Practical Order-Revealing Encryption with Limited Leakage

Nathan Chenette, Kevin Lewi, Stephen A. Weis, David J. Wu

https://crypto.stanford.edu/ore/

Order-Revealing Encryption

An order-revealing encryption (ORE) scheme (introduced by Boneh et al.) is a secret-key encryption scheme that allows anyone to determine the ordering of the ciphertexts.

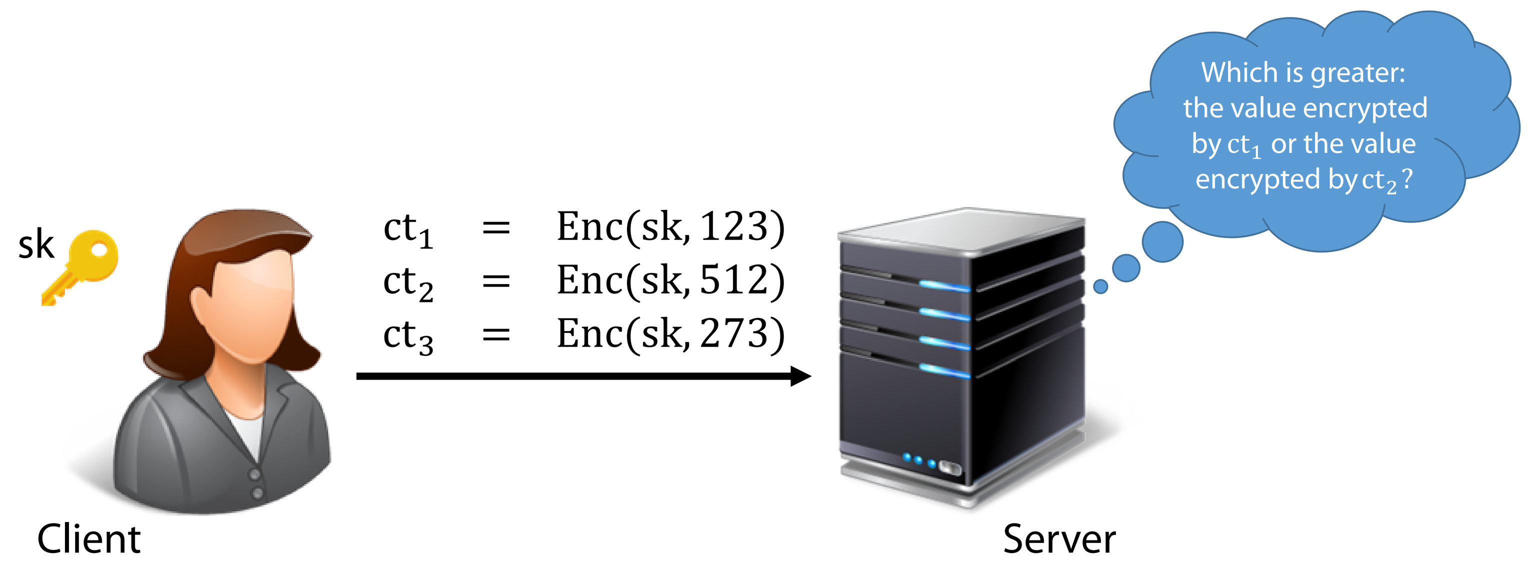

Since comparisons can be performed directly on ciphertexts, order-revealing encryption is useful for sorting and searching over encrypted data.

A closely related notion introduced by Boldyreva et al. is order-preserving encryption (OPE), which has the additional restriction that ciphertexts are numeric and the comparison operation is implemented by numeric comparison of the ciphertexts. In other words,

$$x > y \Leftrightarrow \mathrm{Enc}(\mathrm{sk}, x) > \mathrm{Enc}(\mathrm{sk}, y)$$

In contrast, in an order-revealing encryption scheme, the comparison function can be an *arbitrary* function of the ciphertexts. Thus, order-preserving encryption schemes are a special case of order-revealing encryption.

Defining Security

Best-possible (IND-OCPA) security:

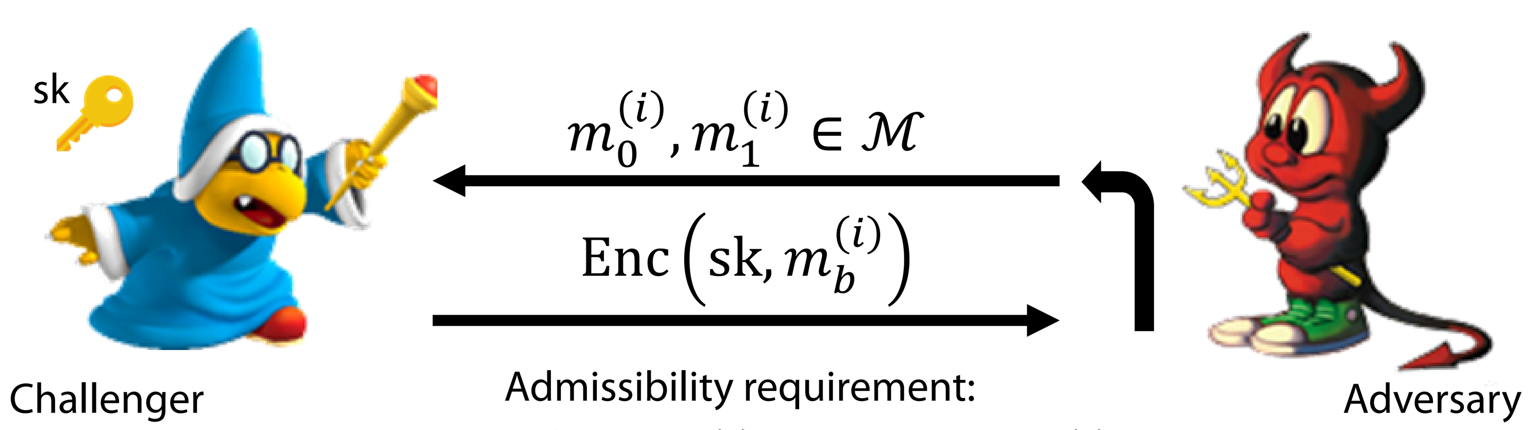

Admissibility requirement:

$$\forall i, j: m_0^{(i)} < m_0^{(j)} \Leftrightarrow m_1^{(i)} < m_1^{(j)}$$

This definition captures the notion that the adversary learns "nothing but the ordering." However, this is a very strong notion of security and seemingly difficult to achieve. In fact, Boldyreva et al. showed a lower bound that no order-preserving encryption scheme can satisfy best-possible security unless the size of the ciphertext space is **exponential** in the size of the plaintext space. This lower bound does not extend to order-revealing encryption, and there do exist candidate constructions of ORE that achieve best-possible security based on indistinguishability obfuscation (iO) or multilinear maps. Unfortunately, these schemes are far from practical.

Since OPE schemes cannot satisfy best-possible security, Boldyreva et al. introduced an alternative notion of security for OPE schemes that compares the outputs of the OPE encryption algorithm to that of a truly random order-preserving function. An OPE scheme is ROPF-CCA secure if no efficient adversary can distinguish encryptions (of messages of the adversary's choosing) from the outputs of a truly random order-preserving encryption evaluated on the adversary's choice of messages.

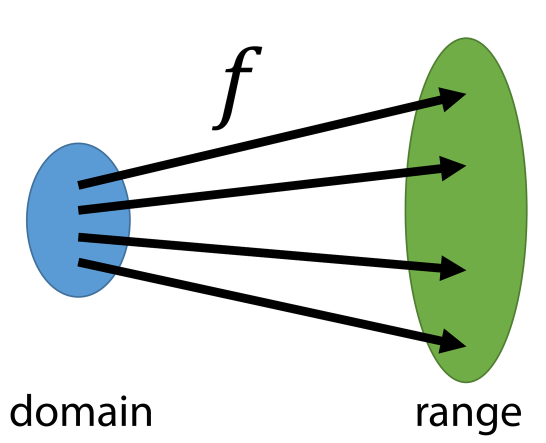

Properties of a truly random order-preserving function:

- Given $f(x)$, can deduce half of the most significant bits of x.
- Given $f(x)$ and $f(y)$, can deduce half of the most significant bits of the distance between x and y.
- No semantic security for even a single message.

The Landscape of OPE/ORE

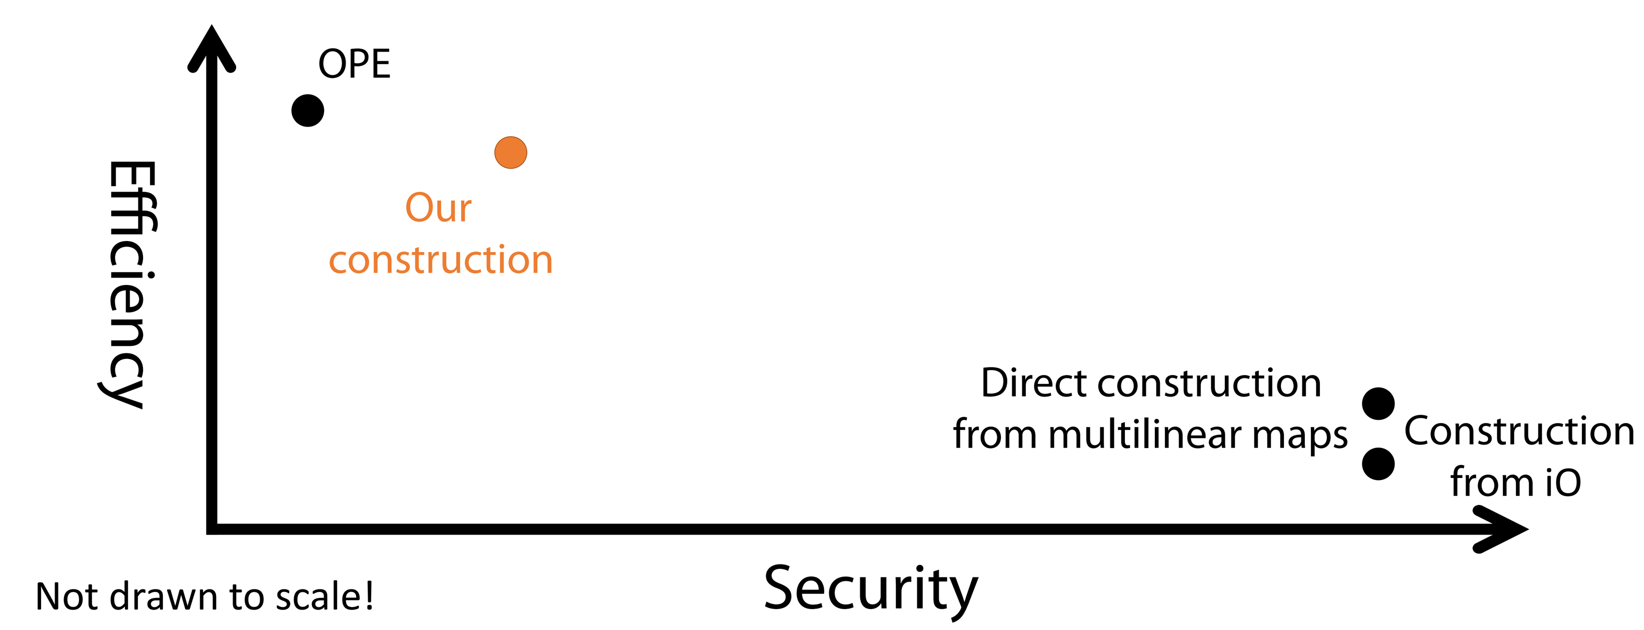

A New Security Notion

Existing practical constructions of ORE (and OPE) do not satisfy best-possible security and leak information about the underlying messages. While some of these schemes can be shown to be secure under some security notion (e.g., ROPF-CCA), these security notions do not give a simple characterization of the leakage of the underlying encryption scheme (without relying on strong assumptions on the message distribution).

We introduce a new simulation-based notion of security with respect to a leakage function to obtain a notion that explicitly specifies the information leakage of the encryption scheme.

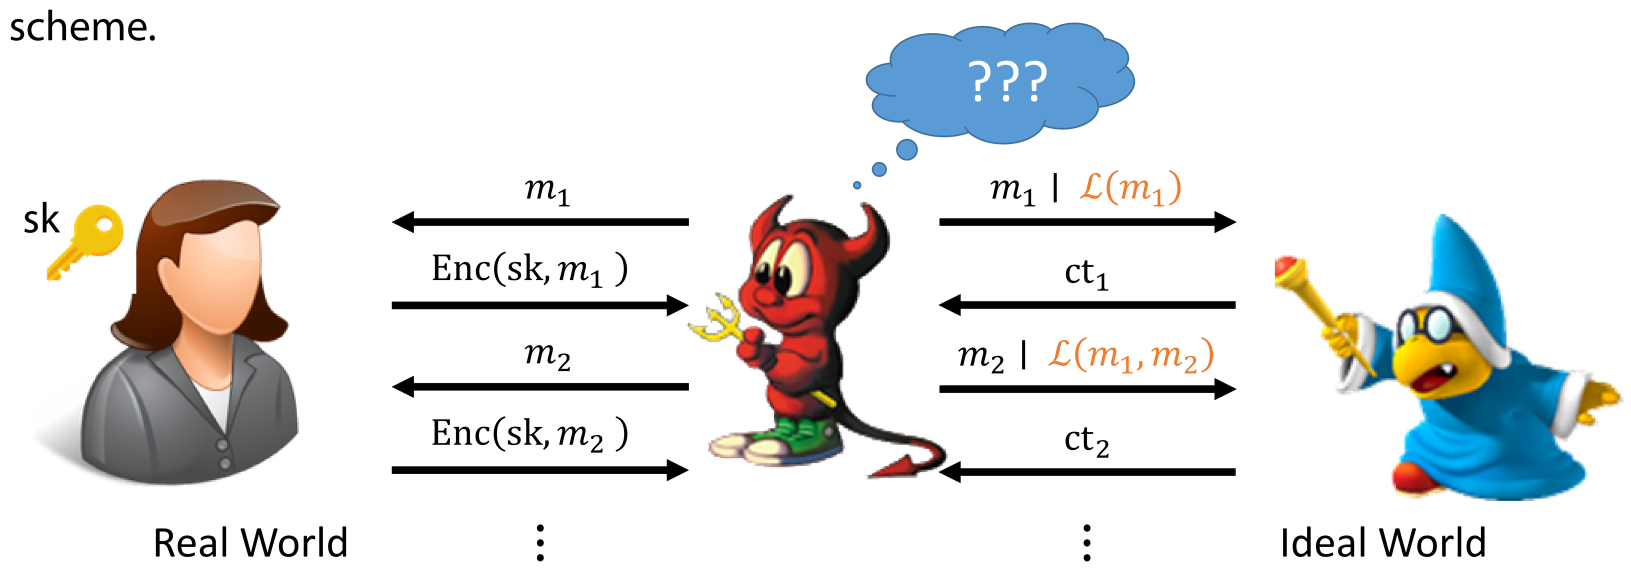

$\mathcal{L}(m_1, \dots, m_q)$: Leakage function on messages $m_1, \dots, m_q$

Best-possible security (nothing is leaked except the ordering):

$$\mathcal{L}(m_1, \dots, m_q) = \{(i, j, \mathbf{1}\{m_i < m_j\}) \mid 1 \le i < j \le q\}$$

Definition states that whatever can be inferred from the ciphertexts can be inferred from the leakage function alone (i.e., the ciphertexts can be *simulated* given just the leakage function evaluated on the messages).

Our Leakage Function

We consider a leakage function that leaks a little more than just the ordering of the messages. This will enable a very efficient construction from psuedorandom functions (PRFs) alone.

Our leakage function:

$$\mathcal{L}(m_1, \dots, m_q) = \left\{\left(i, j, \mathbf{1}\{m_i < m_j\}, \mathrm{ind}_{\mathrm{diff}}(m_i, m_j)\right) \mid 1 \le i < j \le q\right\}$$

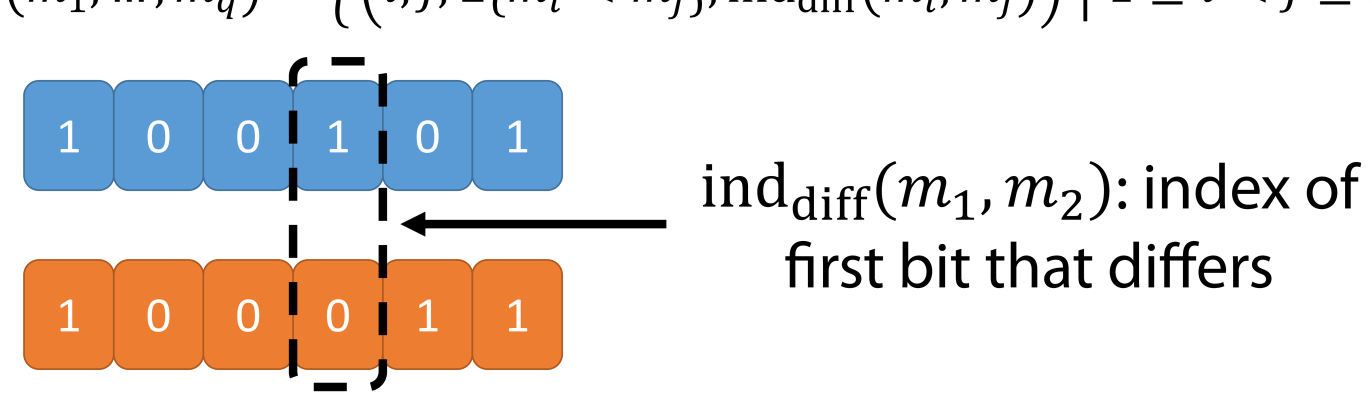

Our leakage function reveals some partial information about the distances between messages.

Our Construction

Basic idea: for each index i, apply a PRF to the first $i-1$ bits, then add the i^{th} bit (mod 3)

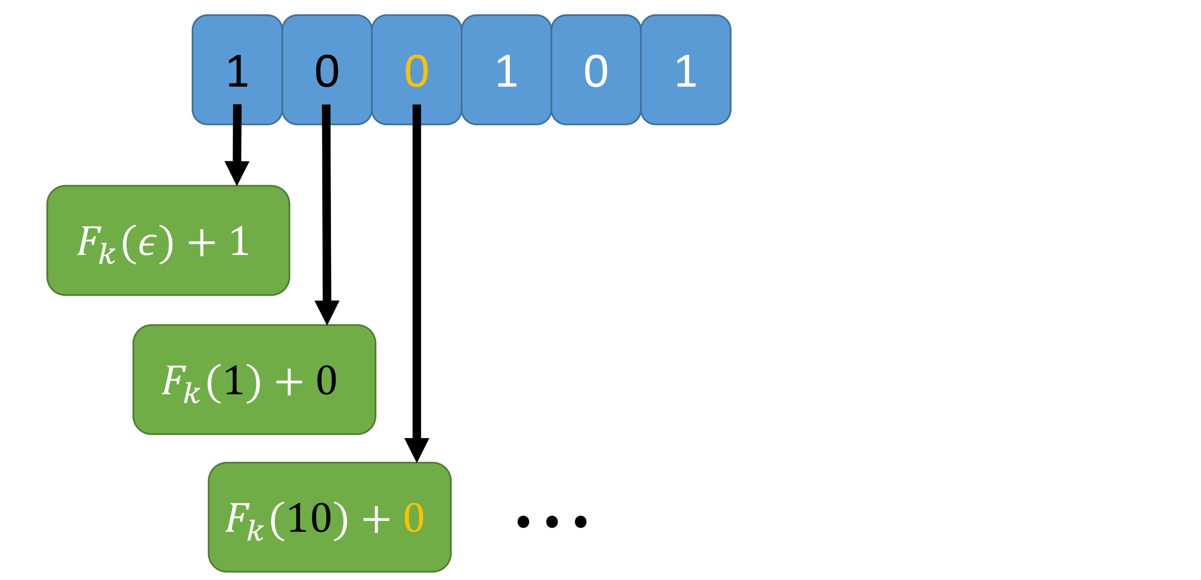

To compare two ciphertexts, find the first block where they differ. The precise ordering can be determined by comparing the values (mod 3).

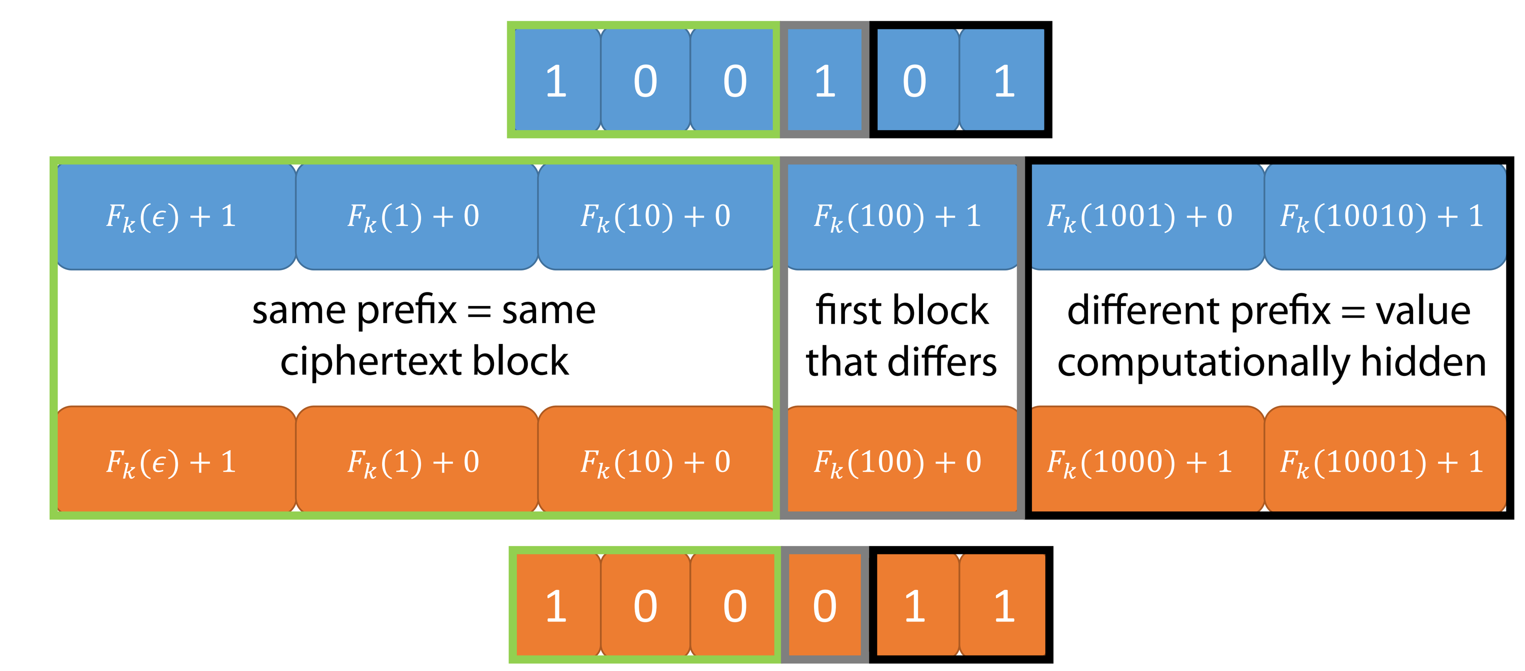

Properties of our scheme:

- Each ciphertext block is an element in $\mathbb{Z}_3$, so for an n-bit message, ciphertexts are approximately $1.6n$ bits long.
- Encryption only requires PRF evaluations while decryption just requires bitwise comparisons.
- Security reduces directly to PRF security.
- Can convert to an OPE scheme by increasing the ciphertext block size.
- Possible to compose OPE with ORE to achieve security at least as strong as the underlying OPE encryption.

Evaluation and Conclusions

One evaluation metric for ORE/OPE is window one-wayness security.

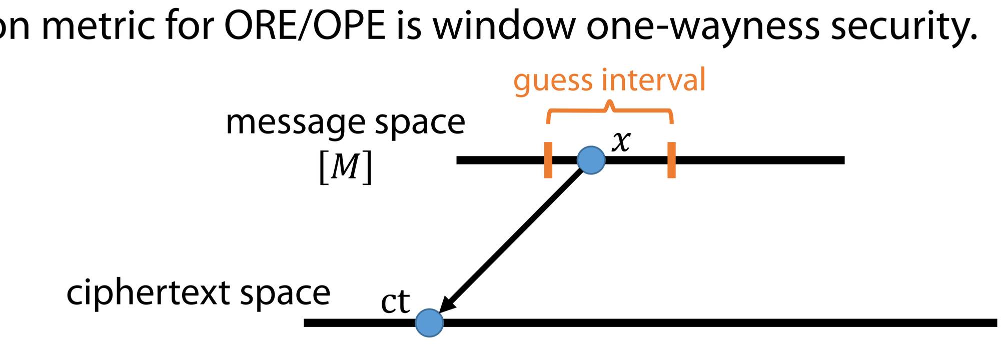

Theorem (Informal) [Boldyreva et al.]: For an ROPF, if the size of the guess interval $r = O(\sqrt{M})$, then there is an efficient adversary whose window one-wayness advantage is close to 1.

Each ciphertext alone reveals half of the most significant bits of the plaintext!

Theorem (Informal). For our OPE scheme, if the size of the guess interval $r = M^{1-\epsilon}$ for any constant $\epsilon > 0$, then for all efficient adversaries, their (generalized) window one-wayness advantage is negligible.

No constant fraction ϵ of the bits of the plaintexts are revealed.

In this work, we introduced a new notion of security for order-revealing encryption that allows for a precise quantification of the leakage of the scheme. We then gave a new and very practical ORE construction from PRFs that leaks slightly more than just the ordering between messages.