How to Encipher Messages on a Small Domain

Deterministic Encryption and the Thorp Shuffle

Ben Morris
University of California, Davis
Dept of Mathematics

Phil Rogaway **Till Stegers**
University of California, Davis
Dept of Computer Science

CRYPTO 2009 — August 18, 2009

ℓ

How to encipher a CCN?

More generally,

How to encipher $\{0,1,\ldots, N\text{-}1\}$?

A special case of ***Format-Preserving Encryption*** **(FPE)** [Brightwell, Smith 97; Spies 08; Bellare, Ristenpart, R, Steger 09]

PRF		PRP
$F: \mathcal{K} \times \{0,1\}^{128} \to \{0,1\}^{128}$	➡	$E: \mathcal{K} \times \{0,1,\ldots, N\text{-}1\} \to : \{0,1,\ldots, N\text{-}1\}$

Known technique	Limitation
• Balanced Feistel [Luby, Rackoff 88; Maurer, Pietrzak 03; Patarin 04] • Benes construction [Aiello, Venkatesan 96; Patarin 08] • Feistel adapted to $Z_a ´ Z_b$ [Black Rogaway 02]	Poor proven bounds for small N
• Induced ordering on $\mathrm{AES}_K(0),\ldots,\mathrm{AES}_K(N-1)$ • "Knuth shuffle"	Preprocessing time $\Omega(N)$
• Cycle walking [Folklore; Black Rogaway02]	For enciphering on $\mathcal{X} \subseteq \mathcal{M}$ when $\lvert\mathcal{X}\rvert / \lvert\mathcal{M}\rvert$ is reasonably large
• *De novo* constructions [Schroeppel 98] • *Ad hoc* modes [FIPS 74: 1981, Brightwell, Smith 97; Mattsson 09]	Provable security not possible
• Wide-block modes [Naor, Reingold 99; Halevi 04]	Starts beyond blockcipher's blocksize
• Granboulan-Pornin construction [GP 07]	Very inefficient

What's wrong with balanced Feistel?

$N = 2^n$

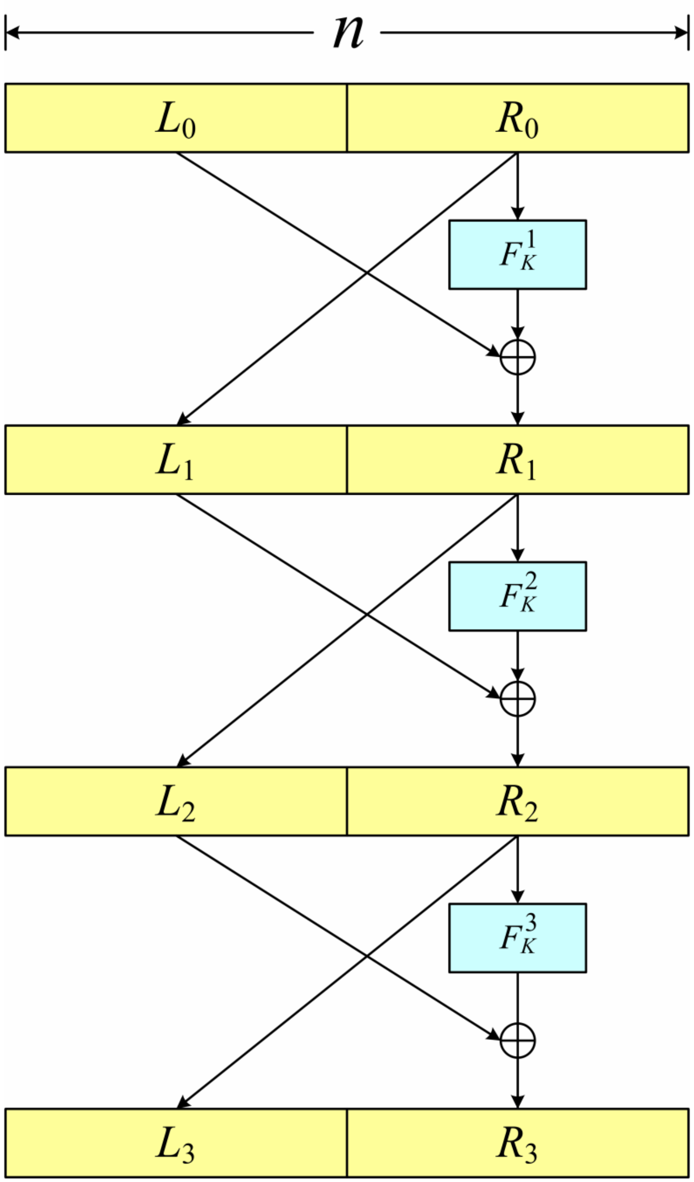

In practice, probably **nothing**.
But, information theoretically,
it only tolerates $2^{n/2}$ queries

Approximate security bounds

[Luby, Rackoff 88] (3 and 4 rounds) $2^{n/4}$

[Maurer, Pietrzak 03] (R rounds) $2^{n/2 - 1/R}$

[Patarin 04] (asymptotic) $2^{n/2 - \varepsilon}$

Attacks

For constant rounds $2^{n/2}$

For R rounds $2^{n/2 + \lg R}$

Encrypting by shuffling

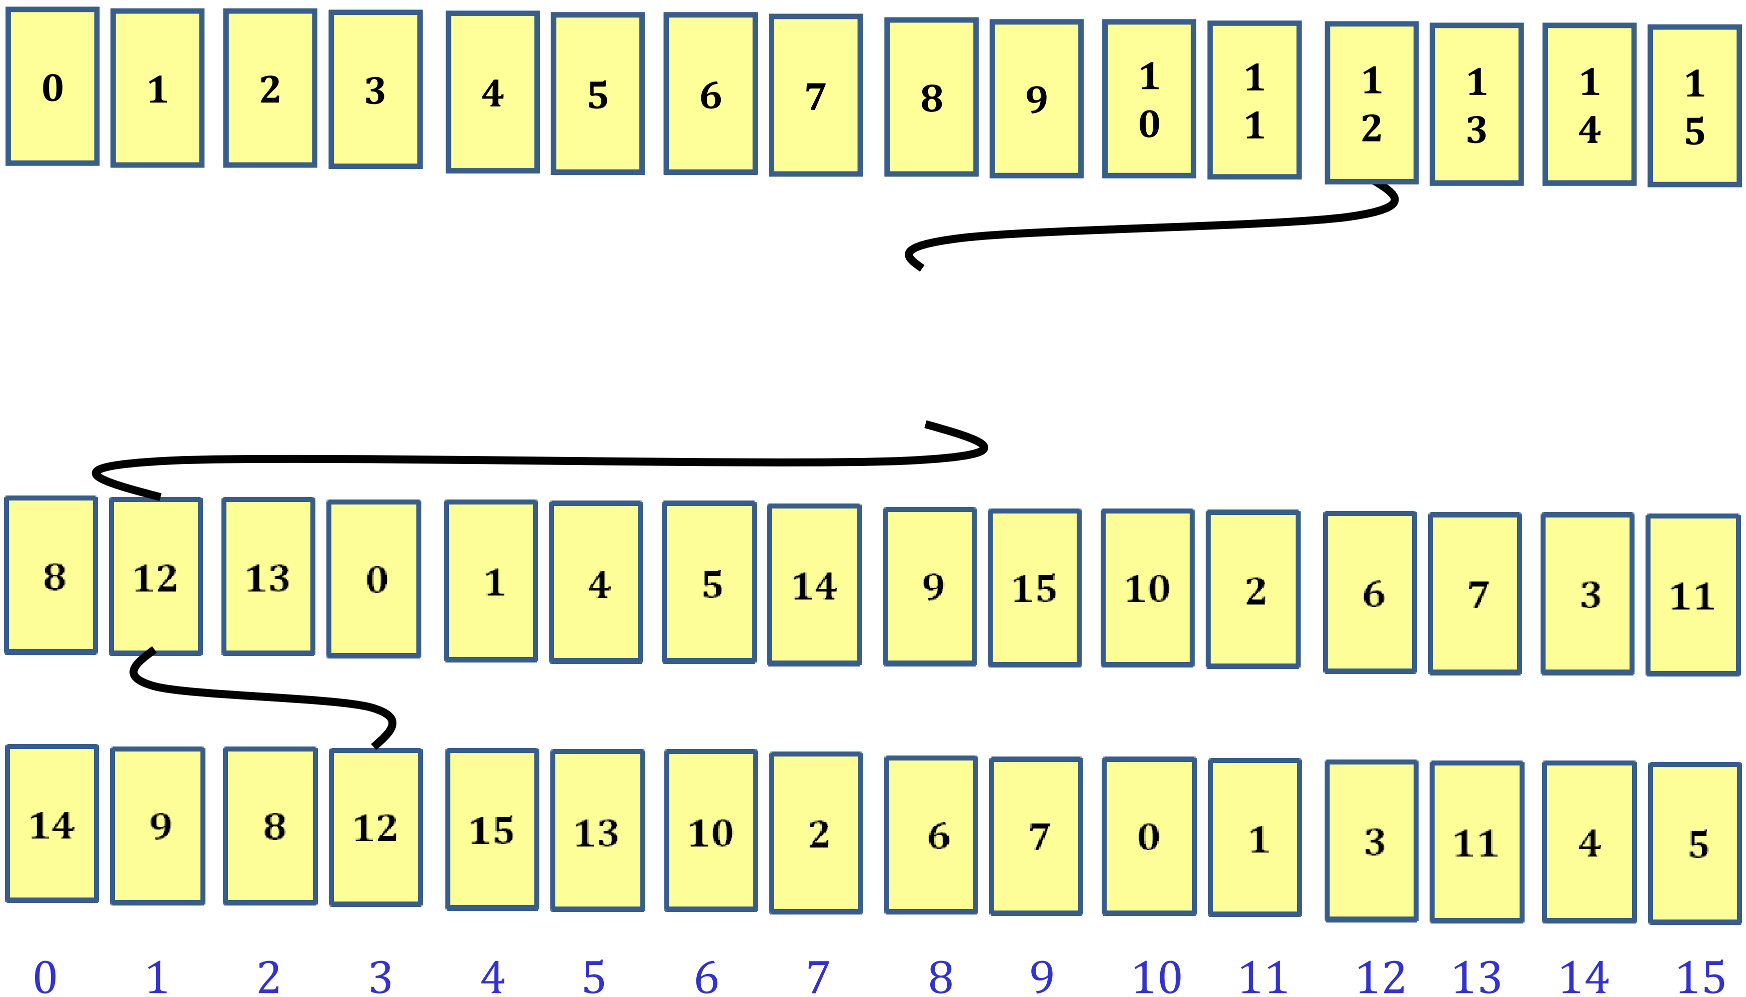

[Naor ~1989] An **oblivious** shuffle: you can follow the path of a card without attending to the other cards. The riffle shuffle is **not** oblivious. The **Thorp shuffle** is.

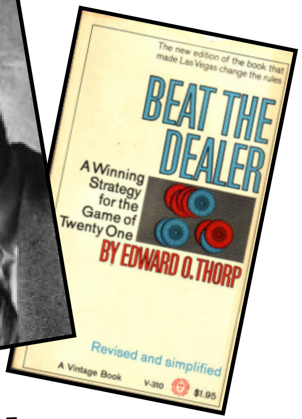

Edward Thorp

[Thorp 73]

Thorp Shuffle

Th[N, R]

To shuffle a deck of N cards (N even):

For round r = 1, 2, …, R do

- Cut the deck exactly in half
- Using a fair coin toss c, drop left-then-right (c=0) or right-then-left (c=1)

One round of the Thorp shuffle

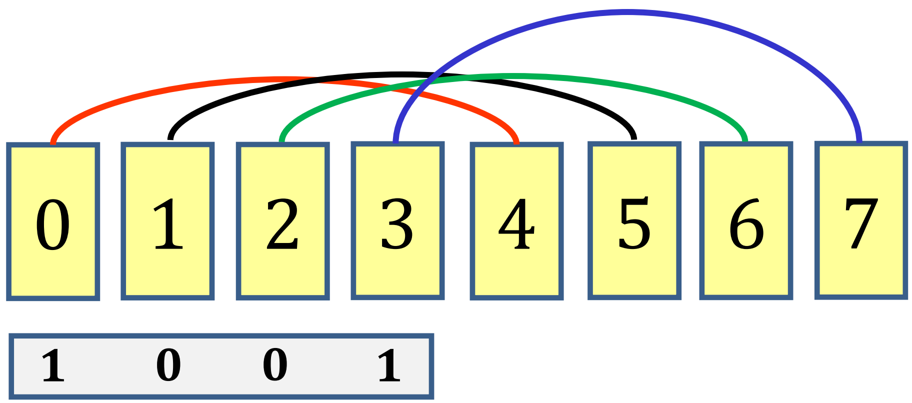

1. Cards at positions x and $x + N/2$ are said to be **adjacent**

2. Flip a coin for each pair of adjacent cards

3. The coins indicate if adjacent cards get moved

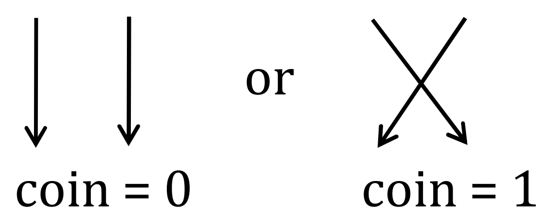

Thorp shuffle = maximally unbalanced Feistel

when $N = 2^n$

At round r, move the card at position $x \in \{0, \ldots, N\text{-}1\}$ to position

$$\begin{cases} 2x + F_K(r, x) & \textbf{if } x < N/2 \\ 2(x - N/2) + (1 - F_K(r, x - N/2)) & \textbf{otherwise} \end{cases}$$

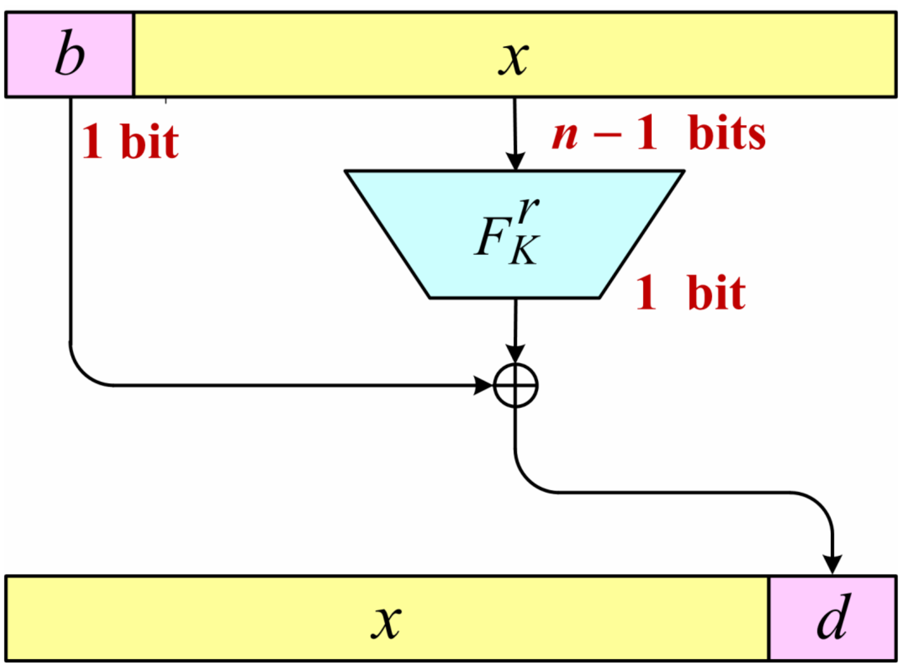

equivalent

Measuring adversarial success

$$\boxed{E = \mathrm{Th}[N, R]}$$

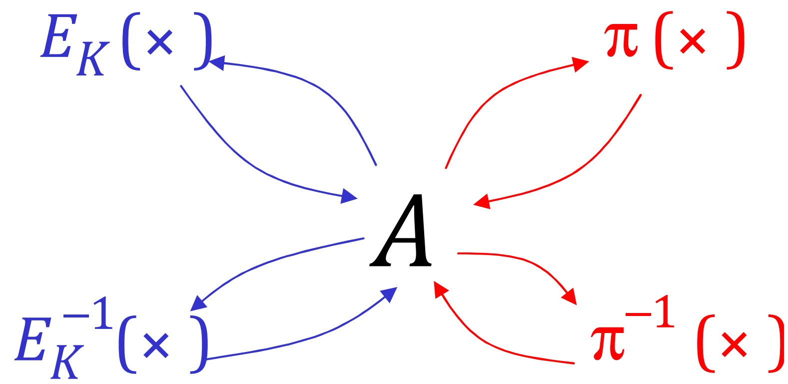

strong PRP

$$\mathbf{Adv}^{\mathrm{cca}}_{N,R}(q) = \max_{A \in \mathrm{CCA}(q)} \Pr[A^{E_K E_K^{-1}} \rightarrow 1] - \Pr[A^{\pi\,\pi^{-1}} \rightarrow 1]$$

nonadaptive PRP

$$\mathbf{Adv}^{\mathrm{ncpa}}_{N,R}(q) = \max_{A \in \mathrm{NCPA}(q)} \Pr[A^{E_K} \rightarrow 1] - \Pr[A^{\pi} \rightarrow 1]$$

What is Known?

$N = 2^n$

For $\boxed{q = N,}$ $\quad \mathbf{Adv}^{\mathrm{ncpa}}_{N,R}(q) \le 2^{-r}$

if $R = O(r \log^{44} N)$ [Morris 05]

$R = O(r \log^{19} N)$ [Montenegro, Tetali 06]

$R = O(r \log^{4} N)$ [Morris 08]

If $\boxed{R = n,}$ $\quad \mathbf{Adv}^{\mathrm{cca}}_{N,R}(q) \le (n+1)\dfrac{q^2}{N}$

(security to about $N^{1/2}$ queries) [Naor, Reingold 99]

(throw in pairwise independent permutations, too)

Main result — Thorp shuffle — CCA

Theorem Let $N = 2^n$ and R=4nr (ie, 4r passes).

$$\mathbf{Adv}^{\text{cca}}_{N,R}(q) \le \frac{2q}{r+1} \left(\frac{4qn}{N} \right)^r$$

Can tolerate
$q = N^{1-1/r}$
queries with
4r passes.

Unbalanced Feistel
provably stronger
than balanced Feistel

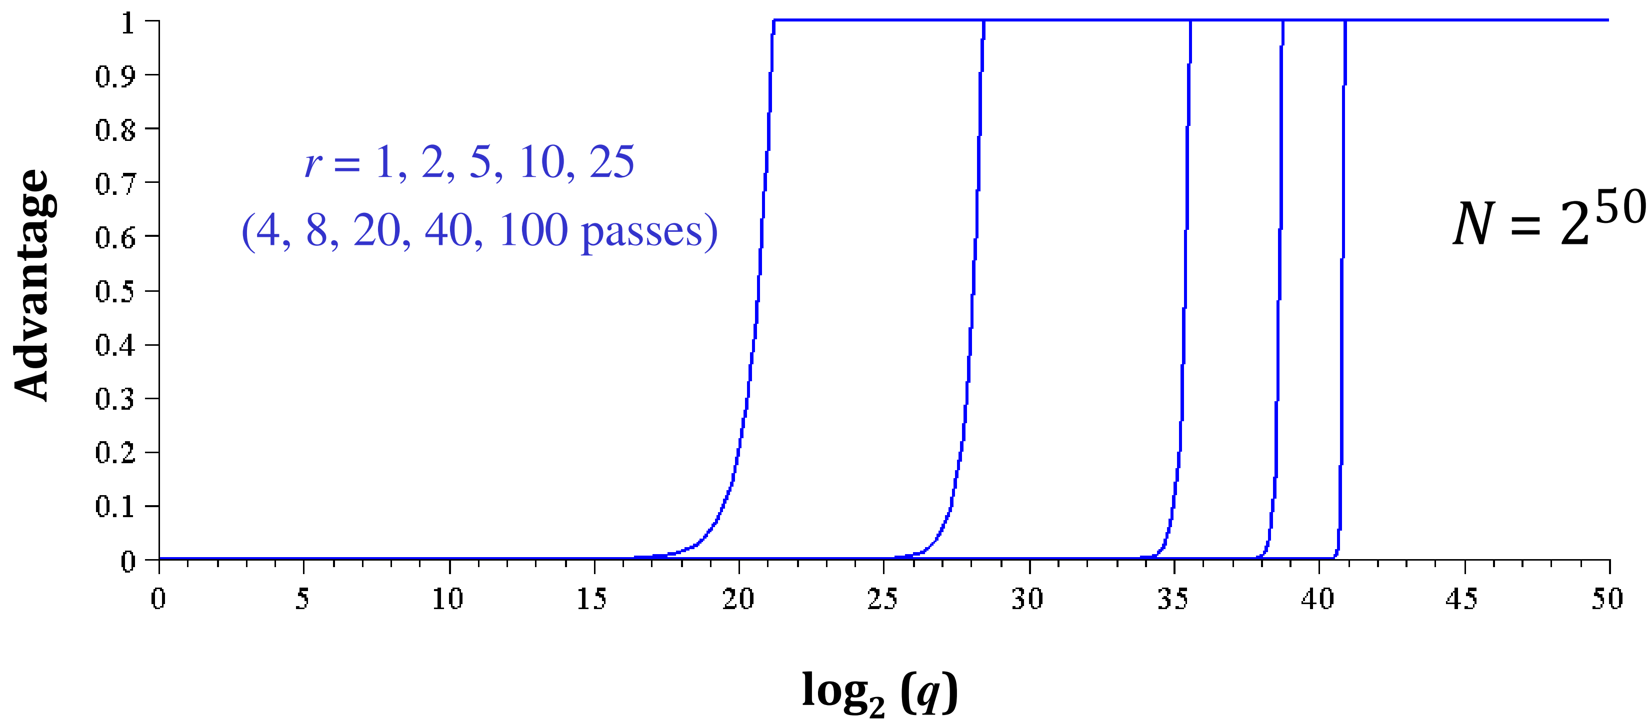

Proving CCA security

1. Prove **NCPA security** of the "projected Thorp shuffle" (and its inverse) using a **coupling argument**

2. Conclude **CCA security** using a wonderful theorem from [Maurer, Pietrzak, Renner 2007] :

$$\boxed{\mathbf{Adv}_{F \circ G^{-1}}^{\mathrm{cca}}(q) \leq \mathbf{Adv}_{F}^{\mathrm{cpa}}(q) + \mathbf{Adv}_{G}^{\mathrm{cpa}}(q)}$$

Notation and basic setup

Fix distinct $z_1, \ldots, z_q \in \mathcal{C} = \{0,1\}^n$ and define:

X_t Positions of cards $z_1, \ldots, z_q$ at time t

$\{X_t\}$ Markov chain — the projected Thorp shuffle

$X_t(i)$ Location of card z_i at time t

τ_t Distribution of $\{X_t\}$

π Stationary distribution of $\{X_t\}$
= Uniform distribution on q-tuples of positions, $\{0,1\}^n$

Want to show : $\|\tau_t - \pi\|$ is small (for t not too big)

Hybrid argument

For $0 \le \ell \le q$, let

X_t^ℓ = Positions of cards $z_1, \ldots, z_q$ at time t assuming cards
$z_1, \ldots, z_\ell$ start in **designated** positions,
$z_{\ell+1}, \ldots, z_q$ start in **random** (uniform, distinct) positions

$X_t^q \quad \cdots \quad X_t^{\ell+1} \quad X_t^\ell \quad \cdots \quad X_t^0$

Designated cards have specified posns.
τ_t - distributed

Fix ℓ

Designated cards have random initial posns.
π-distributed

Then
$$\boxed{\| \tau_t - \pi \| \le \sum_{\ell=0}^{q-1} \| \tau_t^{\ell+1} - \tau_t^\ell \|}$$

[Doeblin 1930s; Aldous 1980s]

Coupling arguments

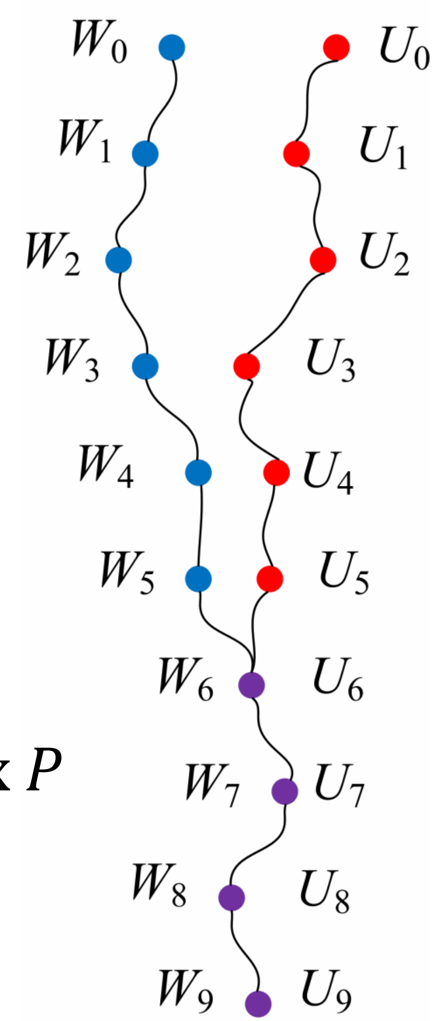

Markov chain $\{ W_t \}$ with transition matrix P
Stationary distribution π

Want to show $\| P^t(x, \cdot) - \pi \|$ is small

Construct a **pair process** , $\{(W_t, U_t)\}$ (defined on a single prob space), the **coupling**, where

- $\{ W_t \}$ and $\{ U_t \}$ are MCs with transition matrix P
- If $W_t = U_t$ then $W_{t+1} = U_{t+1}$
- $W_0 = x$ and $U_0 \sim \pi$

Let $T = \min \{t : W_t = U_t \}$

Coupling time

Then $\| P^t(x, \cdot) - \pi \| \leq \Pr (W_t \neq U_t)$
$= \Pr (T > t)$

What gets coupled

Fix ℓ

First ℓ+1 cards in designated positions.

$\tau_t^{\ell+1}$ distributed

First ℓ cards in designated positions; (ℓ+1)st card at a random position.

τ_t^{ℓ} distributed

$$X_t^q \quad \cdots \quad X_t^{\ell+1} \quad X_t^{\ell} \quad \cdots \quad X_t^0$$

Then $$\| \tau_t - \pi \| \leq \sum_{\ell=0}^{q-1} \| \tau_t^{\ell+1} - \tau_t^{\ell} \|$$

Towards defining our coupling
Re-conceptualizing how our MC evolves

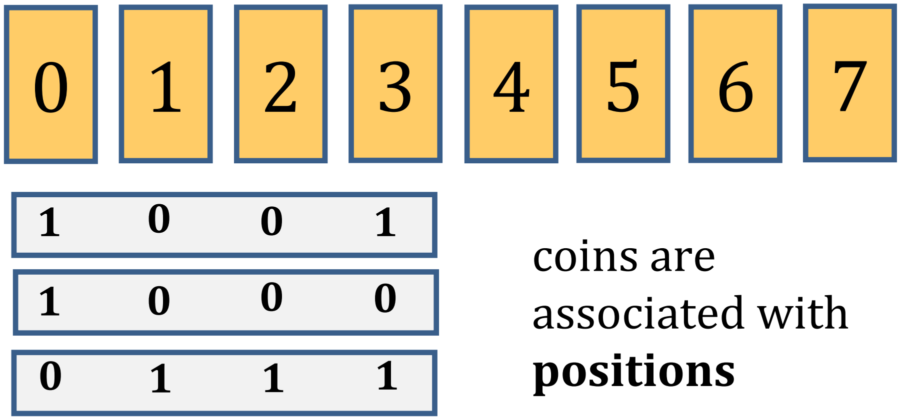

coins are associated with **positions**

Before: a coin $c(r\,,x)$ for each round r and **position** $(x, x + N/2)$. The coin determined if cards went

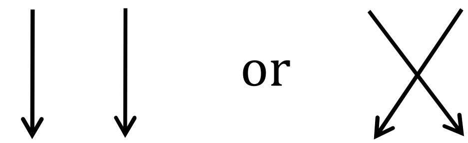

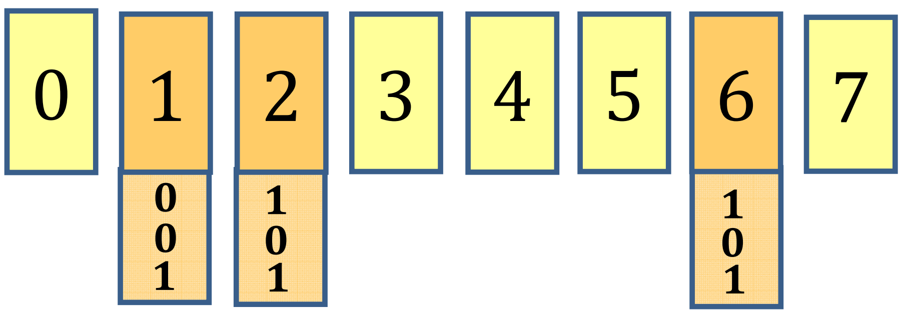

coins are associated with **designated cards**

Now: a coin $c(r, x)$ for each round r and **designated card** x.

Update rule:

- Card z_i adjacent to a non-designated card: use its coin to decide if it goes left (0) or right (1)
- Card z_i adjacent to z_j where $i < j$: use the coin of z_i to decide where it goes … and so where z_j goes, too.

Defining our coupling

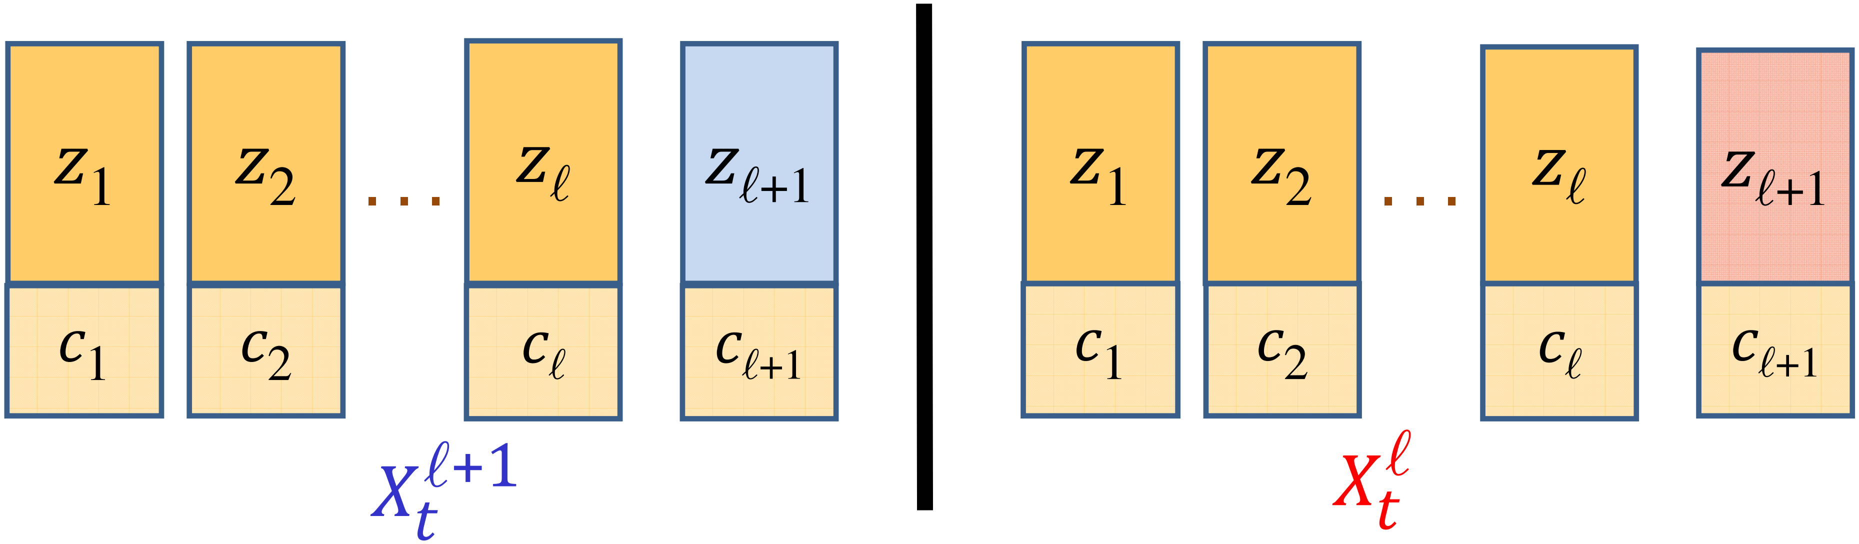

To define the pair process $(X_t^{\ell+1}, X_t^{\ell})$

- Start cards $z_1, \ldots, z_\ell$ in the specified locations for both $X_t^{\ell+1}$ and X_t^{ℓ}
- Start card $z_{\ell+1}$ at specified location in $X_t^{\ell+1}$
- Start card $z_{\ell+1}$ at uniform location in X_t^{ℓ}
- Evolve the process with the same coins and the update rule

Then:

- Cards $z_1, \ldots, z_\ell$ follow the **same** trajectory
- Once $z_{\ell+1}$ and $z_{\ell+1}$ match, they stay the same
- Card $z_{\ell+1}$ is uniform

Waiting for the $(\ell+1)^{\text{st}}$ cards to couple

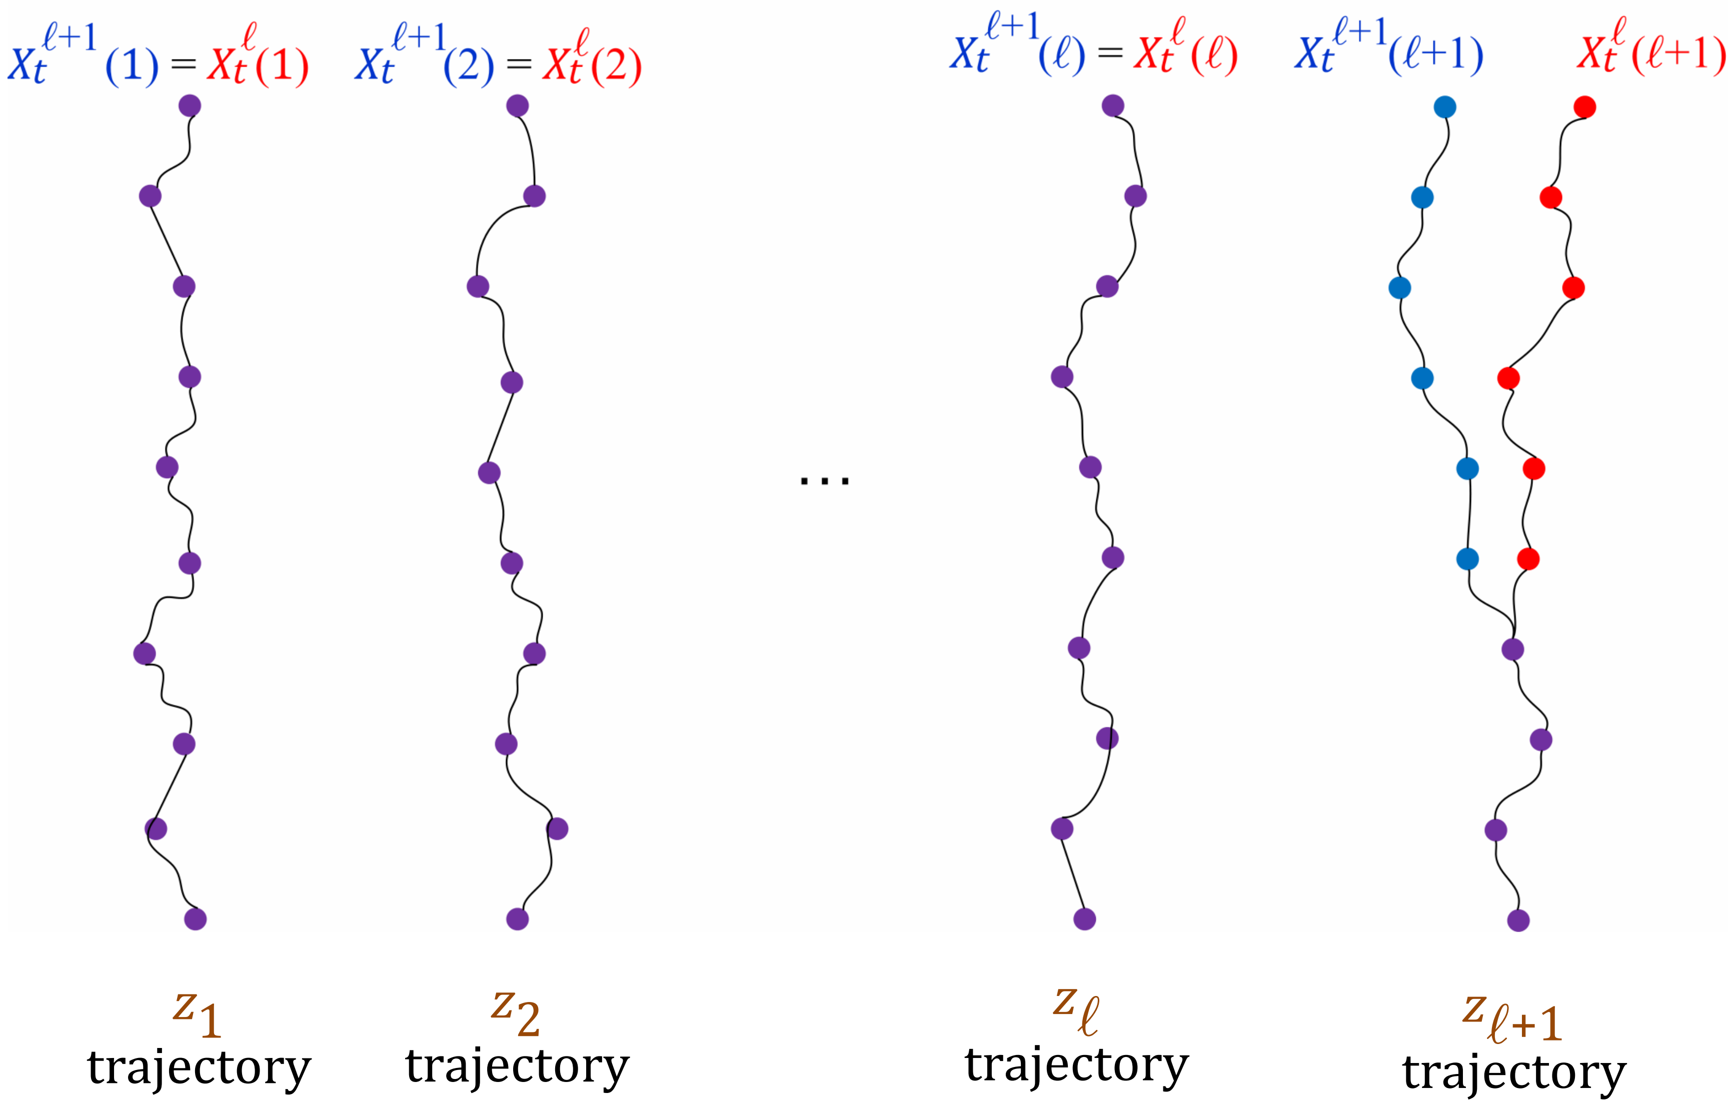

After a "burn-in" period, designated cards are rarely adjacent

Claim: For any pair of cards z_i and z_j and any time $t \geq n-1$,
$\mathbf{P}(z_i \text{ and } z_j \text{ are adjacent at time } t) \leq 1/2^{n-1}$

Reason: The only way for z_i and z_j to end up adjacent at time t is if there were **consistent coin tosses** in in each of the prior $n-1$ steps.
The probability of this is $1/2^{n-1}$.

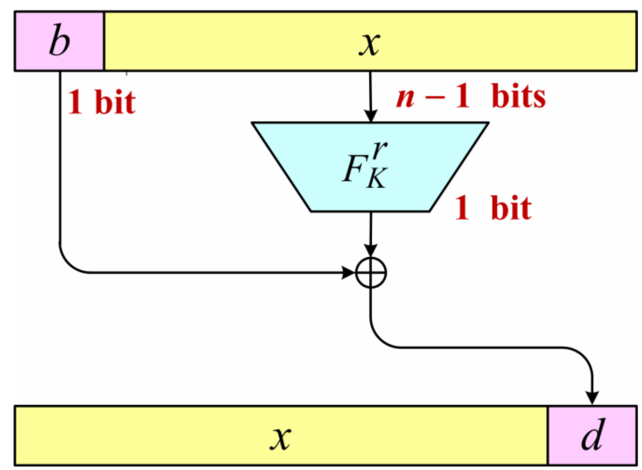

The coupling bound

$$\| \tau_t - \pi \| \le \sum \overbrace{\| \tau_t^{\ell+1} - \tau_t^{\ell} \|}^{\text{Want to show this is small.}}$$

By coupling, it's $\le \mathbf{P}(T > t)$ where T is the coupling time for $X_t^{\ell+1}$ and X_t^{ℓ} :

$$T = \min \{ t\colon \mathbf{P}\big(X_t^{\ell+1} = X_t^{\ell}\big)\}$$

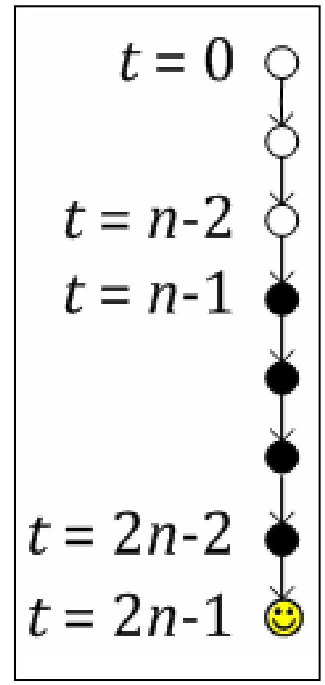

Claim: $\boxed{\mathbf{P}\,(T > 2n - 1) \le 2 \times n \times \ell \times (1 / 2^{n\text{-}1})}$

Cards $z_{\ell+1}$ **fail** to converge only if

$z_{\ell+1}$ is adjacent to some z_i in $X_t^{\ell+1}$ **or**

$z_{\ell+1}$ is adjacent to some z_i in X_t^{ℓ}

for some $i \le \ell$, in one of the last n time steps.

At most $2n\ell$ ways for this to happen. Just showed:

$$\mathbf{P}\big(z_{\ell+1} \text{ and } z_i \text{ are adjacent at time } t \le n+1\big) \le 1/\,2^{n-1}$$

Concluding the result

$$\mathbf{P}\,(T > 2n\text{-}1\,) \;\le\; 2 \times n \times \ell \times 2^{1-n}$$

so $$\mathbf{P}\,(T > r\,(2n\text{-}1)\,) \;\le\; (\,2 \times n \times \ell \times 2^{1-n}\,)^{r}$$

$$\|\,\tau_t - \pi\,\| \;\le\; \sum_{\ell=0}^{q-1} \left(n\ell 2^{2\text{-}n}\right)^r \;\le\; \left(n2^{2\text{-}n}\right)^r \int_0^q x^{\,r}\, dx$$

$$\Vert$$

$$\mathbf{Adv}^{\mathrm{ncpa}}_{N,\,R}\,(q) \qquad\qquad \le\; \frac{q}{r+1}\left(\frac{4qn}{N}\right)^{r}$$

Extensions and directions

- For a weaker security notion, DPA, **two passes** is enough.
- A simple trick lets you do **5 rounds per AES**
- When *N* is **not a power of 2**, things get more complex (in progress; constants increase)
- **NIST submission** ("FFX mode") (with T. Spies) coming soon
- **Coupling technique** generally useful in cryptography. Analyze other unbalanced Feistel schemes with V.T. Hoang.

- **Open:**
  Tiny *N* ?
  CCA security for 2 or 4 passes ?
  Can perfect shuffling (à la [Granboulan, Pornin 07]) be practical?

Thorp shuffle — DPA security

Theorem Let $N = 2^n$ and R=2nr (ie, 2r passes).

$$\mathbf{Adv}^{\text{dpa}}_{N,R}(q) \le \left(\frac{4qn}{N}\right)^r$$

Asymptotically: you can tolerate $q = N^{1-\varepsilon}$ queries with two rounds

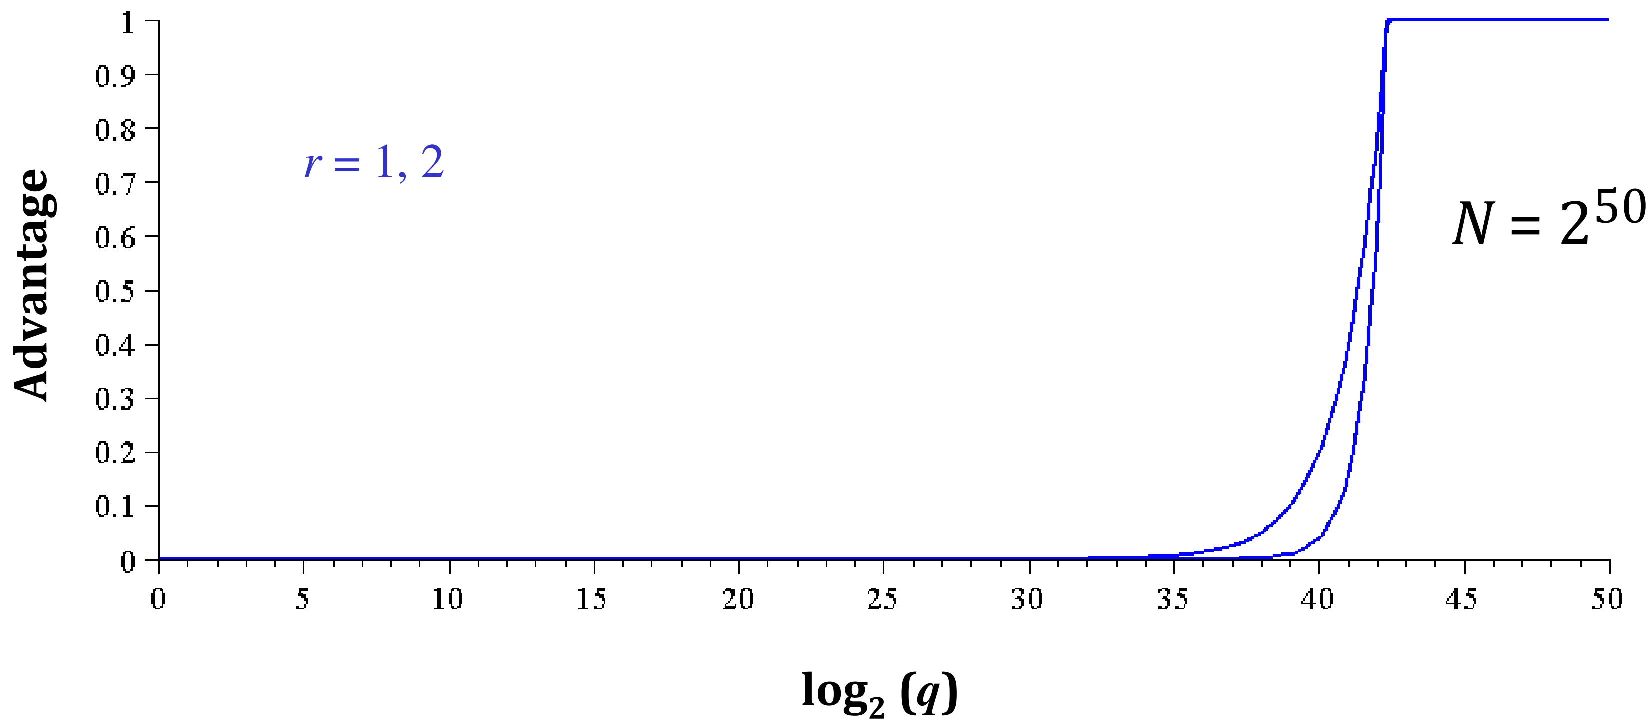

The 5x speedup trick