A Survey of Recent Results on Key-Alternating Ciphers

Yannick Seurin
(based on joint work with
R. Lampe and J. Patarin)

ANSSI

Mathcrypt 2013 — July 5, 2013

Introduction

A key-alternating cipher with r rounds is the following construction:

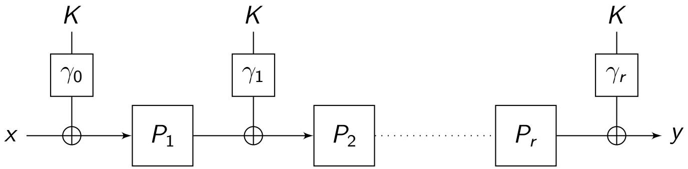

- The P_i's are public permutations on $\{0,1\}^n$
- $K \in \{0,1\}^\ell$ is the (master) key
- The γ_i's are key derivation functions mapping K to n-bit values

Also named Iterated Even-Mansour (IEM) cipher

Introduction

A key-alternating cipher with r rounds is the following construction:

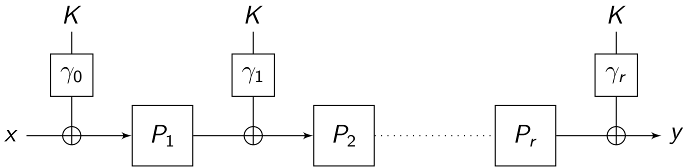

- The P_i's are public permutations on $\{0,1\}^n$
- $K \in \{0,1\}^\ell$ is the (master) key
- The γ_i's are key derivation functions mapping K to n-bit values

Also named Iterated Even-Mansour (IEM) cipher

Introduction

Most (if not all) SPN ciphers can be described as key-alternating ciphers. E.g. for AES-128, one has $r = 10$, the γ_i's are efficiently invertible permutations, and:

$$P_1 = \ldots = P_9 = \texttt{SubBytes} \circ \texttt{ShiftRows} \circ \texttt{MixColumns}$$
$$P_{10} = \texttt{SubBytes} \circ \texttt{ShiftRows}$$

When the P_i's are fixed permutations, one can prove results like:

- the best differential characteristic over $r' < r$ rounds has probability at most p
- the best linear approximation over $r' < r$ rounds has probability at most p'

This gives upper bounds on the success probability of very specific adversaries

Introduction

Most (if not all) SPN ciphers can be described as key-alternating ciphers. E.g. for AES-128, one has $r = 10$, the γ_i's are efficiently invertible permutations, and:

$$P_1 = \ldots = P_9 = \texttt{SubBytes} \circ \texttt{ShiftRows} \circ \texttt{MixColumns}$$
$$P_{10} = \texttt{SubBytes} \circ \texttt{ShiftRows}$$

When the P_i's are fixed permutations, one can prove results like:

- the best differential characteristic over $r' < r$ rounds has probability at most p
- the best linear approximation over $r' < r$ rounds has probability at most p'

This gives upper bounds on the success probability of very specific adversaries

Introduction

Recently, a lot of results have been obtained in the Random Permutation Model: the P_i's are viewed as oracles to which the adversary can make black-box queries (both to P_i and P_i^{-1}).

Interpretation: gives a guarantee against any adversary which do not use particular properties of the P_i's

In fact, this model was already considered 15 years ago by Even and Mansour for $r = 1$ round: they showed that the following cipher is secure up to $\mathcal{O}(2^{n/2})$ queries of the adversary:

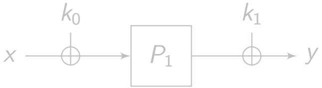

Introduction

Recently, a lot of results have been obtained in the Random Permutation Model: the P_i's are viewed as oracles to which the adversary can make black-box queries (both to P_i and P_i^{-1}).

Interpretation: gives a guarantee against any adversary which do not use particular properties of the P_i's

In fact, this model was already considered 15 years ago by Even and Mansour for $r = 1$ round: they showed that the following cipher is secure up to $\mathcal{O}(2^{n/2})$ queries of the adversary:

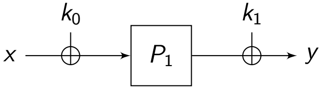

Outline

Outline

1. Indistinguishability
   - Introduction
   - The coupling technique
   - The indistinguishability proof
2. Interlude: tweakable block ciphers
3. Indifferentiability
   - Introduction
   - Indifferentiability of the IEM cipher
   - At least 4 rounds are necessary
   - Indifferentiability proof for 12 rounds

Outline

The IEM cipher with independent keys

We focus in this part on the IEM cipher with independent round keys:

$$K = (k_0, k_1, \ldots, k_r)$$

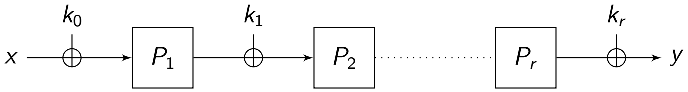

Total key space: $\{0,1\}^{(r+1)n}$

Notation:

$$y = \mathtt{EM}^{P_1,\ldots,P_r}_{(k_0,\ldots,k_r)}(x) \ .$$

Formalizing indistinguishability for the IEM cipher

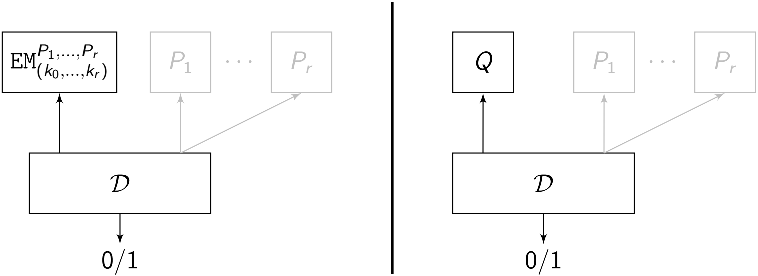

- left: $k_0, \ldots, k_r \leftarrow_\$ \{0,1\}^n$ are randomly chosen keys
- right: Q is a random permutation independent of $P_1, \ldots, P_r$
- we are in the Random Permutation Model: the distinguisher also has oracle access to $P_1, \ldots, P_r$ in both worlds

Formalizing indistinguishability for the IEM cipher

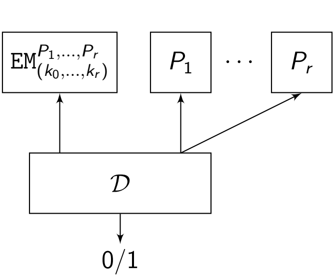

- left: $k_0, \ldots, k_r \leftarrow_{\$} \{0,1\}^n$ are randomly chosen keys
- right: Q is a random permutation independent of $P_1, \ldots, P_r$
- we are in the Random Permutation Model: the distinguisher also has oracle access to $P_1, \ldots, P_r$ in both worlds

Indistinguishability of the IEM cipher: Summary of results

Results for independent round keys $(k_0, k_1, \ldots, k_r)$
Notation: $N = 2^n$

- for $r = 1$ round, EM is secure up to $\mathcal{O}(N^{1/2})$ queries [EM97]
- for $r \geq 2$, EM is secure up to $\mathcal{O}(N^{2/3})$ queries [BKL$^+$12]
- for any even r, EM is secure up to $\mathcal{O}(N^{r/(r+2)})$ queries [LPS12]
- tight result: EM is secure up to $\mathcal{O}(N^{r/(r+1)})$ queries [CS14]

In the following, we focus on the [LPS12] result which uses the coupling technique.

Indistinguishability of the IEM cipher: Summary of results

Results for independent round keys $(k_0, k_1, \ldots, k_r)$
Notation: $N = 2^n$

- for $r = 1$ round, EM is secure up to $\mathcal{O}(N^{1/2})$ queries [EM97]
- for $r \geq 2$, EM is secure up to $\mathcal{O}(N^{2/3})$ queries [BKL$^+$12]
- for any even r, EM is secure up to $\mathcal{O}(N^{r/(r+2)})$ queries [LPS12]
- tight result: EM is secure up to $\mathcal{O}(N^{r/(r+1)})$ queries [CS14]

In the following, we focus on the [LPS12] result which uses the coupling technique.

Outline

Coupling: definition

Definition (Coupling)

Let μ and ν be two probability distributions on Ω. A coupling of μ and ν is a probability dist. λ on $\Omega \times \Omega$ such that:

$$\forall x \in \Omega, \sum_{y \in \Omega} \lambda(x, y) = \mu(x)$$

$$\forall y \in \Omega, \sum_{x \in \Omega} \lambda(x, y) = \nu(y)$$

In other words, λ is a joint probability distribution whose marginal distributions are resp. μ and ν.

Definition (Statistical distance)

$\|\mu - \nu\| = \frac{1}{2} \sum_{x \in \Omega} |\mu(x) - \nu(x)|$.

Coupling: definition

Definition (Coupling)

Let μ and ν be two probability distributions on Ω. A coupling of μ and ν is a probability dist. λ on $\Omega \times \Omega$ such that:

$$\forall x \in \Omega, \sum_{y \in \Omega} \lambda(x, y) = \mu(x)$$

$$\forall y \in \Omega, \sum_{x \in \Omega} \lambda(x, y) = \nu(y)$$

In other words, λ is a joint probability distribution whose marginal distributions are resp. μ and ν.

Definition (Statistical distance)

$\|\mu - \nu\| = \frac{1}{2} \sum_{x \in \Omega} |\mu(x) - \nu(x)|$.

The coupling lemma

Lemma

Let μ and ν be two probability distributions and λ be a coupling. Let $(X, Y) \sim \lambda$. Then:

$$\|\mu - \nu\| \leq \Pr[X \neq Y] \ .$$

- Introduced by Aldous, key tool to study the mixing time of Markov chains
- First used in crypto by Mironov [Mir02] to analyze the shuffle of RC4, later by [MRS09, HR10] to analyze Feistel ciphers

The coupling lemma

Lemma

Let μ and ν be two probability distributions and λ be a coupling. Let $(X, Y) \sim \lambda$. Then:

$$\|\mu - \nu\| \leq \Pr[X \neq Y] \ .$$

- Introduced by Aldous, key tool to study the mixing time of Markov chains
- First used in crypto by Mironov [Mir02] to analyze the shuffle of RC4, later by [MRS09, HR10] to analyze Feistel ciphers

A (very) simple example

Two couplings of the uniform distribution on $\{1, 2, 3, 4\}$ with itself:

X/Y	1	2	3	4
1	$1/16$	$1/16$	$1/16$	$1/16$
2	$1/16$	$1/16$	$1/16$	$1/16$
3	$1/16$	$1/16$	$1/16$	$1/16$
4	$1/16$	$1/16$	$1/16$	$1/16$

$\Pr[X \neq Y] = 3/4$

X/Y	1	2	3	4
1	$1/4$	0	0	0
2	0	$1/4$	0	0
3	0	0	$1/4$	0
4	0	0	0	$1/4$

$\Pr[X \neq Y] = 0$

Not all couplings give good upper bounds on $\|\mu - \nu\|$

NB: there always exists a coupling λ for which equality

$$\|\mu - \nu\| = \Pr[X \neq Y]$$

is achieved (but it may be hard to describe when μ and ν are not efficiently computable)

A (very) simple example

Two couplings of the uniform distribution on $\{1, 2, 3, 4\}$ with itself:

X/Y	1	2	3	4
1	1/16	1/16	1/16	1/16
2	1/16	1/16	1/16	1/16
3	1/16	1/16	1/16	1/16
4	1/16	1/16	1/16	1/16

$\Pr[X \neq Y] = 3/4$

X/Y	1	2	3	4
1	1/4	0	0	0
2	0	1/4	0	0
3	0	0	1/4	0
4	0	0	0	1/4

$\Pr[X \neq Y] = 0$

Not all couplings give good upper bounds on $\|\mu - \nu\|$

NB: there always exists a coupling λ for which equality

$$\|\mu - \nu\| = \Pr[X \neq Y]$$

is achieved (but it may be hard to describe when μ and ν are not efficiently computable)

A simple example

Two coins:

- a perfect one: $p_{\text{head}} = 0.5$
- a biased one: $p'_{\text{head}} = 0.6$

Show that over N tosses, the probability that the biased coin makes k heads is larger than the probability that the perfect coin makes k heads (for any $k \leq N$). Two solutions:

1. compute the binomial law: a bit tedious...
2. couple the two distributions as follows:
   - toss the perfect coin
   - if the perfect coin makes head, the biased coin makes head
   - if the perfect coin makes tail, the biased coin makes head with proba 0.2

$\Rightarrow$ the marginal distributions are correct (simple)

$\Rightarrow$ for any k, the biased coin makes k heads with larger probability than the perfect coin (trivial)

A simple example

Two coins:

- a perfect one: $p_{\mathrm{head}} = 0.5$
- a biased one: $p'_{\mathrm{head}} = 0.6$

Show that over N tosses, the probability that the biased coin makes k heads is larger than the probability that the perfect coin makes k heads (for any $k \leq N$). Two solutions:

1. compute the binomial law: a bit tedious...
2. couple the two distributions as follows:
   - toss the perfect coin
   - if the perfect coin makes head, the biased coin makes head
   - if the perfect coin makes tail, the biased coin makes head with proba 0.2

   $\Rightarrow$ the marginal distributions are correct (simple)

   $\Rightarrow$ for any k, the biased coin makes k heads with larger probability than the perfect coin (trivial)

A simple example

Two coins:

- a perfect one: $p_{\text{head}} = 0.5$
- a biased one: $p'_{\text{head}} = 0.6$

Show that over N tosses, the probability that the biased coin makes k heads is larger than the probability that the perfect coin makes k heads (for any $k \leq N$). Two solutions:

1. compute the binomial law: a bit tedious...
2. couple the two distributions as follows:
   - toss the perfect coin
   - if the perfect coin makes head, the biased coin makes head
   - if the perfect coin makes tail, the biased coin makes head with proba 0.2

$\Rightarrow$ the marginal distributions are correct (simple)

$\Rightarrow$ for any k, the biased coin makes k heads with larger probability than the perfect coin (trivial)

A simple example

Two coins:

- a perfect one: $p_{\text{head}} = 0.5$
- a biased one: $p'_{\text{head}} = 0.6$

Show that over N tosses, the probability that the biased coin makes k heads is larger than the probability that the perfect coin makes k heads (for any $k \leq N$). Two solutions:

1. compute the binomial law: a bit tedious...
2. couple the two distributions as follows:
    - toss the perfect coin
    - if the perfect coin makes head, the biased coin makes head
    - if the perfect coin makes tail, the biased coin makes head with proba 0.2

$\Rightarrow$ the marginal distributions are correct (simple)

$\Rightarrow$ for any k, the biased coin makes k heads with larger probability than the perfect coin (trivial)

Outline

Two types of distinguishers

NB: $\mathcal{D}$ is computationally unbounded and makes at most q queries to each oracle

We define the two following classes of distinguishers:

- NCPA (Non-Adaptive Chosen Plaintext Attacks):
  $\rightarrow$ works in two phases:
  - $\mathcal{D}$ first queries $P_1, \ldots, P_r$ as it wishes (in both directions, adaptively);
  - then it makes q non-adaptive direct queries to $\mathrm{EM}^{P_1,\ldots,P_r}/Q$
- CCA (Chosen Ciphertext Attacks):
  $\rightarrow$ the most general class of distinguisher, can adaptively query all oracles in both directions, in any order

Two types of distinguishers

NB: $\mathcal{D}$ is computationally unbounded and makes at most q queries to each oracle

We define the two following classes of distinguishers:

- NCPA (Non-Adaptive Chosen Plaintext Attacks):
  $\rightarrow$ works in two phases:
  - $\mathcal{D}$ first queries $P_1, \ldots, P_r$ as it wishes (in both directions, adaptively);
  - then it makes q non-adaptive direct queries to $\mathtt{EM}^{P_1,\ldots,P_r}/Q$
- CCA (Chosen Ciphertext Attacks):
  $\rightarrow$ the most general class of distinguisher, can adaptively query all oracles in both directions, in any order

Two types of distinguishers

NB: $\mathcal{D}$ is computationally unbounded and makes at most q queries to each oracle

We define the two following classes of distinguishers:

- NCPA (Non-Adaptive Chosen Plaintext Attacks):
  $\rightarrow$ works in two phases:
  - $\mathcal{D}$ first queries $P_1, \ldots, P_r$ as it wishes (in both directions, adaptively);
  - then it makes q non-adaptive direct queries to $\mathtt{EM}^{P_1,\ldots,P_r}/Q$
- CCA (Chosen Ciphertext Attacks):
  $\rightarrow$ the most general class of distinguisher, can adaptively query all oracles in both directions, in any order

The case of NCPA distinguishers: the result

We will show the following:

Theorem

For any NCPA $\mathcal{D}$ making at most q queries to each oracle, the distinguishing advantage against the IEM with r rounds is at most

$$2^r \frac{q^{r+1}}{N^r} \; .$$

$\rightarrow$ security up to $\mathcal{O}(N^{r/(r+1)})$ queries.

The case of NCPA distinguishers: a matching attack

$\rightarrow$ security up to $\mathcal{O}(N^{r/(r+1)})$ queries.

A matching attack has been described in [BKL$^+$12]:

- make $\mathcal{O}(N^{r/(r+1)})$ queries to the cipher and to each P_i
- for each possible key, find a "contradictory path"
- any wrong key will have a contradictory path with high proba.
- (note: this is just exhaustive key search, but we are interested in the number of queries rather than computational cost)

The case of NCPA distinguishers

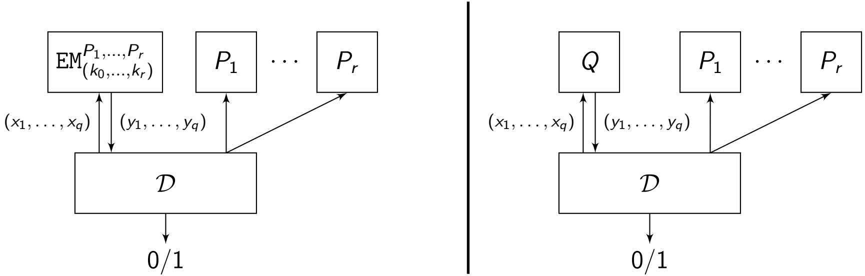

$\mathcal{D}$ first makes q queries to P_1, ..., P_r and obtains equations:

$$P_i(a_{i,j}) = b_{i,j},\ i \in [1,r],\ j \in [1,q]\ .$$

Then it makes q non-adaptive queries $(x_1,\ldots,x_q)$ to $\mathtt{EM}/Q$ and receives answers $(y_1,\ldots,y_q)$

The case of NCPA distinguishers

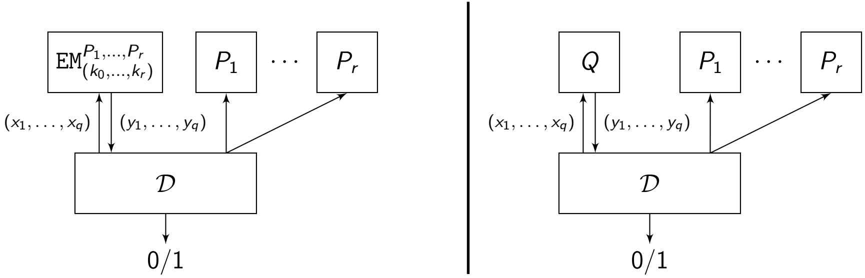

The distribution of $(a_{i,j})$, $(b_{i,j})$ is the same in both worlds
$\rightarrow$ the advantage of $\mathcal{D}$ is given by the statistical distance between the distributions of $(y_1,\dots,y_q)$ in the real and the ideal world

Notation:
μ_q = distribution of $(y_0,\dots,y_q)$ in the real world
μ_0 = distribution of $(y_0,\dots,y_q)$ in the ideal world (uniform)
$\rightarrow$ we want to upper bound $\|\mu_q - \mu_0\|$

The case of NCPA distinguishers

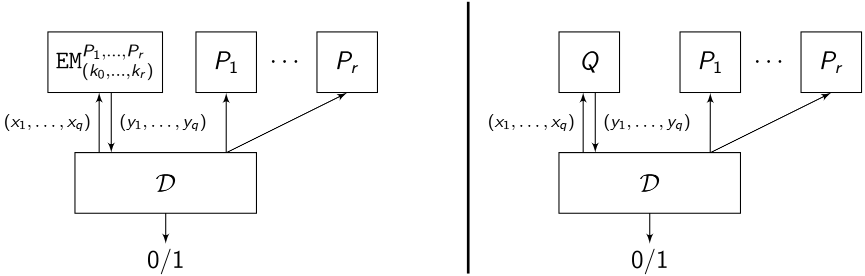

The distribution of $(a_{i,j})$, $(b_{i,j})$ is the same in both worlds
$\rightarrow$ the advantage of $\mathcal{D}$ is given by the statistical distance between the distributions of $(y_1,\ldots,y_q)$ in the real and the ideal world

Notation:
μ_q = distribution of $(y_0,\ldots,y_q)$ in the real world
μ_0 = distribution of $(y_0,\ldots,y_q)$ in the ideal world (uniform)
$\rightarrow$ we want to upper bound $\|\mu_q - \mu_0\|$

The case of NCPA distinguishers

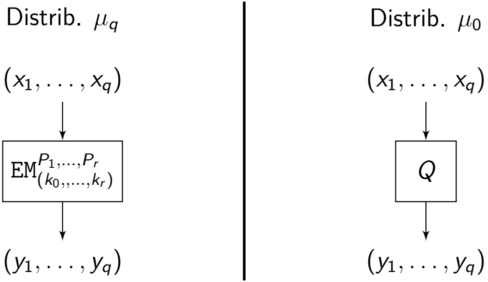

The distribution μ_q in the real world is obtained as follows:

- draw random permutations $P_1, \ldots, P_r$ satisfying $P_i(a_{i,j}) = b_{i,j}$
- draw independent random round keys $(k_0, \ldots, k_r)$
- let $y_i = \mathtt{EM}^{P_1,\ldots,P_r}_{(k_0,\ldots,k_r)}(x_i)$

A hybrid argument

The uniform distribution μ_0 is also obtained by drawing uniformly random (distinct) inputs $(u_1, \ldots, u_q)$ and computing their image through EM

A hybrid argument

The uniform distribution μ_0 is also obtained by drawing uniformly random (distinct) inputs $(u_1, \ldots, u_q)$ and computing their image through EM

A hybrid argument

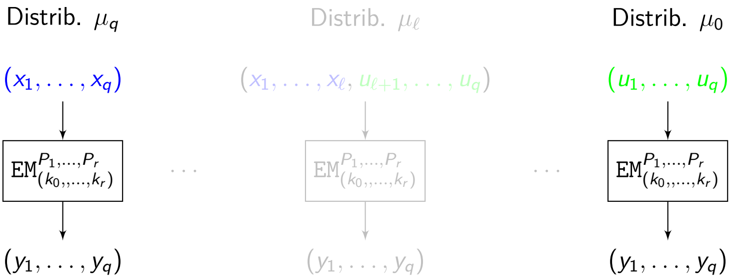

Hybrid distributions μ_ℓ, $\ell \in [0, q]$

$$\|\mu_q - \mu_0\| \leq \sum_{\ell=0}^{q-1} \|\mu_{\ell+1} - \mu_\ell\| \ .$$

$\rightarrow$ We will upper bound $\|\mu_{\ell+1} - \mu_\ell\|$ with a coupling.

A hybrid argument

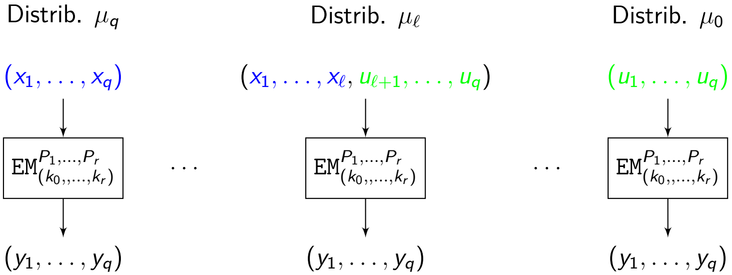

Hybrid distributions μ_ℓ, $\ell \in [0, q]$

$$\|\mu_q - \mu_0\| \leq \sum_{\ell=0}^{q-1} \|\mu_{\ell+1} - \mu_\ell\| \ .$$

$\rightarrow$ We will upper bound $\|\mu_{\ell+1} - \mu_\ell\|$ with a coupling.

Coupling $\mu_{\ell+1}$ and μ_ℓ

Distrib. $\mu_{\ell+1}$

$$(x_1, \dots, x_\ell, x_{\ell+1}, u_{\ell+2}, \dots, u_q)$$

$$\downarrow$$

$$\boxed{\mathtt{EM}^{P_1,\dots,P_r}_{(k_0,,\dots,k_r)}}$$

$$\downarrow$$

$$(y_1, \dots, y_\ell, y_{\ell+1}, y_{\ell+2}, \dots, y_q)$$

Distrib. μ_ℓ

$$(x_1, \dots, x_\ell, u_{\ell+1}, u_{\ell+2}, \dots, u_q)$$

$$\downarrow$$

$$\boxed{\mathtt{EM}^{P_1,\dots,P_r}_{(k_0,,\dots,k_r)}}$$

$$\downarrow$$

$$(y_1, \dots, y_\ell, y_{\ell+1}, y_{\ell+2}, \dots, y_q)$$

$(y_{\ell+2}, \dots, y_q)$ are distributed identically in both cases
$\rightarrow$ can be dropped

Coupling $\mu_{\ell+1}$ and μ_ℓ

$(y_{\ell+2}, \dots, y_q)$ are distributed identically in both cases
$\rightarrow$ can be dropped

Coupling $\mu_{\ell+1}$ and μ_ℓ

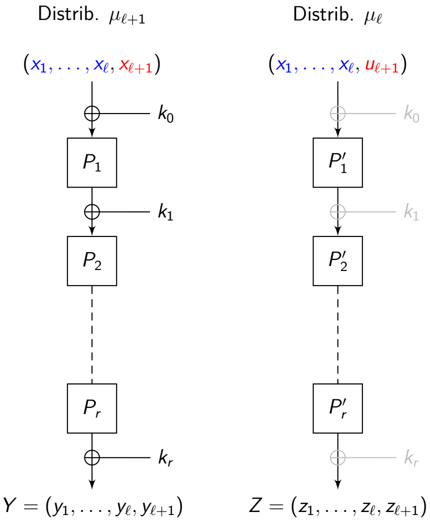

- we will define the second EM cipher (keys and permutations) as a function of the first one in order to have $Y = Z$ with high probability
- first, we choose exactly the same keys

Coupling $\mu_{\ell+1}$ and μ_ℓ

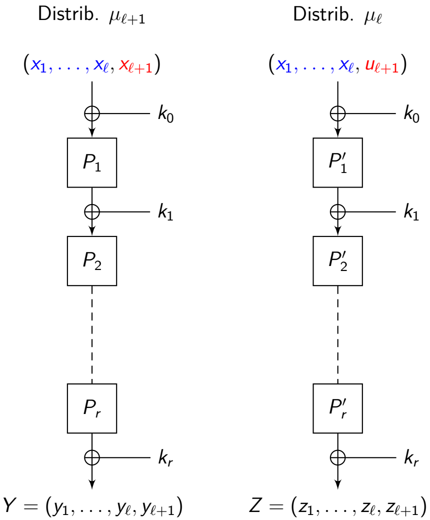

- we will define the second EM cipher (keys and permutations) as a function of the first one in order to have $Y = Z$ with high probability
- first, we choose exactly the same keys

Coupling $\mu_{\ell+1}$ and μ_ℓ

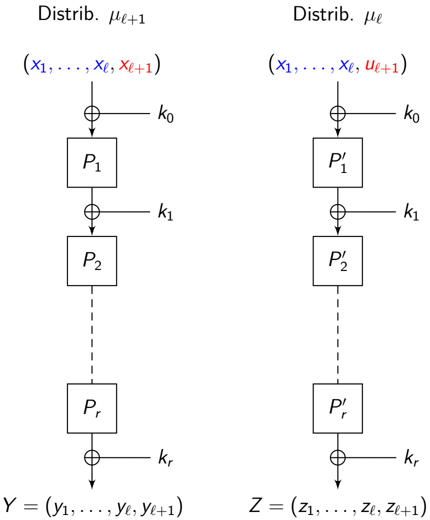

- we will define the permutations P'_i so that $Y = Z$ with high probability
- first, we define $P'_i(\cdot) = P_i(\cdot)$ on all points encountered during the encryption of $x_1, \dots, x_\ell$ $\rightarrow$ this implies $y_1 = z_1, \dots, y_\ell = z_\ell$

Coupling $\mu_{\ell+1}$ and μ_ℓ

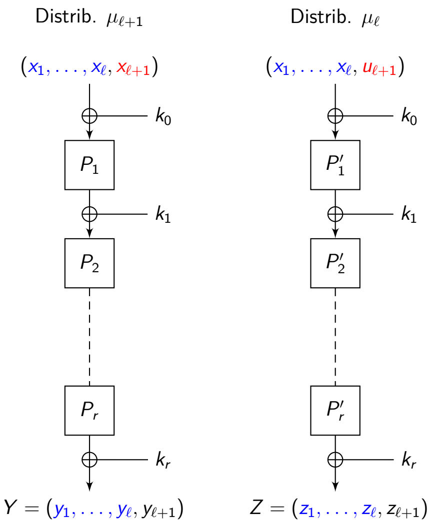

- we will define the permutations P'_i so that $Y = Z$ with high probability
- first, we define $P'_i(\cdot) = P_i(\cdot)$ on all points encountered during the encryption of $x_1, \dots, x_\ell$ $\to$ this implies $y_1 = z_1, \dots, y_\ell = z_\ell$

Coupling $\mu_{\ell+1}$ and μ_ℓ

Distrib. $\mu_{\ell+1}$

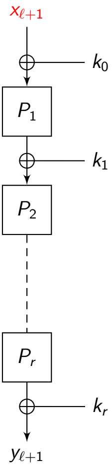

Distrib. μ_ℓ

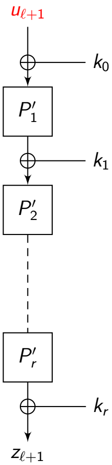

- it remains to equate $y_{\ell+1}$ and $z_{\ell+1}$
- let $x^i_{\ell+1}$, resp. $u^i_{\ell+1}$ denote the input to P_i, resp P'_i, while encrypting $x_{\ell+1}$, resp. $u_{\ell+1}$
- recall: the permutations P_i and P'_i must satisfy the equations $P_i(a_{i,j}) = b_{i,j}$
- we say $x^i_{\ell+1}$, resp. $u^i_{\ell+1}$ is free if it is different from all $a_{i,j}$'s, $j \in [1, q]$

Coupling $\mu_{\ell+1}$ and μ_ℓ

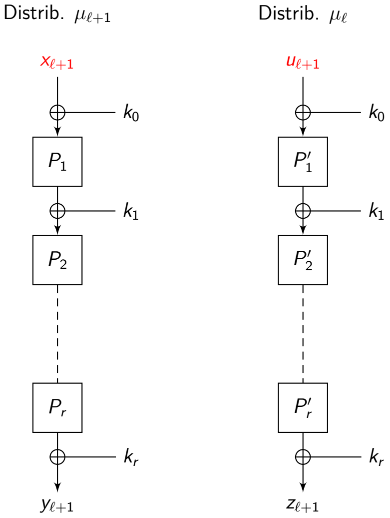

- it remains to equate $y_{\ell+1}$ and $z_{\ell+1}$
- let $x^i_{\ell+1}$, resp. $u^i_{\ell+1}$ denote the input to P_i, resp P'_i, while encrypting $x_{\ell+1}$, resp. $u_{\ell+1}$
- recall: the permutations P_i and P'_i must satisfy the equations $P_i(a_{i,j}) = b_{i,j}$
- we say $x^i_{\ell+1}$, resp. $u^i_{\ell+1}$ is free if it is different from all $a_{i,j}$'s, $j \in [1, q]$

Coupling $\mu_{\ell+1}$ and μ_ℓ

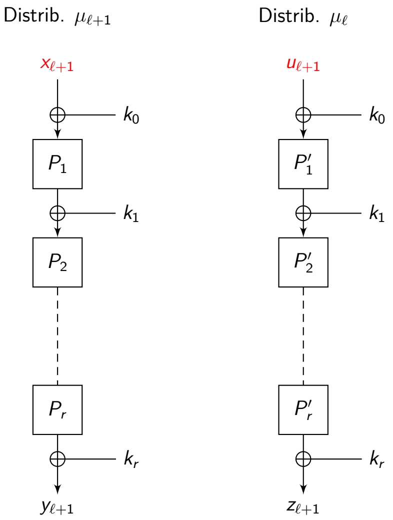

- it remains to equate $y_{\ell+1}$ and $z_{\ell+1}$
- let $x^i_{\ell+1}$, resp. $u^i_{\ell+1}$ denote the input to P_i, resp P'_i, while encrypting $x_{\ell+1}$, resp. $u_{\ell+1}$
- recall: the permutations P_i and P'_i must satisfy the equations $P_i(a_{i,j}) = b_{i,j}$
- we say $x^i_{\ell+1}$, resp. $u^i_{\ell+1}$ is free if it is different from all $a_{i,j}$'s, $j \in [1, q]$

Coupling $\mu_{\ell+1}$ and μ_ℓ

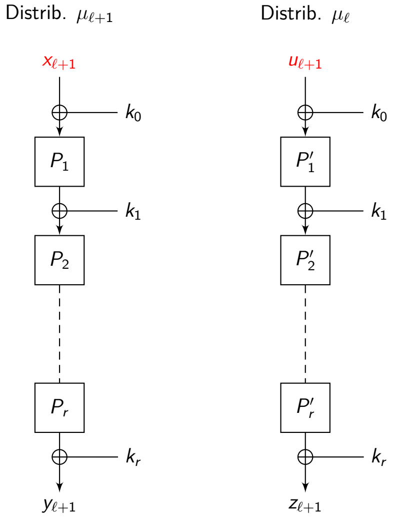

- it remains to equate $y_{\ell+1}$ and $z_{\ell+1}$
- let $x^i_{\ell+1}$, resp. $u^i_{\ell+1}$ denote the input to P_i, resp P'_i, while encrypting $x_{\ell+1}$, resp. $u_{\ell+1}$
- recall: the permutations P_i and P'_i must satisfy the equations $P_i(a_{i,j}) = b_{i,j}$
- we say $x^i_{\ell+1}$, resp. $u^i_{\ell+1}$ is free if it is different from all $a_{i,j}$'s, $j \in [1, q]$

Coupling $\mu_{\ell+1}$ and μ_ℓ

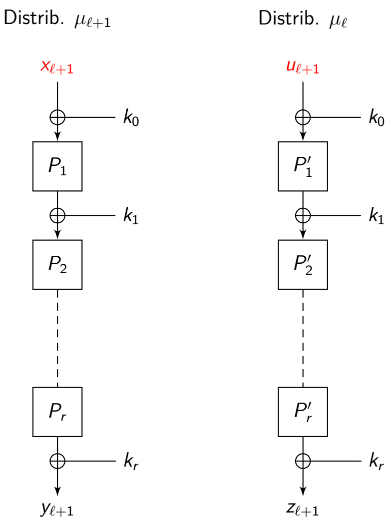

- we proceed iteratively for $i = 1..r$ as follows:
  - if $u^i_{\ell+1}$ is not free, then $P'_i(u^i_{\ell+1})$ is imposed by the equations $P'_i(a_{i,j}) = b_{i,j}$
  - if $u^i_{\ell+1}$ is free but $x^i_{\ell+1}$ is not, we define $P'_i(u^i_{\ell+1})$ uniformly at random among possible values
  - if $u^i_{\ell+1}$ and $x^i_{\ell+1}$ are both free, we define

$$P'_i(u^i_{\ell+1}) = P_i(x^i_{\ell+1})$$

$\rightarrow$ successful coupling, the subsequent outputs remain equal

Coupling $\mu_{\ell+1}$ and μ_ℓ

We have $Y \neq Z$ only if we fail to couple at all rounds $i = 1, \ldots, r$.

Probability to fail to couple at round i
(given that it failed at rounds $1, \ldots, i-1$):
Since $x^i_{\ell+1}$ and $u^i_{\ell+1}$ are randomized by key k_{i-1}, and since $|(a_{i,j})| = q$, the probability that $x^i_{\ell+1}$ or $u^i_{\ell+1}$ is not free is at most $2q/N$.

Hence, the probability to fail to couple at all r rounds and to have $Y \neq Z$ at the output of the two EM ciphers is:

$$\Pr[Y \neq Z] \leq \left(\frac{2q}{N}\right)^r .$$

Coupling $\mu_{\ell+1}$ and μ_ℓ

We have $Y \neq Z$ only if we fail to couple at all rounds $i = 1, \ldots, r$.

Probability to fail to couple at round i
(given that it failed at rounds $1, \ldots, i-1$):
Since $x^i_{\ell+1}$ and $u^i_{\ell+1}$ are randomized by key k_{i-1}, and since $|(a_{i,j})| = q$, the probability that $x^i_{\ell+1}$ or $u^i_{\ell+1}$ is not free is at most $2q/N$.

Hence, the probability to fail to couple at all r rounds and to have $Y \neq Z$ at the output of the two EM ciphers is:

$$\Pr[Y \neq Z] \leq \left(\frac{2q}{N}\right)^r .$$

Coupling $\mu_{\ell+1}$ and μ_ℓ

We have $Y \neq Z$ only if we fail to couple at all rounds $i = 1, \ldots, r$.

Probability to fail to couple at round i
(given that it failed at rounds $1, \ldots, i-1$):
Since $x^i_{\ell+1}$ and $u^i_{\ell+1}$ are randomized by key k_{i-1}, and since $|(a_{i,j})| = q$, the probability that $x^i_{\ell+1}$ or $u^i_{\ell+1}$ is not free is at most $2q/N$.

Hence, the probability to fail to couple at all r rounds and to have $Y \neq Z$ at the output of the two EM ciphers is:

$$\Pr[Y \neq Z] \leq \left(\frac{2q}{N}\right)^r \quad .$$

Concluding the proof

By the coupling lemma

$$\|\mu_{\ell+1} - \mu_\ell\| \leq \Pr[Y \neq Z] \leq \left(\frac{2q}{N}\right)^r \quad .$$

Hence:

$$\|\mu_q - \mu_0\| \leq \sum_{\ell=0}^{q-1} \|\mu_{\ell+1} - \mu_\ell\| \leq 2^r \frac{q^{r+1}}{N^r} \quad .$$

which gives the upper bound on advantage on any NCPA distinguisher.

Concluding the proof

By the coupling lemma

$$\|\mu_{\ell+1} - \mu_\ell\| \leq \Pr[Y \neq Z] \leq \left(\frac{2q}{N}\right)^r \quad .$$

Hence:

$$\|\mu_q - \mu_0\| \leq \sum_{\ell=0}^{q-1} \|\mu_{\ell+1} - \mu_\ell\| \leq 2^r \frac{q^{r+1}}{N^r} \quad .$$

which gives the upper bound on advantage on any NCPA distinguisher.

From NCPA to CCA security

We use the following "two weak make one strong" composition theorem:

Theorem ([MPR07])

Let E and F be two NCPA-secure block ciphers, with the same domain and resp. key spaces $\mathcal{K}_E$ and $\mathcal{K}_F$. Then $E \circ F^{-1}$ is a CCA-secure block cipher with key space $\mathcal{K}_E \times \mathcal{K}_F$.

The IEM cipher with $2r$ rounds is the composition of 2 IEM ciphers with r rounds (splitting the key $k_r = k'_r \oplus k''_r$):

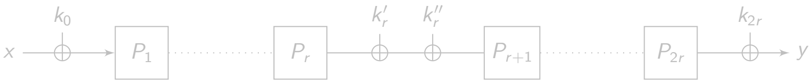

From NCPA to CCA security

We use the following "two weak make one strong" composition theorem:

Theorem ([MPR07])

Let E and F be two NCPA-secure block ciphers, with the same domain and resp. key spaces $\mathcal{K}_E$ and $\mathcal{K}_F$. Then $E \circ F^{-1}$ is a CCA-secure block cipher with key space $\mathcal{K}_E \times \mathcal{K}_F$.

The IEM cipher with $2r$ rounds is the composition of 2 IEM ciphers with r rounds (splitting the key $k_r = k'_r \oplus k''_r$):

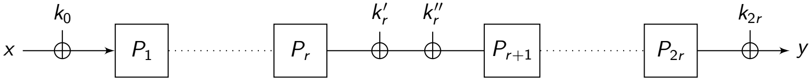

From NCPA to CCA security

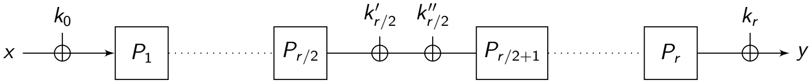

Theorem

For any CCA $\mathcal{D}$ making at most q queries to each oracle, the distinguishing advantage against the IEM with r rounds (r even) is at most

$$\mathcal{O}\left(\frac{q^{r/2+1}}{N^{r/2}}\right) = \mathcal{O}\left(\frac{q^{r+2}}{N^r}\right) .$$

$\rightarrow$ security up to $\mathcal{O}(N^{r/(r+2)})$ queries.

New result [CS14]: in fact, security up to $\mathcal{O}(N^{r/(r+1)})$ queries as well.

From NCPA to CCA security

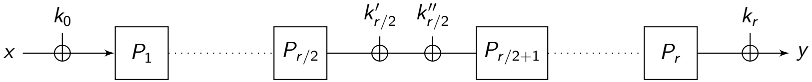

Theorem

For any CCA $\mathcal{D}$ making at most q queries to each oracle, the distinguishing advantage against the IEM with r rounds (r even) is at most

$$\mathcal{O}\left(\frac{q^{r/2+1}}{N^{r/2}}\right) = \mathcal{O}\left(\frac{q^{r+2}}{N^r}\right) .$$

$\rightarrow$ security up to $\mathcal{O}(N^{r/(r+2)})$ queries.

New result [CS14]: in fact, security up to $\mathcal{O}(N^{r/(r+1)})$ queries as well.

Extensions and open problems

- results can be extended to the case where the $(r+1)$ round keys are r-wise independent, e.g.:

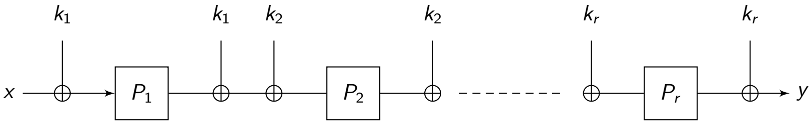

- what about the single-key IEM (all round keys equal)?
  current conjecture: similar bounds to the "independent round keys" case

Extensions and open problems

- results can be extended to the case where the $(r+1)$ round keys are r-wise independent, e.g.:

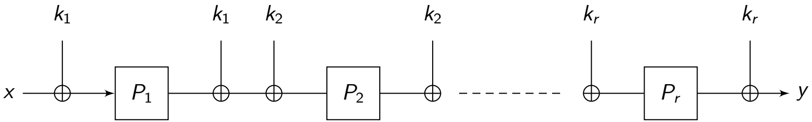

- what about the single-key IEM (all round keys equal)?
  current conjecture: similar bounds to the "independent round keys" case

Outline

Tweakable block ciphers: definition

A tweakable block cipher (TBC) is a family of block ciphers indexed by a tweak $t \in \mathcal{T}$:

$$\widetilde{E} : \mathcal{T} \times \mathcal{K} \times \mathcal{M} \to \mathcal{M}$$

The tweak is a public parameter (under the control of the adversary in the security model)

Introduced by Liskov, Rivest, and Wagner at CRYPTO 2002 [LRW02].

Tweakable block ciphers: definition

A tweakable block cipher (TBC) is a family of block ciphers indexed by a tweak $t \in \mathcal{T}$:

$$\widetilde{E} : \mathcal{T} \times \mathcal{K} \times \mathcal{M} \to \mathcal{M}$$

The tweak is a public parameter (under the control of the adversary in the security model)

Introduced by Liskov, Rivest, and Wagner at CRYPTO 2002 [LRW02].

The original [LRW02] construction

Liskov et al. proposed the following construction of a TBC from an existing blockcipher E:

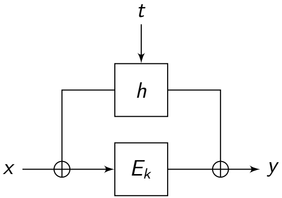

h is an $\varepsilon-\mathrm{AXU}_2$ function: $\Pr_h[h(x) \oplus h(x') = y] \leq \varepsilon$.

[LRW02] proved security (against CCA adversaries) up to $\mathcal{O}(2^{n/2})$ queries (n is the block size of E)

The original [LRW02] construction

Liskov et al. proposed the following construction of a TBC from an existing blockcipher E:

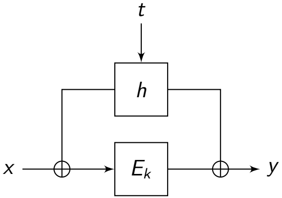

h is an $\varepsilon-\mathrm{AXU}_2$ function: $\Pr_h[h(x) \oplus h(x') = y] \leq \varepsilon$.

[LRW02] proved security (against CCA adversaries) up to $\mathcal{O}(2^{n/2})$ queries (n is the block size of E)

The [LST12] construction

At CRYPTO 2012, Landecker et al. extended the LRW construction as follows:

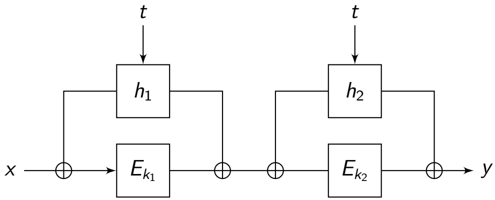

[LST12] proved security (against CCA adversaries) up to $\mathcal{O}(2^{2n/3})$ queries.

Extension to r rounds

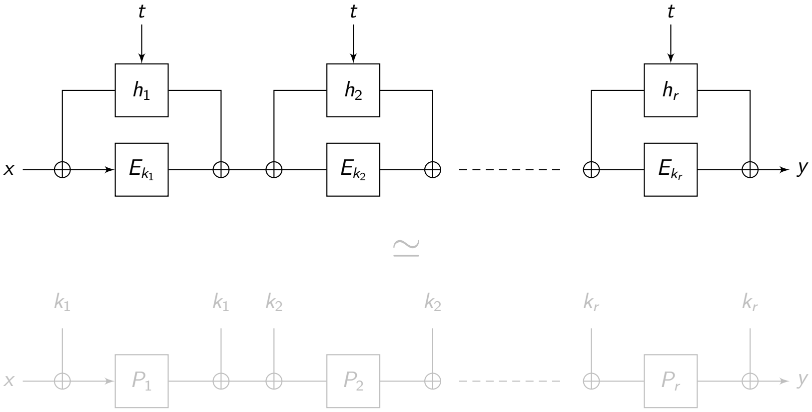

Extension to r rounds

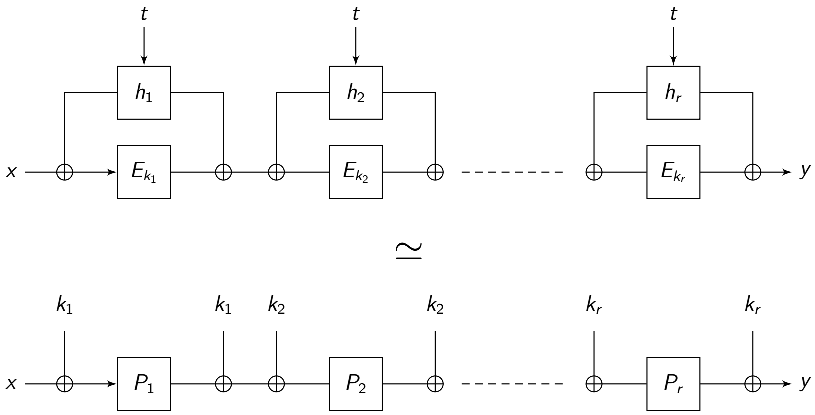

Extension to r rounds

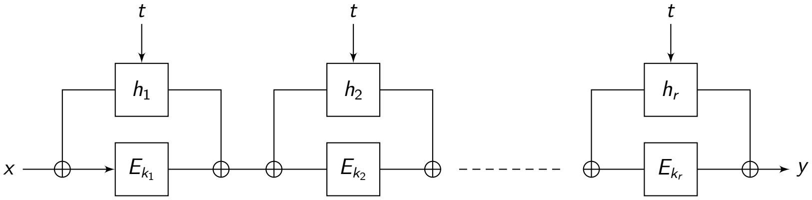

For this TBC construction, one can prove results similar to the ones for the IEM cipher [LS13]:

- secure against NCPA distinguishers up to $\mathcal{O}(2^{rn/(r+1)})$ queries
- secure against CCA distinguishers up to $\mathcal{O}(2^{rn/(r+2)})$ queries

Outline

Outline

From indistinguishability to indifferentiability

Previous results state that the IEM cipher is a (strong) pseudorandom permutation (in the random permutation model)
$=$ usual single, secret key security model

What about related-, known- or chosen-key attacks?
$\rightarrow$ prove the IEM is indifferentiable from an ideal cipher

Ideal cipher: draw an independent random permutation for each key

From indistinguishability to indifferentiability

Previous results state that the IEM cipher is a (strong) pseudorandom permutation (in the random permutation model)
$=$ usual single, secret key security model

What about related-, known- or chosen-key attacks?
$\rightarrow$ prove the IEM is indifferentiable from an ideal cipher

Ideal cipher: draw an independent random permutation for each key

From indistinguishability to indifferentiability

Previous results state that the IEM cipher is a (strong) pseudorandom permutation (in the random permutation model)
$=$ usual single, secret key security model

What about related-, known- or chosen-key attacks?
$\rightarrow$ prove the IEM is indifferentiable from an ideal cipher

Ideal cipher: draw an independent random permutation for each key

A word on the ideal cipher model

- the pseudorandomness security notion for a block cipher is sufficient to prove the security of a lot of applications (encryption modes and MACs)
- however, sometimes it is not sufficient (e.g. for block cipher-based hash functions like Davies-Meyer mode)
- ideally, one expects that a good block cipher "behaves" as an independent random permutation for each key
  $\rightarrow$ ideal cipher model
- similar to the random oracle model for a hash function
  warning: instantiation problems as well (no concrete block cipher can be proved to be an ideal cipher in any reasonable sense)
- though we cannot prove that a block cipher behaves as an ideal cipher in the standard model, we can prove results in idealized models (e.g. the Random Permutation Model that we already used for the IEM cipher)
  $\rightarrow$ indifferentiability notion

A word on the ideal cipher model

- the pseudorandomness security notion for a block cipher is sufficient to prove the security of a lot of applications (encryption modes and MACs)
- however, sometimes it is not sufficient (e.g. for block cipher-based hash functions like Davies-Meyer mode)
- ideally, one expects that a good block cipher "behaves" as an independent random permutation for each key
  → ideal cipher model
- similar to the random oracle model for a hash function
  warning: instantiation problems as well (no concrete block cipher can be proved to be an ideal cipher in any reasonable sense)
- though we cannot prove that a block cipher behaves as an ideal cipher in the standard model, we can prove results in idealized models (e.g. the Random Permutation Model that we already used for the IEM cipher)
  → indifferentiability notion

A word on the ideal cipher model

- the pseudorandomness security notion for a block cipher is sufficient to prove the security of a lot of applications (encryption modes and MACs)
- however, sometimes it is not sufficient (e.g. for block cipher-based hash functions like Davies-Meyer mode)
- ideally, one expects that a good block cipher "behaves" as an independent random permutation for each key
  $\rightarrow$ ideal cipher model
- similar to the random oracle model for a hash function
  warning: instantiation problems as well (no concrete block cipher can be proved to be an ideal cipher in any reasonable sense)
- though we cannot prove that a block cipher behaves as an ideal cipher in the standard model, we can prove results in idealized models (e.g. the Random Permutation Model that we already used for the IEM cipher)
  $\rightarrow$ indifferentiability notion

A word on the ideal cipher model

- the pseudorandomness security notion for a block cipher is sufficient to prove the security of a lot of applications (encryption modes and MACs)
- however, sometimes it is not sufficient (e.g. for block cipher-based hash functions like Davies-Meyer mode)
- ideally, one expects that a good block cipher "behaves" as an independent random permutation for each key
  $\rightarrow$ ideal cipher model
- similar to the random oracle model for a hash function
  warning: instantiation problems as well (no concrete block cipher can be proved to be an ideal cipher in any reasonable sense)
- though we cannot prove that a block cipher behaves as an ideal cipher in the standard model, we can prove results in idealized models (e.g. the Random Permutation Model that we already used for the IEM cipher)
  $\rightarrow$ indifferentiability notion

A word on the ideal cipher model

- the pseudorandomness security notion for a block cipher is sufficient to prove the security of a lot of applications (encryption modes and MACs)
- however, sometimes it is not sufficient (e.g. for block cipher-based hash functions like Davies-Meyer mode)
- ideally, one expects that a good block cipher "behaves" as an independent random permutation for each key
  $\rightarrow$ ideal cipher model
- similar to the random oracle model for a hash function
  warning: instantiation problems as well (no concrete block cipher can be proved to be an ideal cipher in any reasonable sense)
- though we cannot prove that a block cipher behaves as an ideal cipher in the standard model, we can prove results in idealized models (e.g. the Random Permutation Model that we already used for the IEM cipher)
  $\rightarrow$ indifferentiability notion

Indifferentiability: definition

Definition

A construction $\mathcal{C}^{\boldsymbol{F}}$ (here, the IEM cipher $\mathtt{EM}^{P_1,\dots,P_r}$) using an ideal primitive $\boldsymbol{F}$ (here, random permutations $P_1, \dots, P_r$) is said indifferentiable from an ideal primitive $\boldsymbol{G}$ (here, an ideal cipher E) if there exists a polynomial time simulator $\mathcal{S}$ with access to $\boldsymbol{G}$ such that the two systems $(\mathcal{C}^{\boldsymbol{F}}, \boldsymbol{F})$ and $(\boldsymbol{G}, \mathcal{S}^{\boldsymbol{G}})$ are indistinguishable.

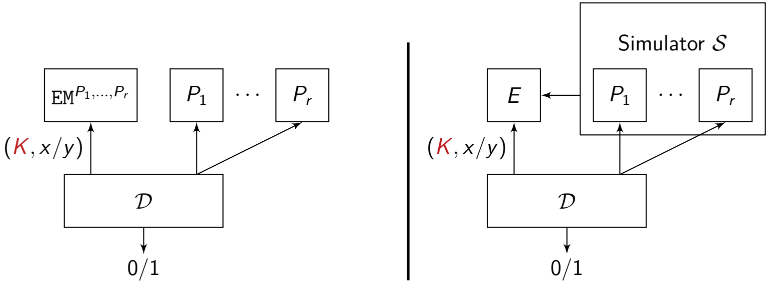

Indifferentiability: definition

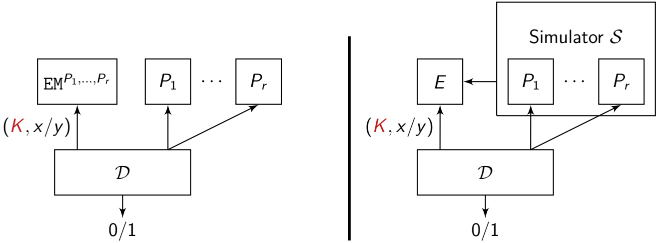

The answers of the simulator $\mathcal{S}$ must be:

- coherent with answers the distinguisher can obtain directly from E
- close in distribution to the answers of a random permutation

NB: The distinguisher specifies the key and the plaintext/ciphertext when querying $\mathtt{EM}^{P_1,\ldots,P_r}$ or E.

Composition theorem

Usefulness of indifferentiability: composition theorem

Theorem

If a cryptosystem Γ is secure when used with an ideal primitive $\boldsymbol{G}$, and if $\mathcal{C}^{\boldsymbol{F}}$ is indifferentiable from $\boldsymbol{G}$, then Γ is also secure when used with $\mathcal{C}^{\boldsymbol{F}}$.

Sketch of the proof:

- assume $\mathcal{C}^{\boldsymbol{F}}$ is indifferentiable from $\boldsymbol{G}$
- assume there is an attacker $\mathcal{A}$ with advantage ε against some cryptosystem Γ using the construction $\mathcal{C}^{\boldsymbol{F}}$
- then one can consider the simulator $\mathcal{S}$ ensured by indifferentiability
- combining $\mathcal{A}$ and $\mathcal{S}$, one obtains an new attacker $\mathcal{A}'$ against cryptosystem Γ used with $\boldsymbol{G}$ with advantage $\simeq \varepsilon$, a contradiction

Composition theorem

Usefulness of indifferentiability: composition theorem

Theorem

If a cryptosystem Γ is secure when used with an ideal primitive $\boldsymbol{G}$, and if $\mathcal{C}^{\boldsymbol{F}}$ is indifferentiable from $\boldsymbol{G}$, then Γ is also secure when used with $\mathcal{C}^{\boldsymbol{F}}$.

Sketch of the proof:

- assume $\mathcal{C}^{\boldsymbol{F}}$ is indifferentiable from $\boldsymbol{G}$
- assume there is an attacker $\mathcal{A}$ with advantage ε against some cryptosystem Γ using the construction $\mathcal{C}^{\boldsymbol{F}}$
- then one can consider the simulator $\mathcal{S}$ ensured by indifferentiability
- combining $\mathcal{A}$ and $\mathcal{S}$, one obtains an new attacker $\mathcal{A}'$ against cryptosystem Γ used with $\boldsymbol{G}$ with advantage $\simeq \varepsilon$, a contradiction

Outline

Independent round keys fails

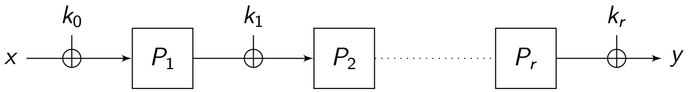

This is not indifferentiable from an ideal cipher with key space $\{0,1\}^{(r+1)n}$ because of the following distinguisher:

- fix a non-zero constant $c \in \{0,1\}^n$
- choose an arbitrary $x \in \{0,1\}^n$ and $k_0 \in \{0,1\}^n$
- define $x' = x \oplus c$ and $k_0' = k_0 \oplus c$
- let $K = (k_0, k_1, \ldots, k_r)$ and $K' = (k_0', k_1, \ldots, k_r)$
- then $\mathtt{EM}(K, x) = \mathtt{EM}(K', x')$
- this holds only with negligible probability for an ideal cipher

Proving indifferentiability for key-alternating ciphers

Independent keys leave too much “freedom” to the adversary.

Two ideas to solve the problem:

1. add a key schedule, and put some cryptographic assumption on it
   $\Rightarrow$ Andreeva et al. CRYPTO 2013 [ABD$^+$13]
2. restrain the key space and correlate the round keys, e.g. $(k, k, \ldots, k)$
   $\Rightarrow$ Lampe and Seurin 2013 (preprint)

The [ABD+13] result

The key-derivation function is modeled as a random oracle from $\{0,1\}^\ell$ to $\{0,1\}^n$ (that the adversary queries in a black-box way)

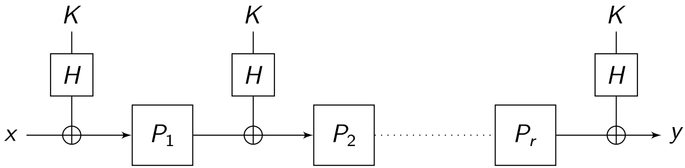

$\rightarrow$ indifferentiable from an ideal cipher with ℓ-bit keys for $r = 5$ ([ABD+13] gives attacks up to 3 rounds)

The assumption about the key derivation is very strong and far from concrete designs (the key-schedule is often invertible)

The [ABD$^+$13] result

The key-derivation function is modeled as a random oracle from $\{0,1\}^\ell$ to $\{0,1\}^n$ (that the adversary queries in a black-box way)

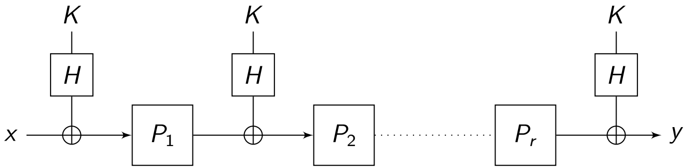

$\rightarrow$ indifferentiable from an ideal cipher with ℓ-bit keys for $r = 5$ ([ABD$^+$13] gives attacks up to 3 rounds)

The assumption about the key derivation is very strong and far from concrete designs (the key-schedule is often invertible)

Our approach

We consider the IEM with a single key:

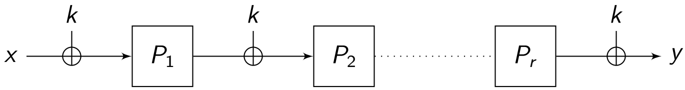

The trivial attack on independent keys does not apply → is it indiff. from an ideal cipher for sufficiently many rounds ?

Main Result

The single-key IEM with $r = 12$ rounds is indifferentiable from an ideal cipher with n-bit blocks and n-bit keys

Also holds when using invertible permutations γ_i for the key derivation (no cryptographic assumption needed).

Our approach

We consider the IEM with a single key:

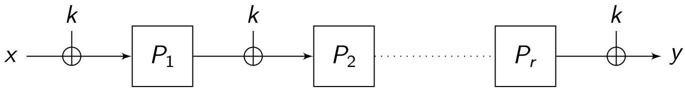

The trivial attack on independent keys does not apply $\rightarrow$ is it indiff. from an ideal cipher for sufficiently many rounds ?

Main Result

The single-key IEM with $r = 12$ rounds is indifferentiable from an ideal cipher with n-bit blocks and n-bit keys

Also holds when using invertible permutations γ_i for the key derivation (no cryptographic assumption needed).

Outline

A simple attack for 1 round

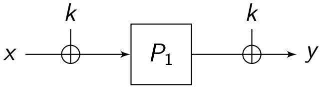

The distinguisher $\mathcal{D}$ proceeds as follows:

- query $P_1(a) = b$ for an arbitrary a
- choose a random key k and define $x = a \oplus k$
- query $E(k, x) = y$ and check whether $y = b \oplus k$ $(*)$

Then:

- when $\mathcal{D}$ interacts with a real EM cipher, $(*)$ always holds
- when $\mathcal{D}$ interacts with $(E, \mathcal{S}^E)$, $(*)$ holds only with negligible probability since $\mathcal{S}$ cannot guess k when answering the query $P_1(a)$

A simple attack for 1 round

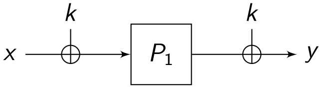

The distinguisher $\mathcal{D}$ proceeds as follows:

- query $P_1(a) = b$ for an arbitrary a
- choose a random key k and define $x = a \oplus k$
- query $E(k, x) = y$ and check whether $y = b \oplus k$ $(*)$

Then:

- when $\mathcal{D}$ interacts with a real EM cipher, $(*)$ always holds
- when $\mathcal{D}$ interacts with $(E, \mathcal{S}^E)$, $(*)$ holds only with negligible probability since $\mathcal{S}$ cannot guess k when answering the query $P_1(a)$

An attack for 3 rounds

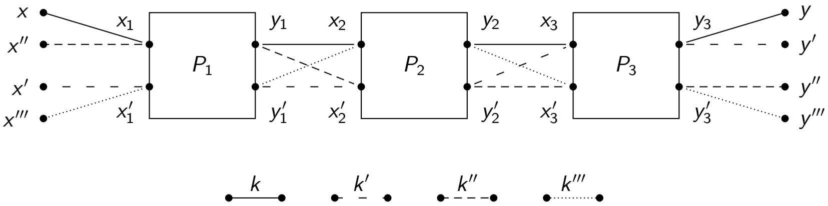

One can (easily) find (x, x', x'', x'''), (y, y', y'', y''') and (k, k', k'', k''') such that $y = \mathtt{EM}^{(P_1,P_2,P_3)}(k, x)$, etc. and:

$$\begin{cases} k \oplus k' \oplus k'' \oplus k''' = 0 \\ x \oplus x' \oplus x'' \oplus x''' = 0 \\ y \oplus y' \oplus y'' \oplus y''' = 0 \end{cases} .$$

This can be showed to be hard for an ideal cipher.

Outline

Simulation: general strategy

The simulator must return answers that are coherent with what the distinguisher can obtain from the ideal cipher E, i.e.:

$$\mathtt{EM}^{P_1,\ldots,P_{12}}(k,x) = E(k,x)$$

For this, the simulator must adapt at least one permutation to "match" what is given by the ideal cipher

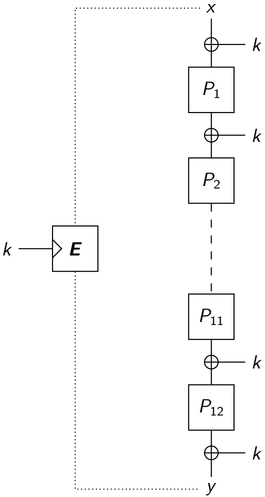

Simulation: general strategy

- the simulator detects and completes "partial chains" = two adjacent queries $P_i(x_i) = y_i$ and $P_{i+1}(x_{i+1}) = y_{i+1}$
- for any partial chain the key is uniquely defined: $k = y_i \oplus x_{i+1}$
- when a partial chain is detected, the simulator completes the missing permutation values randomly, except for one particular permutation which is "adapted" to match the ideal cipher

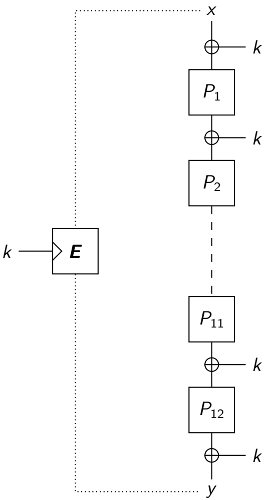

How the simulator works

- the simulator only detects partial chains at very specific places:
  - external chains $(P_1, P_2, P_{11}, P_{12})$ that matches the ideal cipher E
  - central chains (P_6, P_7)
- an external chain can be created only if the distinguisher has made the corresponding query to E $\rightarrow$ only q of them will be completed, which avoids an recursive blow-up of the simulator

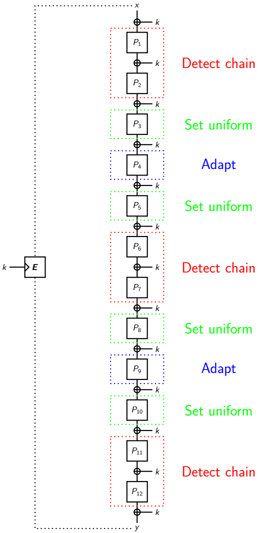

How the simulator works

- the simulator only detects partial chains at very specific places:
    - external chains $(P_1, P_2, P_{11}, P_{12})$ that matches the ideal cipher E
    - central chains (P_6, P_7)
- an external chain can be created only if the distinguisher has made the corresponding query to E $\rightarrow$ only q of them will be completed, which avoids an recursive blow-up of the simulator

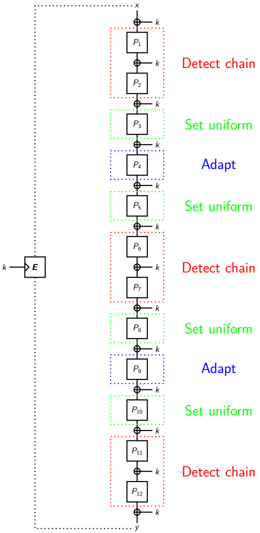

How the simulator works

- the simulator uses specific permutations to adapt chains: P_4 and P_9
- main difficulty: show that the simulator can always adapt (i.e. the permutation has not already been defined on the point needed for adaptation)

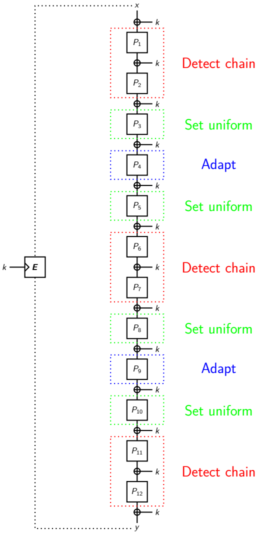

How the simulator works

- the simulator uses specific permutations to adapt chains: P_4 and P_9
- main difficulty: show that the simulator can always adapt (i.e. the permutation has not already been defined on the point needed for adaptation)

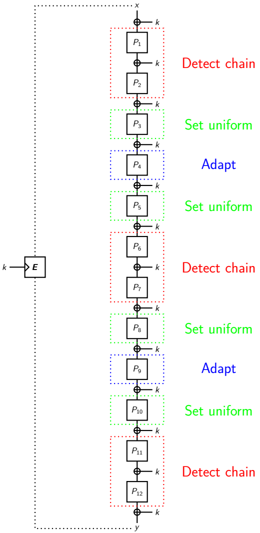

Open problems

The indifferentiability proof requires 12 rounds, but the best attack is only on 3 rounds.

Conjecture

The single-key IEM with $3 < r < 12$ rounds is indifferentiable from an ideal cipher with n-bit keys

$r = 4$ may well be sufficient

Open problems

The indifferentiability proof requires 12 rounds, but the best attack is only on 3 rounds.

Conjecture

The single-key IEM with $3 < r < 12$ rounds is indifferentiable from an ideal cipher with n-bit keys

$r = 4$ may well be sufficient

Conclusion

Summary of results about the IEM cipher:

- pseudorandomness: the IEM cipher with r rounds is indistinguishable from a random permutation up to $\mathcal{O}(N^{r/(r+1)})$ queries
- indifferentiability: the single-key IEM cipher with 12 rounds is indifferentiable from an ideal cipher with n-bit keys

Interpretation of the results:

- shows that the general strategy of building block ciphers from SPNs is sound and may even yield something close to an ideal cipher
- says little about concrete block ciphers: e.g. the permutations P_1, $\ldots$, P_{10} of AES-128 are to simple
- heuristic insurance for e.g. an IEM cipher where the P_i's are instantiated with AES used with fixed keys

Conclusion

Summary of results about the IEM cipher:

- pseudorandomness: the IEM cipher with r rounds is indistinguishable from a random permutation up to $\mathcal{O}(N^{r/(r+1)})$ queries
- indifferentiability: the single-key IEM cipher with 12 rounds is indifferentiable from an ideal cipher with n-bit keys

Interpretation of the results:

- shows that the general strategy of building block ciphers from SPNs is sound and may even yield something close to an ideal cipher
- says little about concrete block ciphers: e.g. the permutations P_1, $\ldots$, P_{10} of AES-128 are to simple
- heuristic insurance for e.g. an IEM cipher where the P_i's are instantiated with AES used with fixed keys

The end. . .

Thanks for your attention!

Comments or questions?

References I

Elena Andreeva, Andrey Bogdanov, Yevgeniy Dodis, Bart Mennink, and John P. Steinberger.

On the Indifferentiability of Key-Alternating Ciphers.

In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 (Proceedings, Part I)*, volume 8042 of *Lecture Notes in Computer Science*, pages 531–550. Springer, 2013.

Full version available at `http://eprint.iacr.org/2013/061`.

Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, François-Xavier Standaert, John P. Steinberger, and Elmar Tischhauser.

Key-Alternating Ciphers in a Provable Setting: Encryption Using a Small Number of Public Permutations - (Extended Abstract).

In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology - EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 45–62. Springer, 2012.

References II

Shan Chen and John Steinberger.
Tight Security Bounds for Key-Alternating Ciphers.
In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 327–350. Springer, 2014.
Full version available at `http://eprint.iacr.org/2013/222`.

Shimon Even and Yishay Mansour.
A Construction of a Cipher from a Single Pseudorandom Permutation.
Journal of Cryptology, 10(3):151–162, 1997.

Viet Tung Hoang and Phillip Rogaway.
On Generalized Feistel Networks.
In Tal Rabin, editor, *Advances in Cryptology - CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 613–630. Springer, 2010.

References III

Rodolphe Lampe, Jacques Patarin, and Yannick Seurin.

An Asymptotically Tight Security Analysis of the Iterated Even-Mansour Cipher.

In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology - ASIACRYPT 2012*, volume 7658 of *Lecture Notes in Computer Science*, pages 278–295. Springer, 2012.

Moses Liskov, Ronald L. Rivest, and David Wagner.

Tweakable Block Ciphers.

In Moti Yung, editor, *Advances in Cryptology - CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 31–46. Springer, 2002.

Rodolphe Lampe and Yannick Seurin.

Tweakable Blockciphers with Asymptotically Optimal Security.

In *Fast Software Encryption - FSE 2013*, 2013.

To appear.

References IV

Will Landecker, Thomas Shrimpton, and R. Seth Terashima.
Tweakable Blockciphers with Beyond Birthday-Bound Security.
In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 14–30. Springer, 2012.

Ilya Mironov.
(Not So) Random Shuffles of RC4.
In Moti Yung, editor, *Advances in Cryptology - CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 304–319. Springer, 2002.

Ueli M. Maurer, Krzysztof Pietrzak, and Renato Renner.
Indistinguishability Amplification.
In Alfred Menezes, editor, *Advances in Cryptology - CRYPTO 2007*, volume 4622 of *Lecture Notes in Computer Science*, pages 130–149. Springer, 2007.
Full version available at `http://eprint.iacr.org/2006/456`.

References V

Ben Morris, Phillip Rogaway, and Till Stegers.
How to Encipher Messages on a Small Domain.
In Shai Halevi, editor, *Advances in Cryptology - CRYPTO 2009*, volume 5677 of *Lecture Notes in Computer Science*, pages 286–302. Springer, 2009.