Tweakable Blockciphers with Asymptotically Optimal Security

Rodolphe Lampe[1] Yannick Seurin[2]

[1]University of Versailles, France
Financial support of DGA and ANR PRINCE

[2]ANSSI, Paris, France

11 March 2013

Introduction

Tweakable blockcipher: A family of blockcipher indexed with a tweak (a public parameter) : $\widetilde{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{D} \rightarrow \mathcal{D}$.

Introduced by Liskov-Rivest-Wagner at CRYPTO 2002

Introduction

Tweakable blockcipher: A family of blockcipher indexed with a tweak (a public parameter) : $\widetilde{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{D} \to \mathcal{D}$.

Introduced by Liskov-Rivest-Wagner at CRYPTO 2002

We consider constructions of tweakable blockciphers from an existing blockcipher.

One of the original construction of Liskov-Rivest-Wagner.

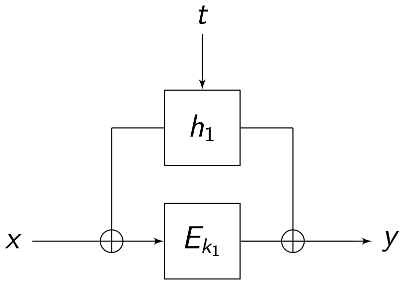

Figure: One of the original construction of Liskov-Rivest-Wagner.

h_1 is randomly chosen in $\mathcal{H}$ a family of $\varepsilon - AXU_2$ functions.

Secure up to $2^{n/2}$ queries against CCA attacks (n being the size of the blocks).

Definition of $\varepsilon - AXU_2$

Definition

Let S be an arbitrary set. A family of functions $\mathcal{H}$ from S to $\{0,1\}^n$ is said to be ε-almost-2-XOR-universal (ε-AXU$_2$) if for all distinct $x, x' \in S$ and all $y \in \{0,1\}^n$, one has

$$\Pr\left[h \leftarrow_{\$} \mathcal{H} \,:\, h(x) \oplus h(x') = y\right] \leq \varepsilon \ .$$

The construction of Landecker-Shrimpton-Terashima (CRYPTO 2012).

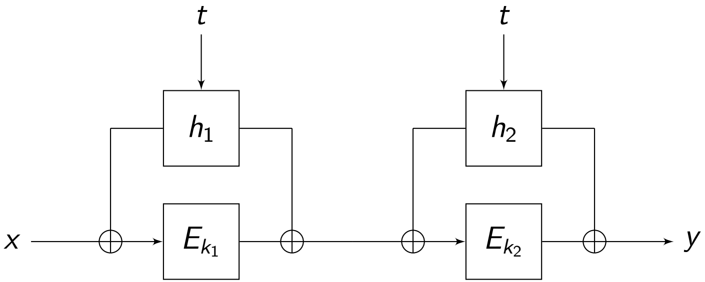

Figure: The construction of Landecker-Shrimpton-Terashima (CRYPTO 2012).

Secure up to $2^{\frac{2n}{3}}$ queries against CCA attacks (n being the size of the blocks).

Definition of r-CLRW

What if we increase the number of rounds ?

What if we increase the number of rounds ?

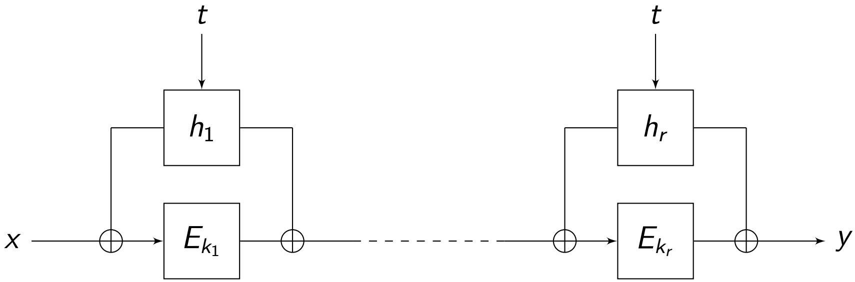

Figure: The $\mathtt{CLRW}^{r,E,\mathcal{H}}$ tweakable blockcipher construction.

Result

Theorem

Let $\mathcal{K}, \mathcal{T}$ be sets, $E \in \mathrm{BC}(\mathcal{K}, n)$ be a blockcipher, and $\mathcal{H}$ be a ε-AXU$_2$ family of functions from $\mathcal{T}$ to $\{0,1\}^n$. Then one has:

$$\mathbf{Adv}^{\widetilde{\mathrm{ncpa}}}_{\mathtt{CLRW}^r,E,\mathcal{H}}(q,\tau) \leq r \cdot \mathbf{Adv}^{\mathrm{ncpa}}_{E}(q,\tau+rqT) + \frac{q^{r+1}}{r+1}(2\varepsilon)^r$$

and

$$\mathbf{Adv}^{\widetilde{\mathrm{cca}}}_{\mathtt{CLRW}^r,E,\mathcal{H}}(q,\tau) \leq r \cdot \mathbf{Adv}^{\mathrm{cca}}_{E}(q,\tau+rqT) + \frac{4\sqrt{2}}{\sqrt{r+2}} q^{(r+2)/4}(2\varepsilon)^{r/4}$$

where T is the time to compute E or E^{-1}.

Secure up to $2^{\frac{r}{r+1}n}$ queries for NCPA attacks.
Secure up to $2^{\frac{r}{r+2}n}$ queries for CCA attacks.

Introducing intermediate worlds

Real World

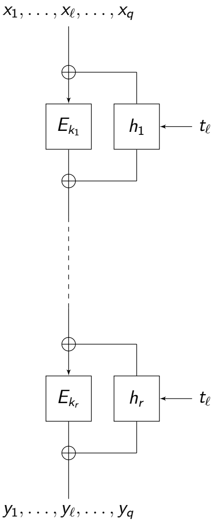

Introducing intermediate worlds

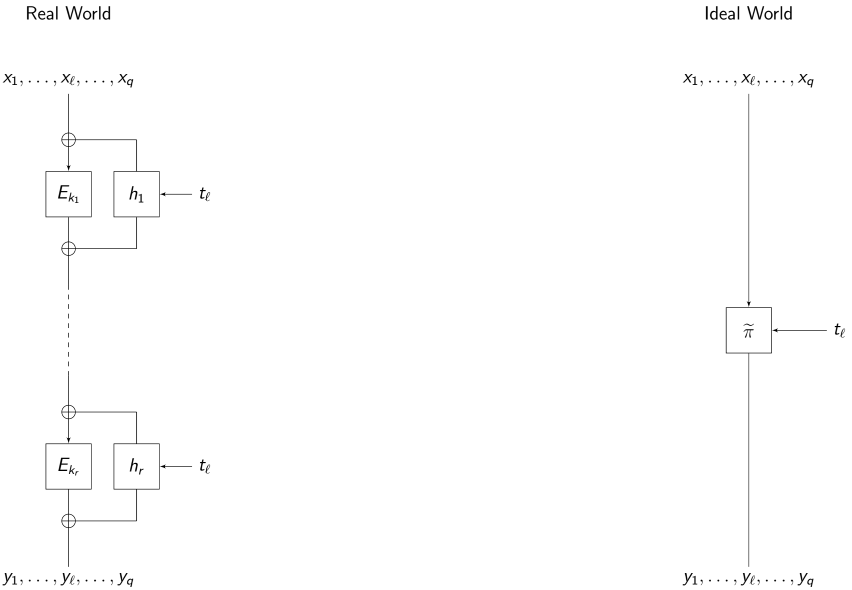

Introducing intermediate worlds

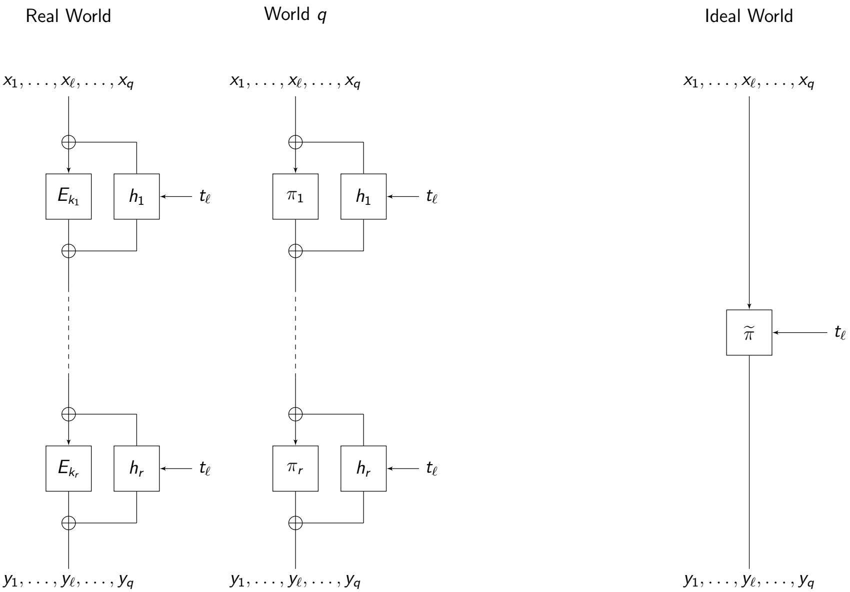

Introducing intermediate worlds

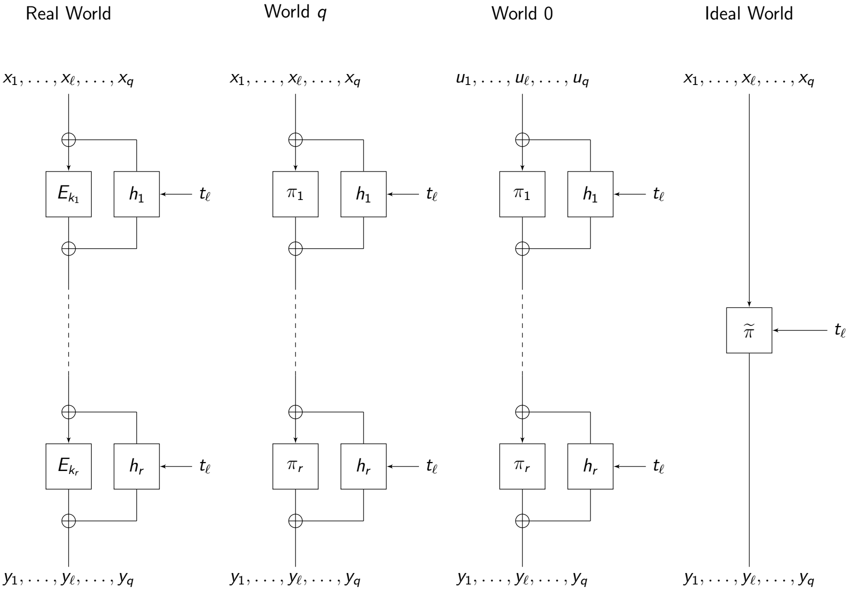

Introducing intermediate worlds

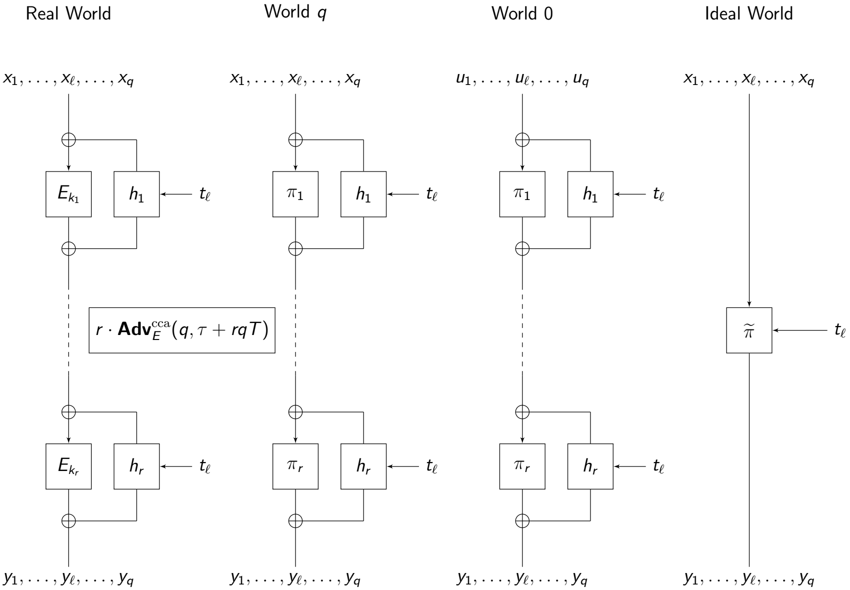

Introducing intermediate worlds

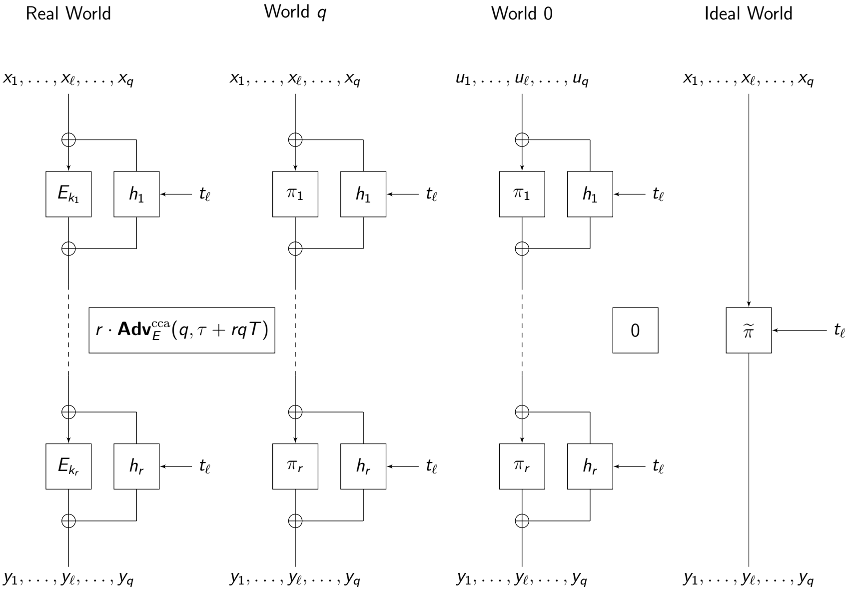

Introducing intermediate worlds

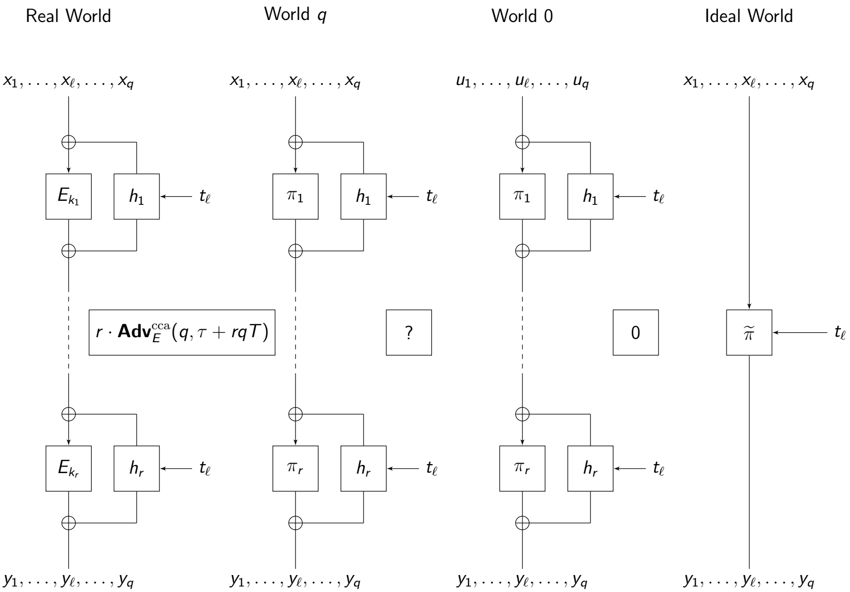

Still introducing intermediate worlds using different queries

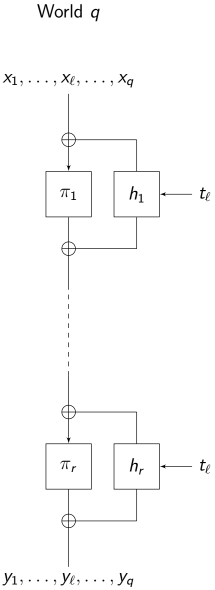

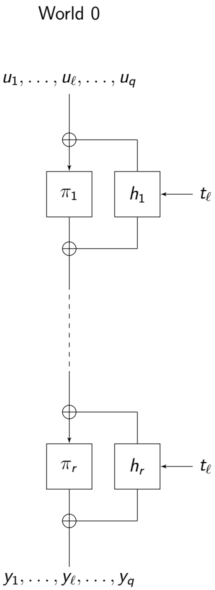

Still introducing intermediate worlds using different queries

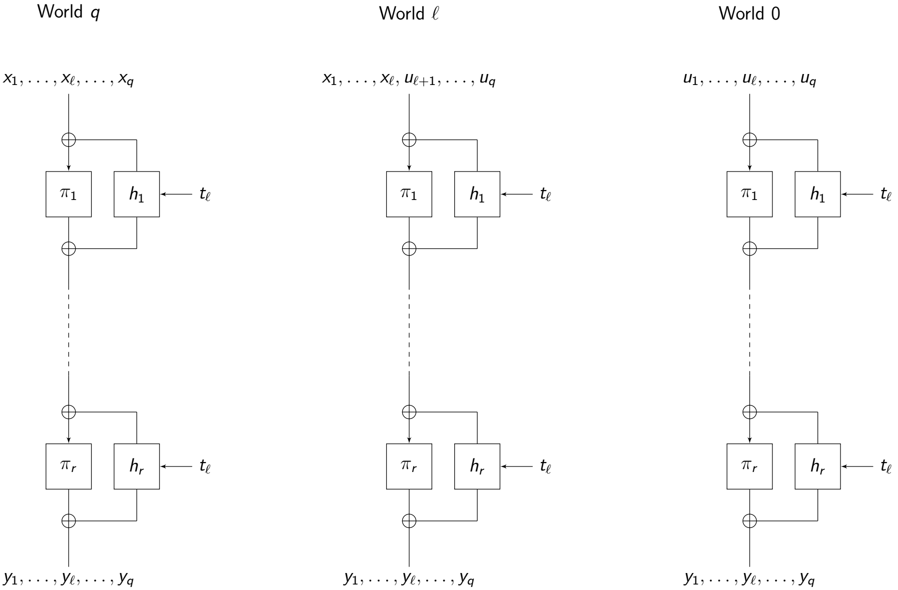

Still introducing intermediate worlds using different queries

World $\ell + 1$

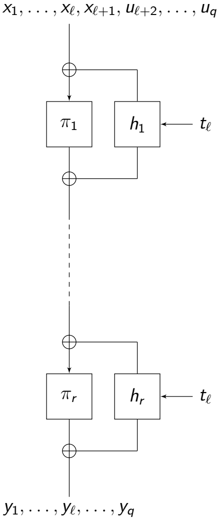

World ℓ

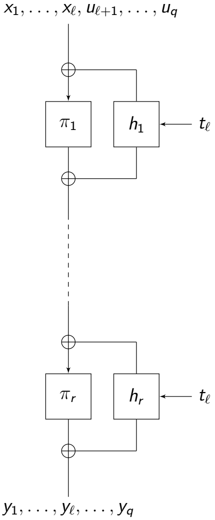

Still introducing intermediate worlds using different queries

World $\ell + 1$

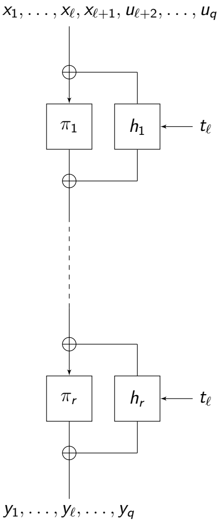

World ℓ

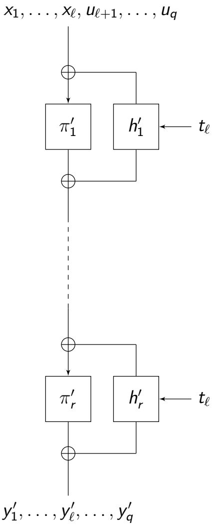

The last $q - \ell - 1$ outputs have the same distributions

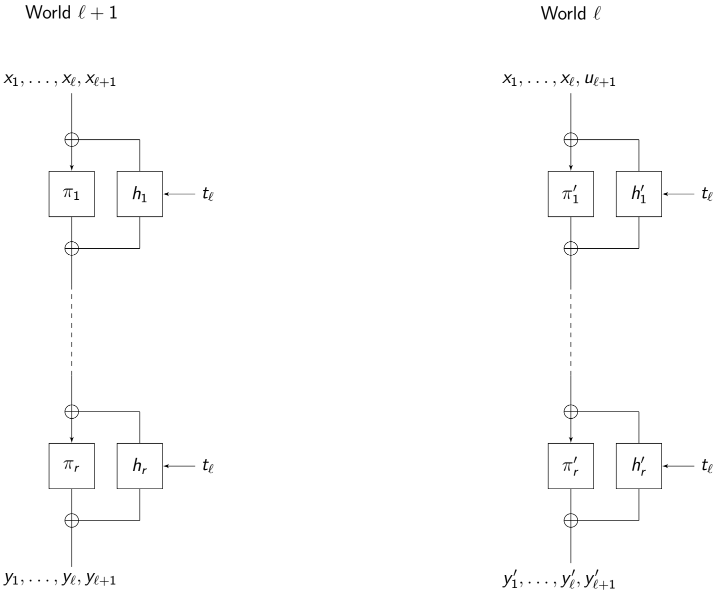

Coupling

A *coupling* of μ and ν is a distribution λ on $\Omega \times \Omega$ such that for all $x \in \Omega$, $\sum_{y\in\Omega} \lambda(x, y) = \mu(x)$ and for all $y \in \Omega$, $\sum_{x\in\Omega} \lambda(x, y) = \nu(y)$. In other words, λ is a joint distribution whose marginal distributions are resp. μ and ν.

Coupling

A *coupling* of μ and ν is a distribution λ on $\Omega \times \Omega$ such that for all $x \in \Omega$, $\sum_{y \in \Omega} \lambda(x, y) = \mu(x)$ and for all $y \in \Omega$, $\sum_{x \in \Omega} \lambda(x, y) = \nu(y)$. In other words, λ is a joint distribution whose marginal distributions are resp. μ and ν.

Lemma (Coupling Lemma)

Let μ and ν be probability distributions on a finite event space Ω, let λ be a coupling of μ and ν, and let $(X, Y) \sim \lambda$ (i.e. (X, Y) is a random variable sampled according to distribution λ). Then $\|\mu - \nu\| \leq \Pr[X \neq Y]$.

Examples

Let $0 \leq p_1 \leq p_2 \leq 1$ and C_1, C_2 be two coins such that C_1 makes a head with probability p_1 and C_2 makes a head with probability p_2.

Examples

Let $0 \leq p_1 \leq p_2 \leq 1$ and C_1, C_2 be two coins such that C_1 makes a head with probability p_1 and C_2 makes a head with probability p_2. What's the advantage to distinguish the two coins ?

Let $0 \leq p_1 \leq p_2 \leq 1$ and C_1, C_2 be two coins such that C_1 makes a head with probability p_1 and C_2 makes a head with probability p_2.
What's the advantage to distinguish the two coins ?
We can couple them:

p_1	C_1 and C_2 make head
$p_2 - p_1$	C_1 makes tail and C_2 makes head
$1 - p_2$	C_1 and C_2 make tail

Examples

Let $0 \leq p_1 \leq p_2 \leq 1$ and C_1, C_2 be two coins such that C_1 makes a head with probability p_1 and C_2 makes a head with probability p_2.
What's the advantage to distinguish the two coins ?
We can couple them:

p_1	C_1 and C_2 make head
$p_2 - p_1$	C_1 makes tail and C_2 makes head
$1 - p_2$	C_1 and C_2 make tail

The advantage is upperbounded by $p_2 - p_1$.

Application of the Coupling Technique

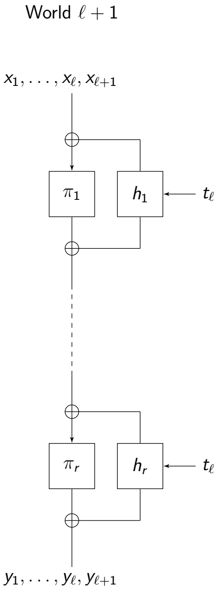

Coupling on the ℓ first queries

- Pick $h_1, \ldots, h_r$ in $\mathcal{H}$.
- Define $h'_1 = h_1, \ldots, h'_r = h_r$.

Application of the Coupling Technique

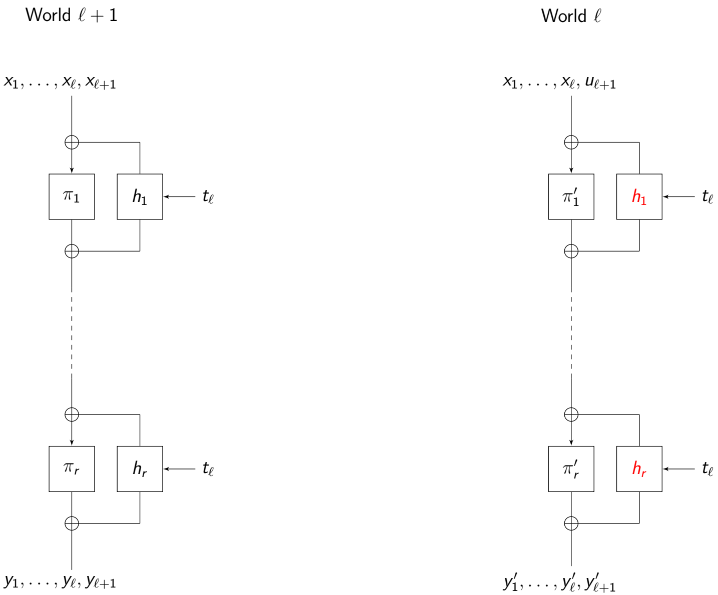

Coupling on the ℓ first queries

- Pick $h_1, \dots, h_r$ in $\mathcal{H}$.
- Define $h'_1 = h_1, \dots, h'_r = h_r$.

Coupling on the ℓ first queries

- Pick $h_1, \ldots, h_r$ in $\mathcal{H}$.
- Define $h'_1 = h_1, \ldots, h'_r = h_r$.
- Pick $\pi_1, \ldots, \pi_r$ uniformly random.

Coupling on the ℓ first queries

- Pick $h_1, \ldots, h_r$ in $\mathcal{H}$.
- Define $h'_1 = h_1, \ldots, h'_r = h_r$.
- Pick $\pi_1, \ldots, \pi_r$ uniformly random.
- For every $i \leq \ell$, π'_1 acts like π_1 when computing x_i.

Coupling on the ℓ first queries

- Pick $h_1, \dots, h_r$ in $\mathcal{H}$.
- Define $h'_1 = h_1, \dots, h'_r = h_r$.
- Pick $\pi_1, \dots, \pi_r$ uniformly random.
- For every $i \leq \ell$, π'_1 acts like π_1 when computing x_i.
- Same process for $\pi'_2, \dots, \pi'_r$.

Coupling on the ℓ first queries

- Pick $h_1, \ldots, h_r$ in $\mathcal{H}$.
- Define $h'_1 = h_1, \ldots, h'_r = h_r$.
- Pick $\pi_1, \ldots, \pi_r$ uniformly random.
- For every $i \leq \ell$, π'_1 acts like π_1 when computing x_i.
- Same process for $\pi'_2, \ldots, \pi'_r$.

$$\Rightarrow \forall i \leq \ell, y'_i = y_i.$$

Application of the Coupling Technique

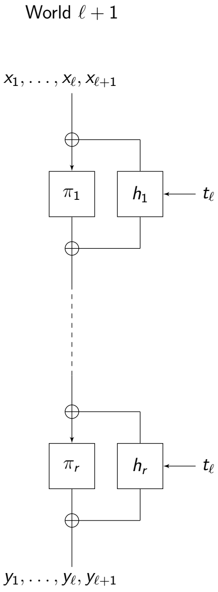

Coupling of the $\ell + 1$−th query

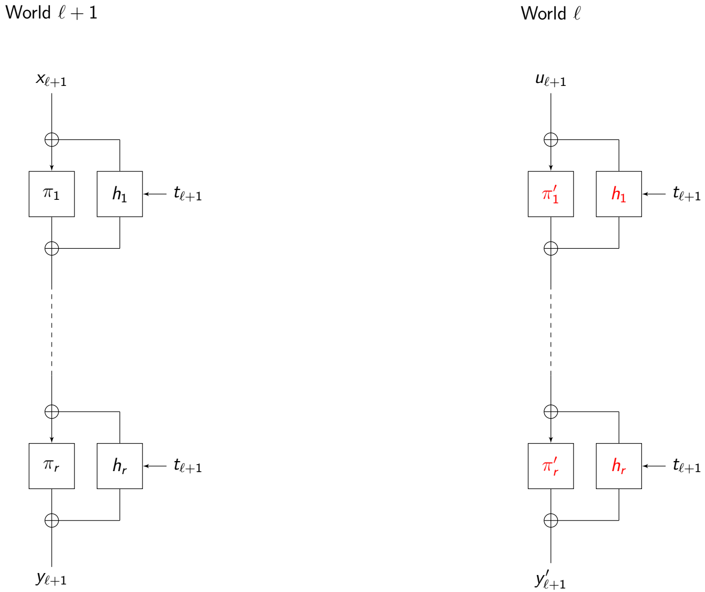

Coupling of the $\ell + 1-$th query

If $\pi_1(x_{\ell+1} \oplus h_1(t_{\ell+1}))$ and $\pi'_1(u_{\ell+1} \oplus h_1(t_{\ell+1}))$ are not already defined, we can couple them by choosing the same randomness for both, we define:

Coupling of the $\ell + 1-$th query

If $\pi_1(x_{\ell+1} \oplus h_1(t_{\ell+1}))$ and $\pi'_1(u_{\ell+1} \oplus h_1(t_{\ell+1}))$ are not already defined, we can couple them by choosing the same randomness for both, we define:

$$\pi'_1(u_{\ell+1} \oplus h_1(t_{\ell+1})) := \pi_1(x_{\ell+1} \oplus h_1(t_{\ell+1})).$$

Coupling of the $\ell + 1-$th query

If $\pi_1(x_{\ell+1} \oplus h_1(t_{\ell+1}))$ and $\pi'_1(u_{\ell+1} \oplus h_1(t_{\ell+1}))$ are not already defined, we can couple them by choosing the same randomness for both, we define:

$$\pi'_1(u_{\ell+1} \oplus h_1(t_{\ell+1})) := \pi_1(x_{\ell+1} \oplus h_1(t_{\ell+1})).$$

If $\pi_1(x_{\ell+1} \oplus h_1(t_{\ell+1}))$ or $\pi'_1(u_{\ell+1} \oplus h_1(t_{\ell+1}))$ is already defined

Coupling of the $\ell + 1-$th query

If $\pi_1(x_{\ell+1} \oplus h_1(t_{\ell+1}))$ and $\pi'_1(u_{\ell+1} \oplus h_1(t_{\ell+1}))$ are not already defined, we can couple them by choosing the same randomness for both, we define:

$$\pi'_1(u_{\ell+1} \oplus h_1(t_{\ell+1})) := \pi_1(x_{\ell+1} \oplus h_1(t_{\ell+1})).$$

If $\pi_1(x_{\ell+1} \oplus h_1(t_{\ell+1}))$ or $\pi'_1(u_{\ell+1} \oplus h_1(t_{\ell+1}))$ is already defined (due to a collision of the form $x_{\ell+1} \oplus h_1(t_{\ell+1}) = x_i \oplus h_1(t_i)$ or $u_{\ell+1} \oplus h_1(t_{\ell+1}) = x_i \oplus h_1(t_i)$

Coupling of the $\ell + 1-$th query

If $\pi_1(x_{\ell+1} \oplus h_1(t_{\ell+1}))$ and $\pi'_1(u_{\ell+1} \oplus h_1(t_{\ell+1}))$ are not already defined, we can couple them by choosing the same randomness for both, we define:

$$\pi'_1(u_{\ell+1} \oplus h_1(t_{\ell+1})) := \pi_1(x_{\ell+1} \oplus h_1(t_{\ell+1})).$$

If $\pi_1(x_{\ell+1} \oplus h_1(t_{\ell+1}))$ or $\pi'_1(u_{\ell+1} \oplus h_1(t_{\ell+1}))$ is already defined (due to a collision of the form $x_{\ell+1} \oplus h_1(t_{\ell+1}) = x_i \oplus h_1(t_i)$ or $u_{\ell+1} \oplus h_1(t_{\ell+1}) = x_i \oplus h_1(t_i)$), we can't couple and we try to couple on the next round.

Probability of not coupling at round 1

The probability for not coupling on the first round is upperbounded by the sum over $i \leq \ell$ of the events

$$x_{\ell+1} \oplus h_1(t_{\ell+1}) = x_i \oplus h_1(t_i) \text{ or } u_{\ell+1} \oplus h_1(t_{\ell+1}) = x_i \oplus h_1(t_i)$$

Probability of not coupling at round 1

The probability for not coupling on the first round is upperbounded by the sum over $i \leq \ell$ of the events

$$x_{\ell+1} \oplus h_1(t_{\ell+1}) = x_i \oplus h_1(t_i) \text{ or } u_{\ell+1} \oplus h_1(t_{\ell+1}) = x_i \oplus h_1(t_i)$$

which is equivalent to $h_1(t_{\ell+1}) \oplus h_1(t_i)$ equals $x_{\ell+1} \oplus x_i$ or $u_{\ell+1} \oplus x_i$.

Probability of not coupling at round 1

The probability for not coupling on the first round is upperbounded by the sum over $i \leq \ell$ of the events

$$x_{\ell+1} \oplus h_1(t_{\ell+1}) = x_i \oplus h_1(t_i) \text{ or } u_{\ell+1} \oplus h_1(t_{\ell+1}) = x_i \oplus h_1(t_i)$$

which is equivalent to $h_1(t_{\ell+1}) \oplus h_1(t_i)$ equals $x_{\ell+1} \oplus x_i$ or $u_{\ell+1} \oplus x_i$.
Since $\max_{x,x',y} \Pr\left[h \leftarrow_{\$} \mathcal{H} \,:\, h(x) \oplus h(x') = y\right] \leq \varepsilon$, the probability of not coupling at round 1 is upperbounded by $\ell \times 2\varepsilon$.

Probability of not coupling at the next rounds

Using the same reasoning, the probability of coupling at each round is upperbounded by $2\ell\varepsilon$ and since each round functions are independent, the probability of coupling nowhere is upperbounded by $(2\ell\varepsilon)^r$.

Probability of not coupling at the next rounds

Using the same reasoning, the probability of coupling at each round is upperbounded by $2\ell\varepsilon$ and since each round functions are independent, the probability of coupling nowhere is upperbounded by $(2\ell\varepsilon)^r$.

$$\sum_{\ell=0}^{q-1}(2\ell\varepsilon)^r \leq \frac{q^{r+1}}{r+1}(2\varepsilon)^r$$

Result

Theorem

Let $\mathcal{K}, \mathcal{T}$ be sets, $E \in \mathrm{BC}(\mathcal{K}, n)$ be a blockcipher, and $\mathcal{H}$ be a ε-AXU$_2$ family of functions from $\mathcal{T}$ to $\{0,1\}^n$. Then one has:

$$\mathbf{Adv}^{\widetilde{\mathrm{ncpa}}}_{\mathtt{CLRW}^{r},E,\mathcal{H}}(q,\tau) \leq r \cdot \mathbf{Adv}^{\mathrm{ncpa}}_{E}(q, \tau + rqT) + \frac{q^{r+1}}{r+1}(2\varepsilon)^r$$

where T is the time to compute E or E^{-1}.

From NCPA to CCA

To obtain CCA security, we show that composing two NCPA-secure tweakable blockciphers (with the same tweak) yields a CCA-secure tweakable blockcipher.

From NCPA to CCA

To obtain CCA security, we show that composing two NCPA-secure tweakable blockciphers (with the same tweak) yields a CCA-secure tweakable blockcipher.

Applying this result to the $\texttt{CLRW}^{r,E,\mathcal{H}}$ construction yield the following result.

Result

Theorem

Let $\mathcal{K}, \mathcal{T}$ be sets, $E \in \mathrm{BC}(\mathcal{K}, n)$ be a blockcipher, and $\mathcal{H}$ be a ε-AXU$_2$ family of functions from $\mathcal{T}$ to $\{0,1\}^n$. Then one has:

$$\mathbf{Adv}^{\widetilde{\mathrm{cca}}}_{\mathtt{CLRW}^r, E, \mathcal{H}}(q, \tau) \leq r \cdot \mathbf{Adv}^{\mathrm{cca}}_{E}(q, \tau + rqT) + \frac{4\sqrt{2}}{\sqrt{r+2}} q^{(r+2)/4} (2\varepsilon)^{r/4}$$

where T is the time to compute E or E^{-1}.

Open question: Prove security up to $2^{\frac{r}{r+1}n}$ queries against CCA attacks.

Thank you

Any question ?