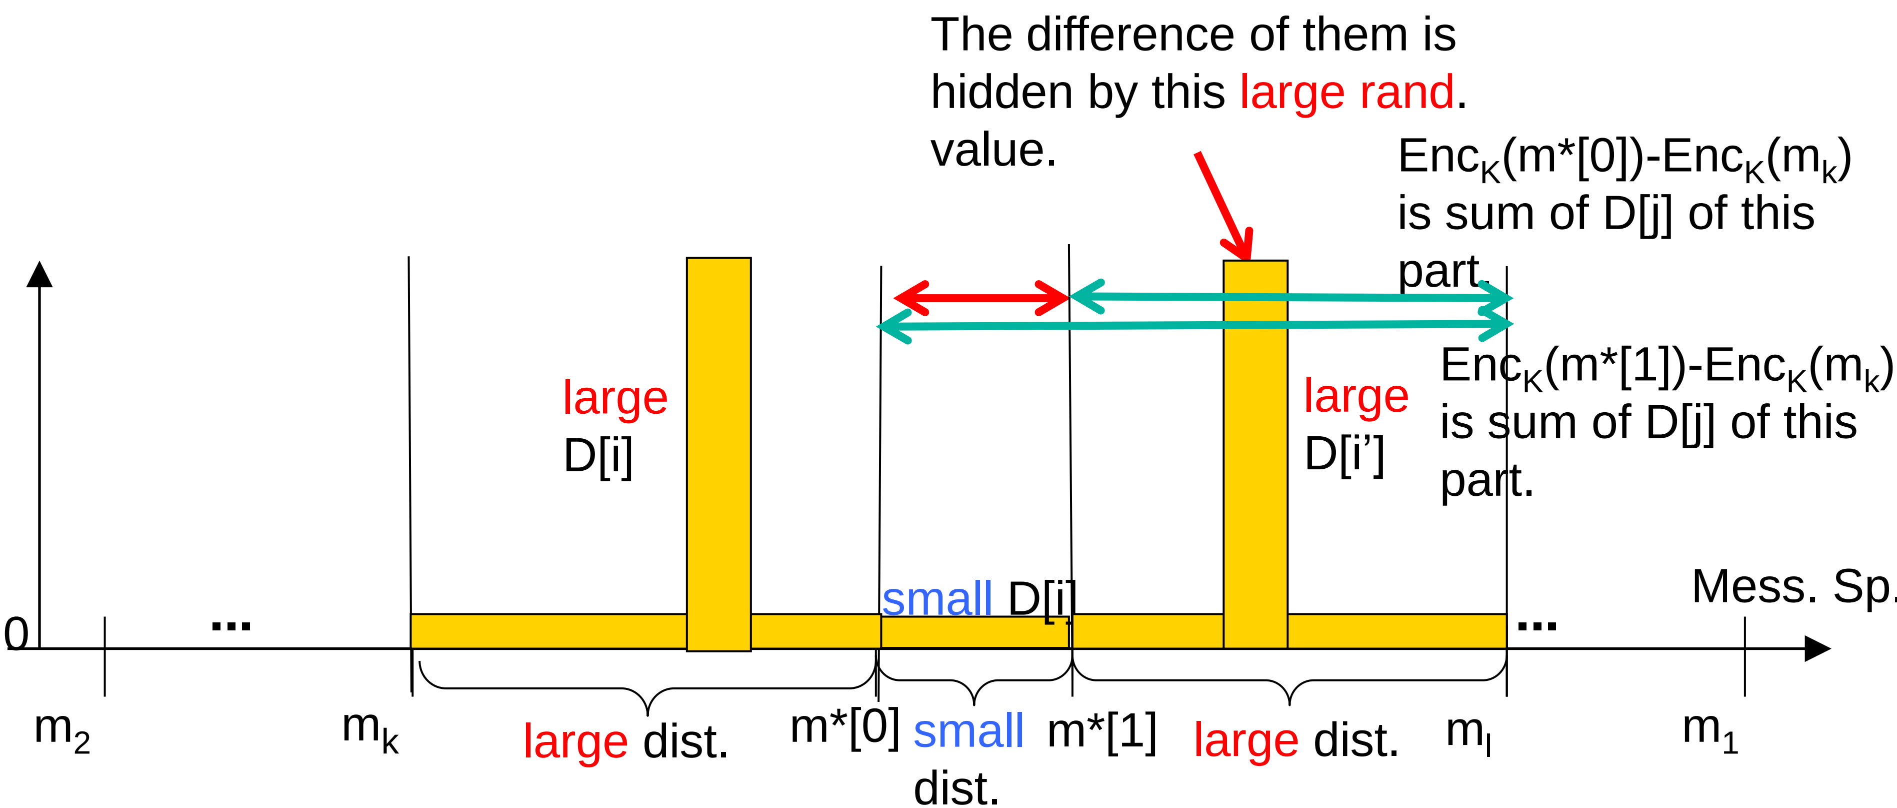

Conclusion

- OPE is very powerful for encrypted database
- but so far, security for it is poorly understood beyond just onewayness the encryption
- We proposed a new indistinguishability notion for OPE.
- This notion can ensure secrecy of lower bits of a plaintext.
- We construct a new OPE scheme which satisfies our new ind. notion.
- In some application hidden lower bits is significant security property like physical measurement, may be trade secret.
- Many question are remaining open.

Thank you