On the Construction of Pseudo-Random Permutations: Luby-Rackoff Revisited

Moni Naor & Omer Reingold

Department of Applied Mathematics
and Computer Science,
The Weizmann Institute, Rehovot, Israel.

WEIZMANN
INSTITUTE
OF SCIENCE

Presentation Outline

Introduction:

- Block Ciphers, Pseudo-Random Permutations & Functions.
- The Luby-Rackoff Construction.
- Pair-Wise Independent Permutations.

This Work:

- Main Construction: Reducing the Cryptographic Work of the LR-Construction.
- Reducing the Distinguisher's Success Probability.
- From Pseudo-Random Functions with Small Input Size to Pseudo-Random Permutations with Large Input Size.

Block-Ciphers

Private-key encryption schemes such that:
The encryption of every plain-text block is a *single* cipher-text block *of the same length*.

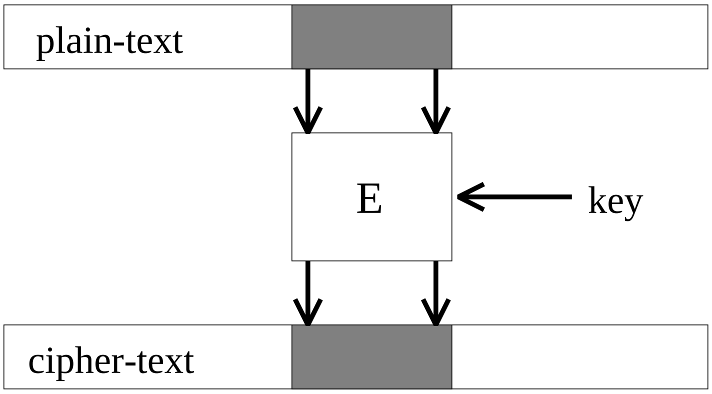

Advantages:

- Saves up on memory and communication bandwidth.
- Easy to incorporate within existing systems.

Main Disadvantage:

Every block is always encrypted in the same way.

Important Example:

Data Encryption Standard (DES).

Pseudo-Random Permutations [LR86]

Distribution of Permutations that is

- **Efficient:** easy to sample, easy to compute.
- **Pseudo-Random:** cannot be *efficiently distinguished* from the uniform distribution with *bias* $\geq \frac{1}{poly}$

The **Distinguisher** Power:

- Pseudo-Random Permutations:

Chosen Plain-Text Attack

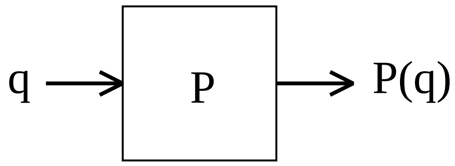

- Strong Pseudo-Random Permutations:

Chosen Plain-Text & *Chosen Cipher-Text* Attack

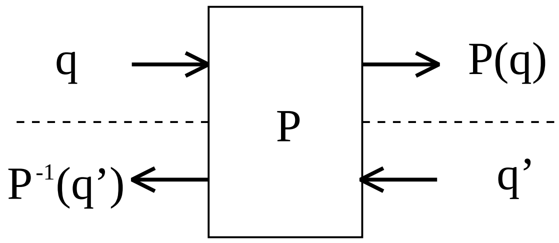

... Pseudo-Random Permutations [LR86]

An **Efficient** Distinguisher:

Polynomial Time Machine or Polynomial Size Circuit.

The **Bias** M achieves in distinguishing between two distributions F and G is

- for an oracle-machine of the first type:

$$|Prob_{f \in F}[M^f \ outputs \ 1] - Prob_{g \in G}[M^g \ outputs \ 1]|$$

- for an oracle-machine of the second type:

$$|Prob_{f \in F}[M^{f,f^{-1}} \ outputs \ 1] - Prob_{g \in G}[M^{g,g^{-1}} \ outputs \ 1]|$$

Pseudo-Random Functions [GGM84]

Distribution of Functions that is

- **Efficient:** easy to sample, easy to compute.
- **Pseudo-Random:** cannot be efficiently *distinguished* from the uniform distribution with bias $\geq \frac{1}{poly}$

The **Distinguisher** Power:

Black-Box Attack $\qquad q \longrightarrow \boxed{F} \longrightarrow F(q)$

Constructions based on other cryptographic primitives:

- Original construction, based on any pseudo-random *generator* [Goldreich, Goldwasser & Micali '84].
- A parallel construction based on pseudo-random *synthesizers* [Naor & Reingold '95].

All the constructions involve "heavy" computations.

Feistel (DES-Like) Permutations

A Feistel permutation for a function $f : \{0,1\}^n \mapsto \{0,1\}^n$ is a permutation on $\{0,1\}^{2n}$ defined by $D_f(L,R) \stackrel{\text{def}}{=} (R, L \oplus f(R))$, where $|L| = |R| = n$.

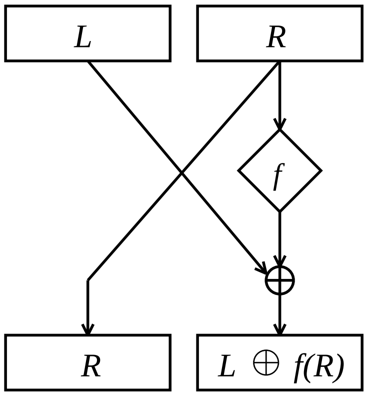

Feistel permutations are as easy to invert as to compute: $D_f^{-1}(L,R) = (R \oplus f(L), L)$.

DES is a composition of 16 Feistel permutations for functions that are determined by the 56 bits key.

Can a composition of (a few) Feistel permutations be pseudo-random?

Yes! [Luby & Rackoff]

The Luby-Rackoff Construction

Composing four (or three) Feistel permutations using pseudo-random functions $(f_1, \ldots, f_4)$.

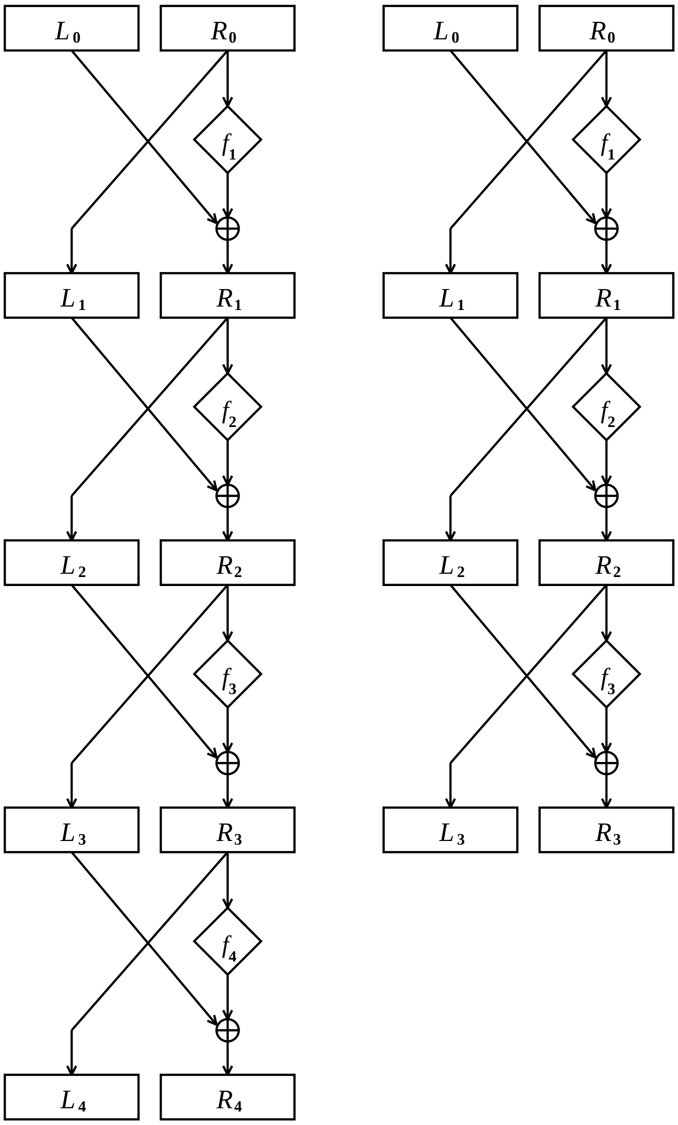

Pair-Wise Independent Permutations

A family of functions $\mathcal{F} = \{f : A \mapsto B\}$ is *pair-wise independent* if $\forall x \neq y \in A$ the value $< f(x), f(y) >_{f \in \mathcal{F}}$ is uniformly distributed over B^2 [Carter & Wegman '79 & '81].

Similarly: a family *of permutations* $\mathcal{F} = \{f : A \mapsto A\}$ is pair-wise independent if $\forall x \neq y \in A$ the value $\langle f(x), f(y) \rangle_{f \in \mathcal{F}}$ is uniformly distributed over $\{(a,b) \in A^2 : a \neq b\}$.

Example ($A = \{0,1\}^n$):
$\{f_{a,b} : a, b \in \{0,1\}^n \ and \ a \neq 0\}$, where
$f_{a,b}(x) \stackrel{\text{def}}{=} a \cdot x + b$ with operations over $GF(2^n)$.

A substantial amount of research deals with definitions of weaker function families and their construction [CW79, WC81, Sti94, Kra94, Rog95, ...].

Related Work

- Alternative proofs of the LR-Construction [Patarin '90, Maurer '92].

- Variants of the construction [Sch, Ohi, ZMI, Rue, Pie, Pat90, Pat91, SadPie91, SadPie92, Pat92].

- A hash function can replace the pseudo-random function in the first round of the three-round LR-Construction [Lucks '96].

- Several pseudo-random permutations from a single one [Even & Mansour '91].
  Protecting against exhaustive key-search attacks [Kilian & Rogaway '96].

- An "LR-Construction" with six rounds is more secure [Patarin '91].

This Work

Main Observation: The different rounds of the LR-Construction serve significantly different roles. The first and last rounds can be replaced by pair-wise independent permutations. Using this we can:

- Reduce the "cryptographic work" and the "cryptographic depth" by a factor of 2.
- Simplify the proof of security.
- Derive generalizations of the construction with almost identical proofs of security:
  1. With $t \geq 2$ (generalized) rounds the adversary's advantage can be reduced from $\frac{m^2}{2^{\ell/2}}$ to $\frac{m^2 \cdot t/2}{2^{(1-1/t)\ell}}$ where the permutation is on ℓ bits and m is the number of queries.
  2. Pseudo-random functions on *small* blocks can be used for a simple construction of strong pseudo-random permutation on *large* blocks.

Main Construction

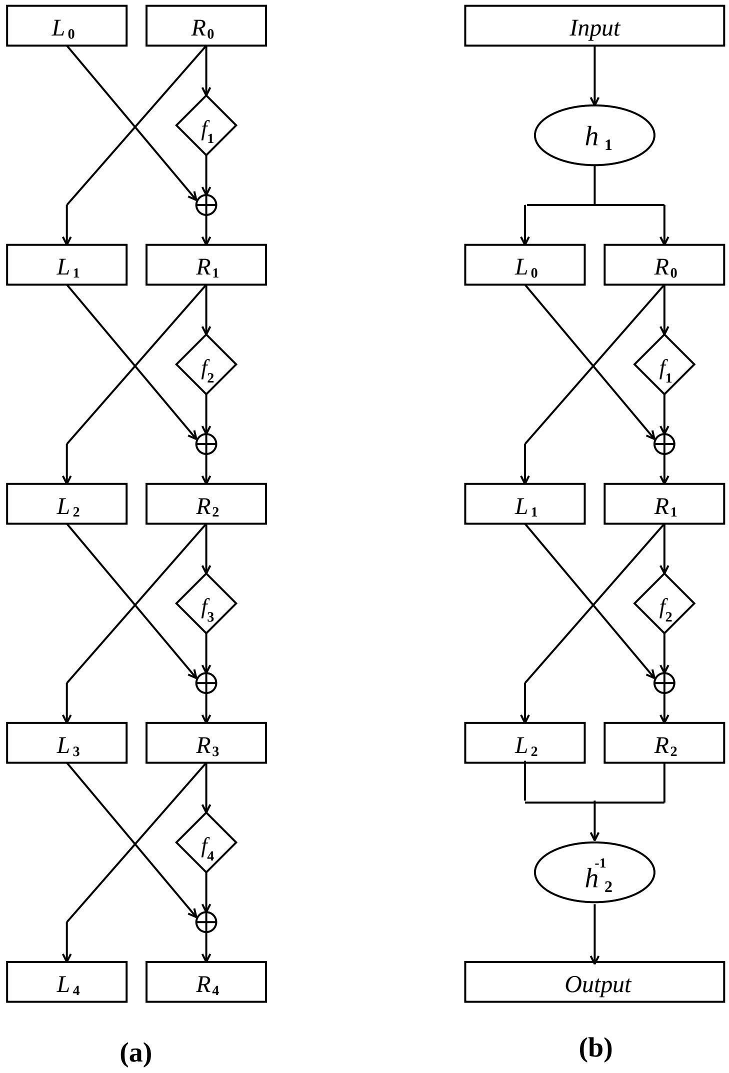

(a) Original LR-Construction. (b) Revised Construction. h_1 & h_2 are pair-wise independent permutations.

The Security of the Construction

Let P be a permutation on $2n$ bits as in the revised construction, where f_1 & f_2 are *random functions*.

Let R be a random permutation on $2n$ bits.

Theorem 1. *For* any *oracle machine* M *(not necessarily efficient) that makes at most* m *queries the bias* M *achieves in distinguishing between* R & P *is at most* $\frac{m^2}{2^n} + \frac{m^2}{2^{2n}}$

Reminder: The bias is defined to be
$|Prob[M^{P,P^{-1}}\ outputs\ 1] - Prob[M^{R,R^{-1}}\ outputs\ 1]|$

Remarks on Security Theorem

- Got indistinguishability in an Information-Theoretic sense. *Computational* indistinguishability of the construction, in case f_1 & f_2 are *pseudo-random* functions, follows easily.
- Theorem 1 and its proof hold for the original LR-Construction.

Intuition: Two-Round Construction P'

Let P' be a two-round permutation:

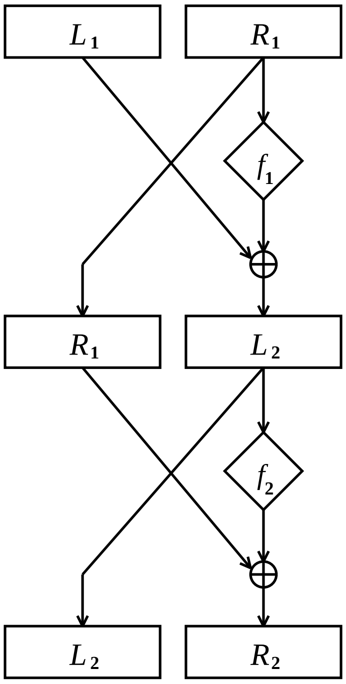

Fix L_1, R_1, L_2, R_2. $P'(L_1, R_1) = (L_2, R_2)$ iff

$$f_1(R_1) = L_1 \oplus L_2 \ \& \ f_2(L_2) = R_1 \oplus R_2.$$

The constraints are independent $\Rightarrow$ probability of this event is 2^{-2n}

... Intuition

Consider a sequence $\{(L_1^i, R_1^i), (L_2^i, R_2^i)\}_{i=1}^m$

If for every $i \neq j$ both $R_1^i \neq R_1^j$ and $L_2^i \neq L_2^j$ then:

$$Prob[\forall i\ P'(L_1^i, R_1^i) = (L_2^i, R_2^i)] = 2^{-2nm}$$

Conclusion: Although P' is not pseudo-random [Luby-Rackoff], it is secure against a "random" attack.

Proof of Security

Assumptions:

- The distinguisher, M, is *deterministic* since M has unlimited computational powers.
- M makes exactly m different queries.

Let F be the permutation accessible to M (either P or R).

Definition 1. [The Transcript T]
For the ith query M makes, define the query-answer pair $< x_i, y_i >$, where:

- *M asked for $F(x_i)$ and got y_i as an answer, or*
- *M asked for $F^{-1}(y_i)$ and got x_i as an answer.*

T is the sequence $\{< x_1, y_1 >, ..., < x_m, y_m >\}$.

Notice:

- T is fully determined by the answers M gets.
- M's output is a deterministic function of T.

... Proof of Security

Consider an additional distribution on the answers to M's queries - *the random process* $\tilde{R}$ that on the ith query of M answers as follows:

1. If M asks for $F(x)$ and for some $1 \leq j < i$ the jth query-answer pair is $< x, y >$, then $\tilde{R}$'s answer is y.

2. If M asks for $F^{-1}(y)$ and for some $1 \leq j < i$ the jth query-answer pair is $< x, y >$, then $\tilde{R}$'s answer is x.

3. If neither 1 nor 2 holds, then $\tilde{R}$'s answer is a uniformly chosen $2n$-bits string.

Note: $\tilde{R}$ may provide answers that are not consistent with *any* permutation:

- T is **inconsistent** if it contains two answer-query pairs of the form (x, y) and (x', y) for $x \neq x'$ or (x, y) and (x, y') for $y \neq y'$.
- Otherwise, T is **consistent.**

… Proof of Security

$\tilde{R}$ is "close" to R:

Proposition 2.

$$\begin{aligned} & |Prob[M^{\tilde{R}} \; outputs \; 1] - Prob[M^{R,R^{-1}} \; outputs \; 1]| \\ \leq \quad & Prob_{\tilde{R}}[T \; is \; inconsistent] \end{aligned}$$

Proposition 3.

$$Prob_{\tilde{R}}[T \; is \; inconsistent] \leq \binom{m}{2} 2^{-2n} < \frac{m^2}{2^{2n+1}}$$

It remains to bound:

$$|Prob[M^{\tilde{R}} \; outputs \; 1] - Prob[M^{P,P^{-1}} \; outputs \; 1]|$$

... Proof of Security - BAD

A Methodological Assumption: h_1 & h_2 are chosen *even when* M*'s queries are answered by* $\tilde{R}$. In this case T is independent of h_1 & h_2

Definition 2. [BAD] *If* T *is consistent, we say that* T *is BAD for* $< h_1, h_2 >$ *if there exist* $1 \leq i < j \leq m$ *such that either* $h_1(x_i)_{|_r} = h_1(x_j)_{|_r}$ *or* $h_2(y_i)_{|_l} = h_2(y_j)_{|_l}$
(where $(L, R)_{|_l} = L$ *&* $(L, R)_{|_r} = R$*)*

Let BAD be the event "either T *is inconsistent or* T *is BAD for* $< h_1, h_2 >$*"*

Since h_1 & h_2 are pair-wise independent and from Proposition 3 we get:

Proposition 4. $Prob_{\tilde{R}}[BAD] < \frac{m^2}{2^n} + \frac{m^2}{2^{2n+1}}$

... Proof of Security - Key Lemma

Lemma 5. *Let S be any possible value of T,*

$$Prob_P[T = S \wedge (\neg BAD)] = Prob_{\tilde{R}}[T = S \wedge (\neg BAD)]$$

Proof. Can assume that $S = \{< x_1, y_1 >, ..., < x_m, y_m >\}$ is consistent.

It is enough to show that for every h_1 & h_2, such that S is not BAD for $< h_1, h_2 >$:

$$\begin{aligned} & Prob_P[T = S \wedge (\neg BAD) | h_1 \,\&\, h_2 \text{ were chosen}] \\ = \;& 2^{-2nm} \\ = \;& Prob_{\tilde{R}}[T = S \wedge (\neg BAD) | h_1 \,\&\, h_2 \text{ were chosen}] \end{aligned}$$

The second equality holds since each answer of $\tilde{R}$ is independent and uniform $2n$-bits string; $\tilde{R}$ should answer "correctly" on all of the m queries.

... Key Lemma

For the first equality recall the two-round construction and the definition of BAD:

If M's queries are answered by P & P^{-1}, then $T = S$ iff $\forall i, y_i = P(x_i)$.

Let $(L_1^i, R_1^i) = h_1(x_i)$ and $(L_2^i, R_2^i) = h_2(y_i)$.
By the definition of P, $y_i = P(x_i)$ iff

$$f_1(R_1^i) = L_1^i \oplus L_2^i \;\&\; f_2(L_2^i) = R_1^i \oplus R_2^i.$$

For every $1 \leq i < j \leq m$ both $R_1^i \neq R_1^j$ and $L_2^i \neq L_2^j$ (o.w. S is BAD for $< h_1, h_2 >$).

Since f_1 & f_2 are random we can conclude:

$$Prob_P[T = S | h_1 \ and \ h_2 \ were \ chosen] = 2^{-2nm}$$

$\square$

... Proof of Security (the end)

Proof. (of Theorem 1) By summing Lemma 5 on all possible values of T:

$$\begin{aligned} & Prob_P[\neg BAD] = Prob_{\tilde{R}}[\neg BAD] \\ \Rightarrow \quad & Prob_P[BAD] = Prob_{\tilde{R}}[BAD] \end{aligned}$$

Summing Lemma 5 on all T's for which M outputs 1:

$$\begin{aligned} & Prob[(M^{P,P^{-1}} \; outputs \; 1) \wedge (\neg BAD)] \\ = \quad & Prob[(M^{\tilde{R}} \; outputs \; 1) \wedge (\neg BAD)] \end{aligned}$$

Therefore,

$$\begin{aligned} & |Prob[M^{P,P^{-1}} \; outputs \; 1] - Prob[M^{\tilde{R}} \; outputs \; 1]| \\ \leq \quad & Prob_{\tilde{R}}[BAD] \\ < \quad & \frac{m^2}{2^n} + \frac{m^2}{2^{2n+1}} \end{aligned}$$

The Theorem follows by Propositions 2 & 3. $\Box$

Relaxing The Construction

The conditions of Theorem 1 can be relaxed:

1. A single function f can replace both f_1 & f_2

2. The real requirements h_1 & h_2 have to obey are that for every $x \neq y$:
   $Prob[h_1(x)_{|_r} = h_1(y)_{|_r}] \leq 2^{-n}$ and
   $Prob[h_2(x)_{|_l} = h_2(y)_{|_l}] \leq 2^{-n}$
   (where $(L, R)_{|_l} = L$ & $(L, R)_{|_r} = R$)

   In particular, h_1 (or h_2^{-1}) can be a *Feistel permutation* D_g, where for every $x \neq y$ and every z $Prob[g(x) \oplus g(y) = z] \leq 2^{-n}$

   Importance:

   - Get the original LR-Construction as a special case.
   - Can use more efficient hash functions.

Reducing the Distinguishing Probability

Motivation:
Assume we want a pseudo-random permutation on *few* bits. For example, in order to reduce the size of its hardware-circuit.

For the LR-Construction of a permutation on ℓ bits the upper bound on the distinguisher's success is $\approx \frac{m^2}{2^{\ell/2}}$

This bound is tight (up to a constant):
Using $O(2^{\ell/4})$ queries the distinguisher has a very high success probability [Patarin '91].

$\Rightarrow$ Cannot use this construction if ℓ is small (e.g. 80).

Generalized Construction:
The first and last rounds as before.
The two Feistel permutations are replaced with t *generalized Feistel permutations*.

The distinguisher's success probability for the adjusted construction is bounded by $\approx \frac{m^2 \cdot t/2}{2^{(1-1/t)\ell}}$

Generalized Feistel Permutations

The Feistel permutation for a function $f : \{0,1\}^{\ell-\ell/t} \mapsto \{0,1\}^{\ell/t}$, is defined by:
$D_f(L,R) \stackrel{\text{def}}{=} (R, L \oplus f(R))$, where $|L| = \ell/t$ and $|R| = \ell - \ell/t$.

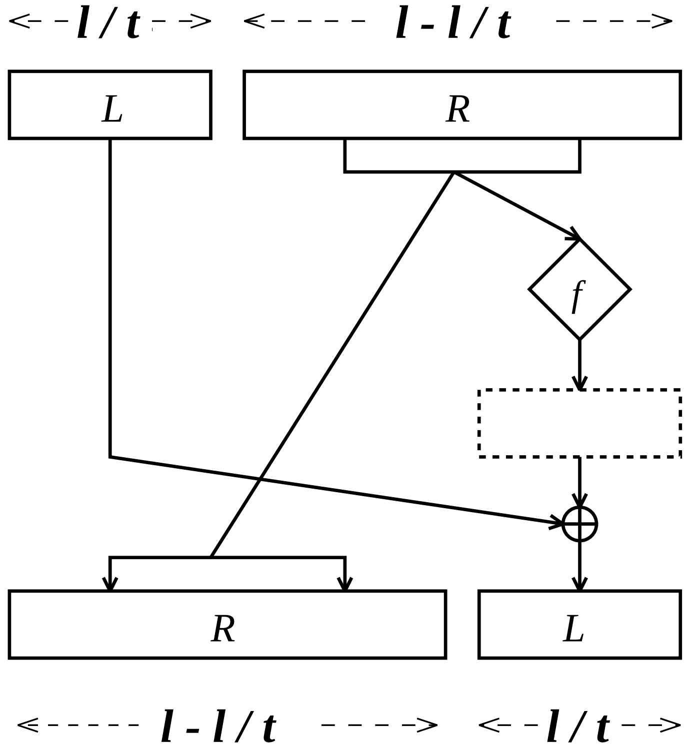

Each one of the t rounds will consist of such a permutation for a pseudo-random function f.

The t-Round Construction for $t = 3$

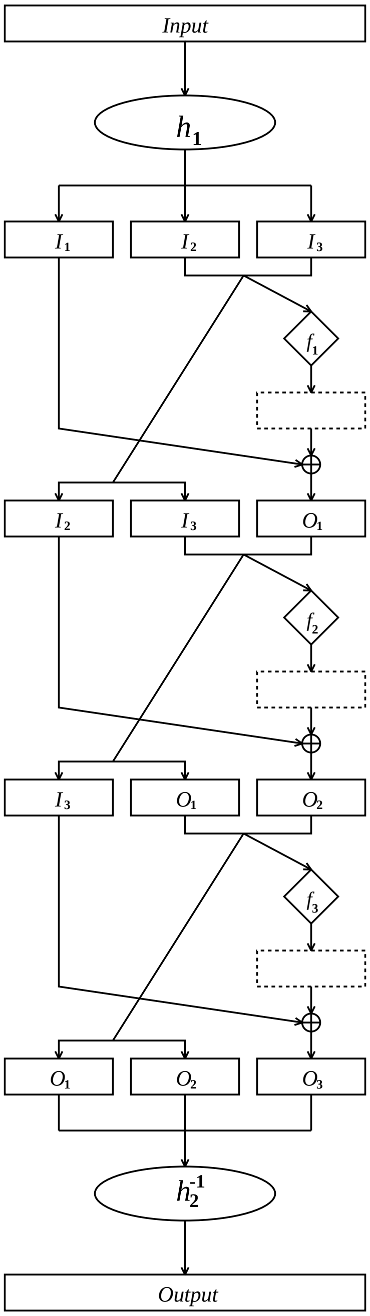

Proof Intuition

Let P be the permutation constructed for $t = 3$. Consider a sequence $\{< x_1, y_1 >, ..., < x_m, y_m >\}$ and define:
$(I_1^i, I_2^i, I_3^i) = h_1(x_i)$ & $(O_1^i, O_2^i, O_3^i) = h_2(y_i)$.

$P(x_i) = y_i$ iff

$f_1(I_2^i, I_3^i) = I_1^i \oplus O_1^i$,

$f_2(I_3^i, O_1^i) = I_2^i \oplus O_2^i$ and

$f_3(O_1^i, O_2^i) = I_3^i \oplus O_3^i$

Therefore, the **BAD** case is when for some $i \neq j$:

$(I_2^i, I_3^i) = (I_2^j, I_3^j)$ or

$(I_3^i, O_1^i) = (I_3^j, O_1^j)$ or

$(O_1^i, O_2^i) = (O_1^j, O_2^j)$

...Proof Intuition

If BAD does not happen then

$$Prob[\forall i \ P(x_i) = (y_i)] = 2^{-\ell \cdot n}$$

and the probability of BAD is:

$$Prob_{h_1,h_2}[BAD] < \frac{3m^2/2}{2^{2\ell/3}}$$

A Simple Construction of Pseudo-Random Permutations on $2\ell \cdot n$ Bits Using Pseudo-Random Functions on n Bits.

Example ($\ell = 3$):

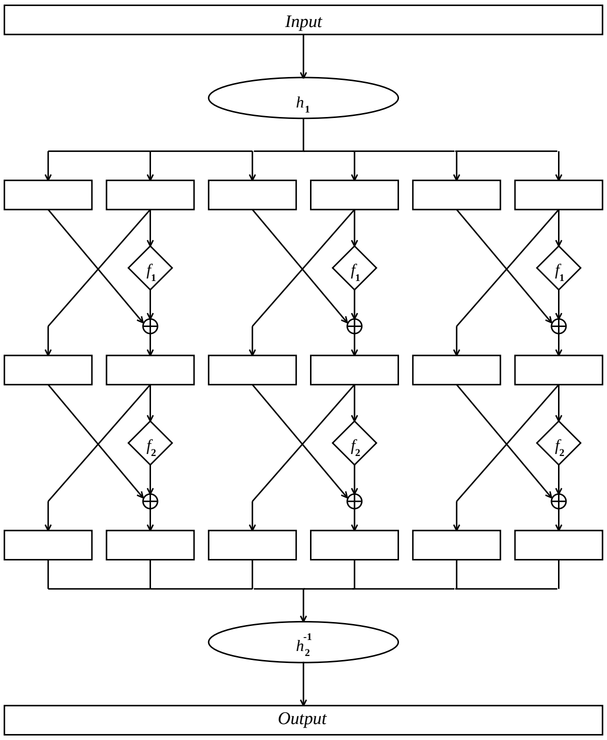

The distinguisher's success probability is bounded by $\frac{m^2 \cdot \ell^2}{2^n} + \frac{m^2}{2^{2n\ell}}$

Conclusion & Further Work

- The revised construction achieves optimal "cryptographic work": the total number of bits on which the cryptographic functions are applied is exactly the number of bits in the input.

  Therefore, *efficient* block-ciphers may be obtained by constructing efficient pseudo-random functions.

- The "cryptographic depth" however, is twice the depth of the cryptographic functions.
  Can we do better?

  Reducing the depth is even more significant in the case of the t-round construction.

- Can we decrease the possible bias of the distinguisher to an arbitrarily small value? (The t-round construction gets us as far as $\approx \frac{\ell \cdot m^2}{2^\ell}$)
  Especially interesting when we think of the LR-Construction as a constructions of k-wise δ-dependent permutations.