Key Recovery Attacks on 3-round Even-Mansour, 8-step LED-128, and Full AES^2

Itai Dinur[1], Orr Dunkelman[2,4],
Nathan Keller[3] and Adi Shamir[4]

[1]École normale supérieure, France

[2]University of Haifa, Israel

[3]Bar-Ilan University, Israel

[4]The Weizmann Institute, Israel

Summary

- The **Even-Mansour** scheme is simple construction of a block cipher proposed in 1991
- The scheme has been generalized to **iterated Even-Mansour** schemes
  - Extensively studied in the last few years

- We study the security of **iterated Even-Mansour** schemes
  - Attack schemes that were previous assumed to be secure
  - Present applications to **concrete** designs

The Even-Mansour Scheme (1991)

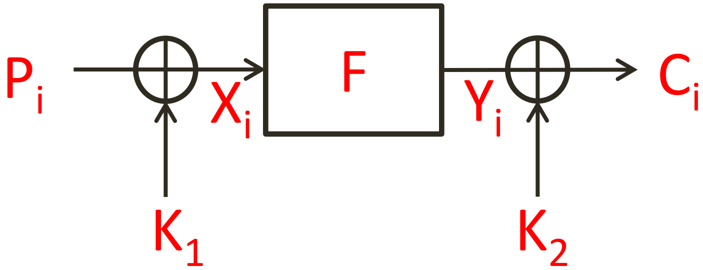

- A simple construction of a block cipher using 2 keys of n bits and a **public** permutation F
- **Information-theoretic** security lower bound:
  - Assume that F is **randomly chosen**
  - Assume that we obtain D plaintext-ciphertext pairs (P_i,C_i)
  - Then, any successful key-recovery attack that evaluates F on T inputs X must satisfy $TD \geq 2^n$

The SlideX Attack [DKS ‘12]

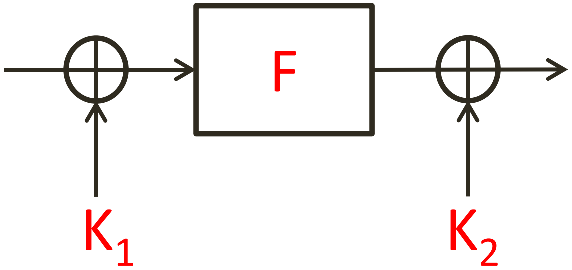

- Security: $TD=2^n$ using the **SlideX** attack

(DKS, Eurocrypt ‘12)

- Given $D=2^{n/2}$ the scheme can be broken in $T=2^{n/2}$

SlideX on EM with 1 Key [DKS ‘12]

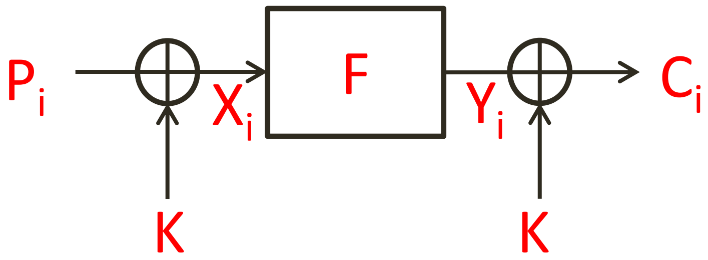

- P_i+K=X_i and C_i+K=Y_i → $P_i+C_i = X_i+Y_i$
- For each (P_i,C_i):
  - Calculate P_i+C_i and store it in a sorted table next to P_i
- For arbitrary values X_j:
  - Calculate $Y_j=F(X_j)$ and search X_j+Y_j in the table
  - For each match, test the suggestion for $K=P_i+X_j$

P_i+C_i	P_i
⋮	⋮

SlideX on EM with 1 Key: Analysis

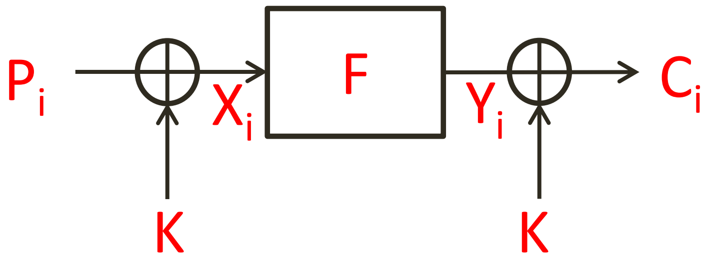

- In order to obtain **w.h.p** a pair (P_i,X_j) such that $K=P_i+X_j$ we need about 2^n such pairs, i.e. $TD=2^n$

The Iterated EM Scheme

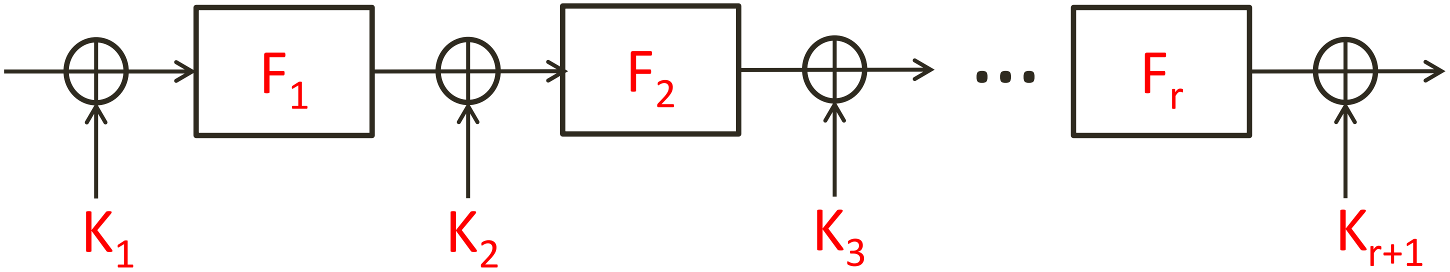

- EM-based schemes are a **very hot** research area
  - Over 10 papers in major crypto conferences since 2011

- There are many possible **key schedules**

2-Round Iterated EM with 1 Key

- Does not provide n-bit security as shown at FSE 2013 [NWW '13]

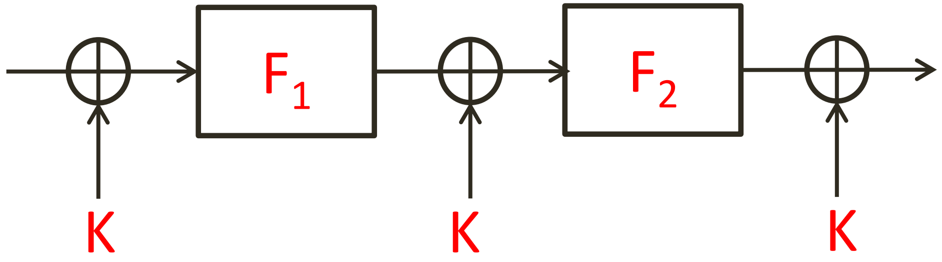

A Variant of the Previous Attack [NWW '13] : Main Idea

- $P_i+V_i=X_i+Y_i$ → $X_1+Y_1= X_2+Y_2=\ldots=X_t+Y_t=\Delta$ then $P_1+V_1= P_2+V_2=\ldots=P_t+V_t=\Delta$
- A t-way collision on the **public** $F'_1(X)=X+F_1(X)$ gives a t-way collision on P_i+V_i with the **same** value Δ
- Given Δ and a random P_i, then $V_i=P_i+\Delta$ with probability $t/2^n>1/2^n$

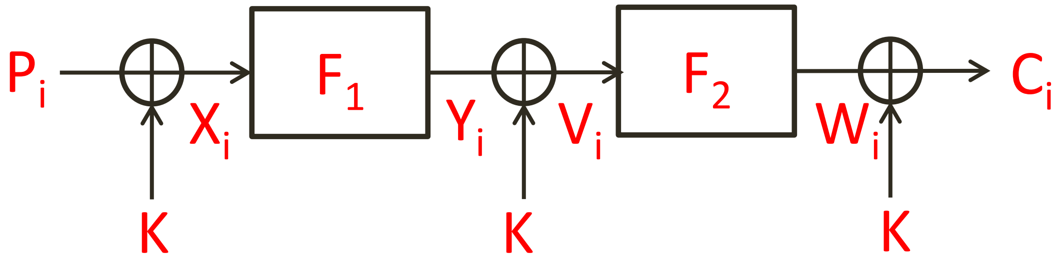

A Variant of the Previous Attack [NWW '13]

- **Preprocessing**: Evaluate F_1 on arbitrary inputs X, find a t-way collision on $F'_1(X)=X+F_1(X)$ and denote the colliding value by Δ
- **Online**: For each (P_i, C_i):
  - Assume that $V_i=P_i+\Delta$ and compute $W_i=F_2(V_i)$
  - Compute a suggestion for $K=W_i+C_i$ and test it

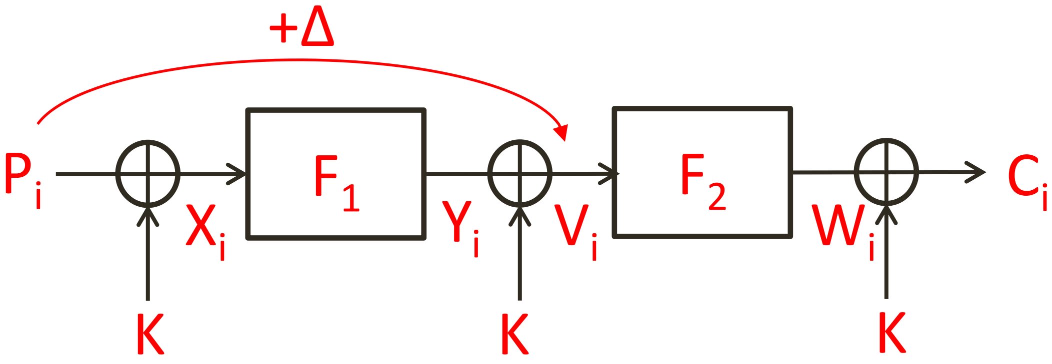

A Variant of the Previous Attack [NWW ‘13] : Analysis

- The data complexity is $D=2^n/t$
  - in order to find a P_i such that $V_i=P_i+\Delta$ and recover K
- The **online** time complexity is also $2^n/t$
- What is the complexity of the preprocessing?

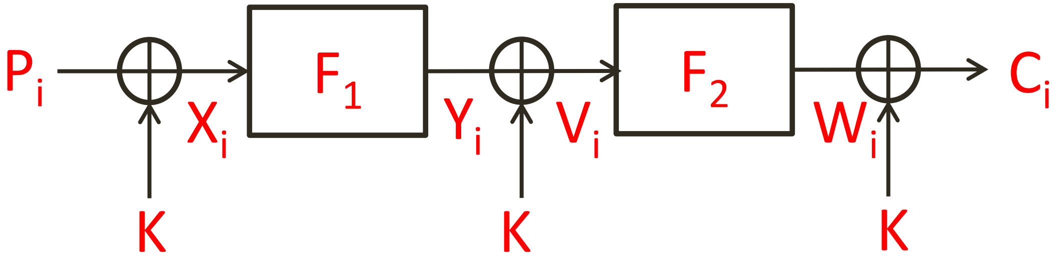

A Variant of the Previous Attack [NWW '13] : Analysis

- If we evaluate F'_1 on **all** 2^n inputs, the attack will not be faster than **exhaustive search**
- We evaluate F'_1 on a $\lambda<1$ fraction of the inputs
- The **preprocessing** time complexity is $\lambda 2^n$
  - in which we find a t-way collision

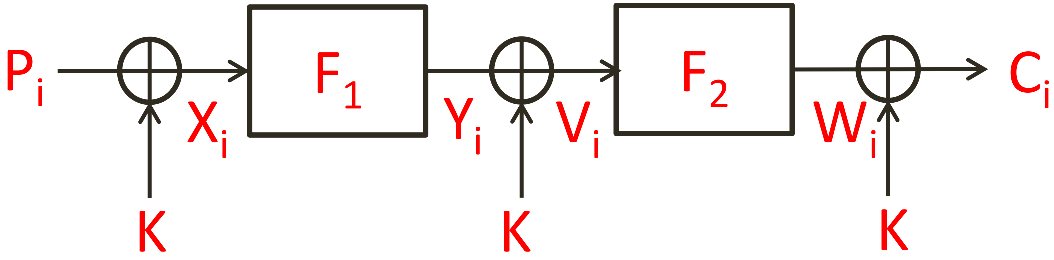

A Variant of the Previous Attack [NWW ‘13] : Analysis

- The **total** time complexity is $\lambda 2^n+2^n/t$
- To calculate the **optimal** time complexity, we need to understand the **tradeoff** between λ and t
- What is the largest t-way collision we expect when evaluating a λ fraction of inputs for F'_1?

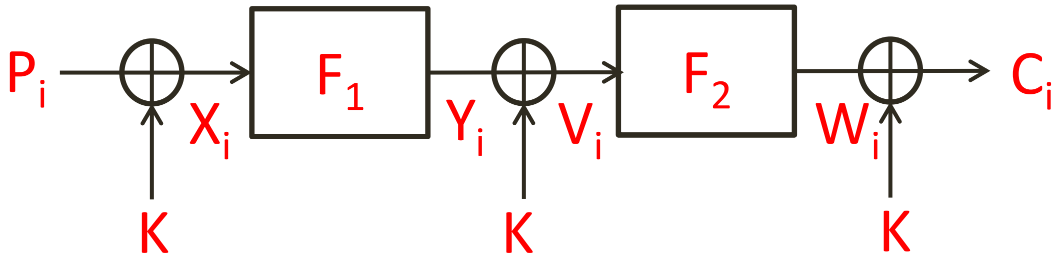

A Variant of the Previous Attack [NWW '13] : Analysis

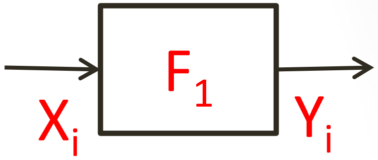

- $F'_1(X)=X+F_1(X)$ is a function from n bits to n bits
- If we evaluate $F'_1(X)$ on a λ fraction of the inputs the expected number of t-way collisions is $(2^n\lambda^t e^{-\lambda})/t!$
  - Assuming standard randomness assumptions on F_1

A Variant of the Previous Attack [NWW '13] : Analysis

- The **tradeoff** between λ and t is enforced by $(2^n\lambda^t e^{-\lambda})/t! \geq 1$
- Taking $\lambda \approx 1/n$ gives $t \approx 1/\lambda \approx n$ and **minimizes** $T \approx 2^n/n$
  - This is faster than **exhaustive search** by a factor of about n, which grows to **infinity** with n
- For $n=64$ → $T \approx 2^{64}/64 \approx 2^{60}$ and also $D \approx 2^{60}$, $M \approx 2^{60}$

Our First Optimization: Reducing the Data Complexity - Main Idea

- Once we take λ and t for which $(2^n\lambda^t e^{-\lambda})/t! \geq 1$, and **slightly** reduce t, the number of t-way collisions grows **rapidly**

Our First Optimization: Reducing the Data Complexity - Analysis

- For $n=64$ and 2^{60} inputs we expect:
  - **4** 10-way collisions
  - **95** 9-way collisions
  - Over **100,000** 8-way collisions
- We can exploit all these in the attack
- For $n=64$ we **greatly reduce** the data complexity from 2^{60} to 2^{45}
  - by taking all collisions with $t \geq 8$ rather than $t \geq 10$
  - The time and memory complexities slightly increase but remain about 2^{60}

3-Round Iterated EM with 1 Key

- The attack on 2-round EM was already somewhat marginal
- We show that 3-round EM **does not** provide n-bit security as well!

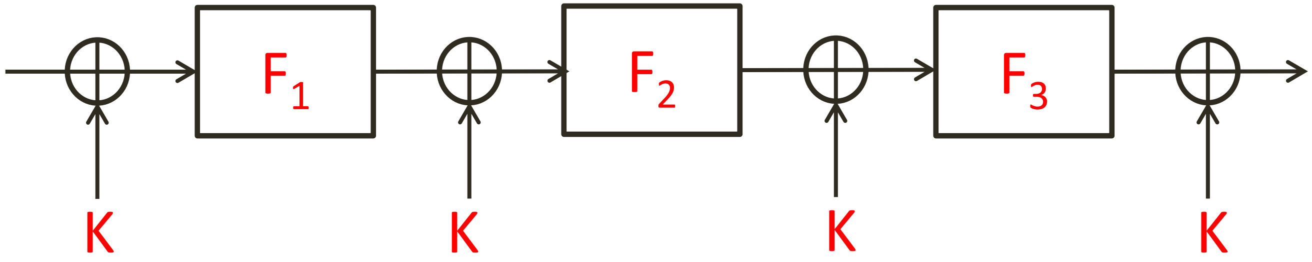

The Main Idea of our New Attack

- We know how to **predict** W_i with a higher probability than a random guess
- Given W_i and C_i we remain with a 1-round EM with 1 key and can apply the **SlideX** attack

- The time complexity increases to $T \approx 2^n/\sqrt{n}$
  - Faster than **exhaustive search** only by a factor of $\sqrt{n}$

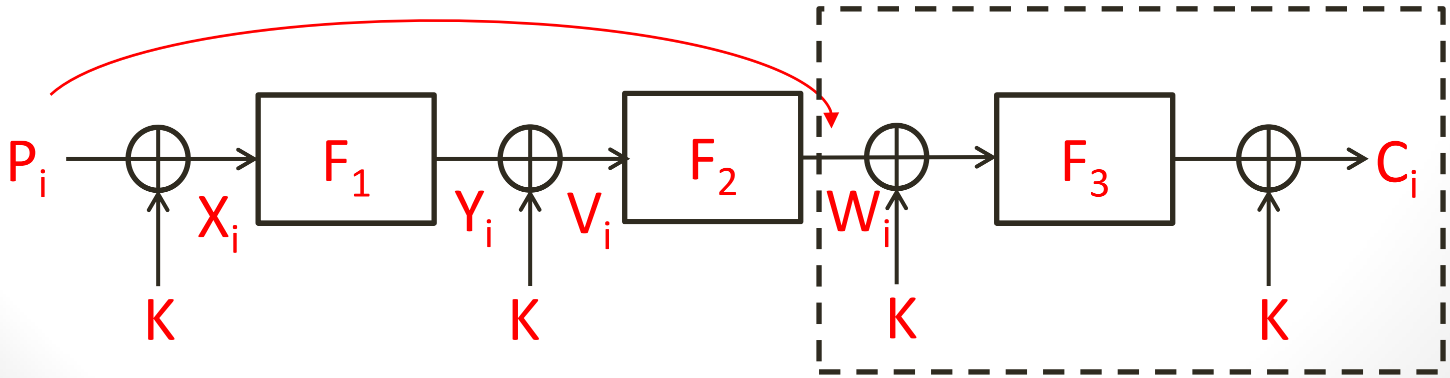

Optimizing our 3-Round Attack

- Apply the same optimization as in the 2-round attack to reduce the **data complexity**
- Use the **freedom** to choose the inputs on which we evaluate F_1 and F_3 in order to **immediately filter** most uninteresting (P_i,C_i)
- The optimization gives us $T \approx 2^n/n$
- This is about the **same** time complexity as the 2-round attack!

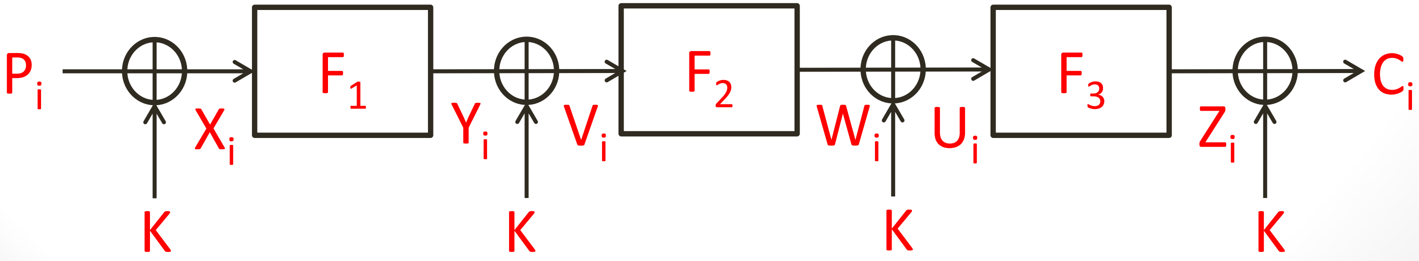

Application to (Original) Zorro

- Zorro is a 128-bit lightweight block cipher presented at CHES 2013 by Gérard et al.
- The **original** cipher was a 3-round EM scheme with 1 key
- The authors **changed** the design due to our results

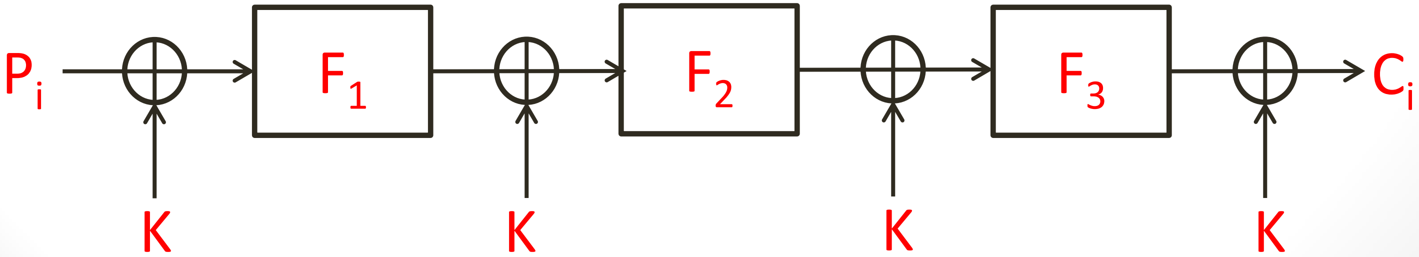

Application to LED-64

- LED is a 64-bit lightweight block cipher presented at CHES 2011 by Guo et al.
- Two main versions: LED-64 and LED-128
- LED-64 is an 8-round EM scheme with 1 key
- Previous attacks on LED-64 **could** only attack 2 rounds

- We can directly apply our attack to 3-round LED-64 with $T \approx 2^{60}$, $M \approx 2^{60}$ and $D = 2^{49}$

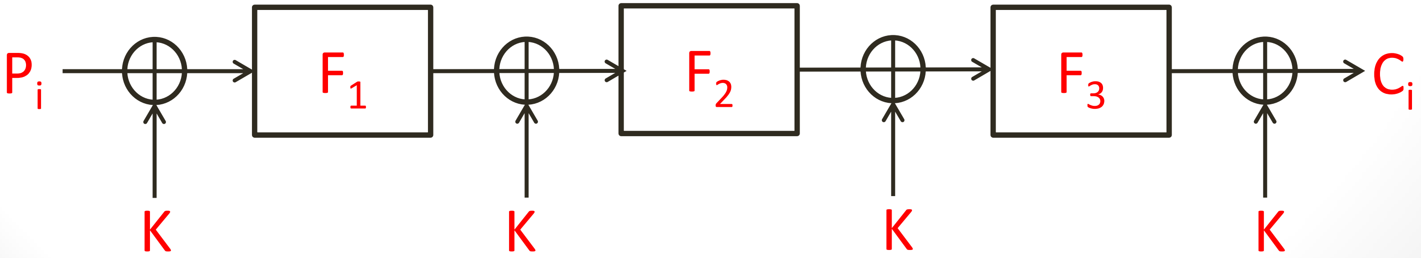

Application to LED-128

- LED-128 uses 2 alternating keys and has 12 rounds
- The best previous attack [NWW '13] could attack 6 rounds

- We use the new techniques to attack 8 rounds!

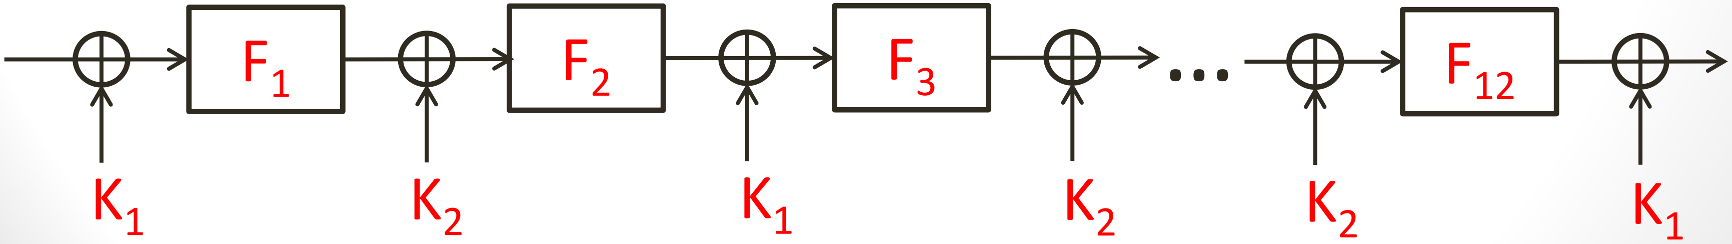

Application to LED-128

- As several previous attacks we guess K_1 in an outer loop
- We remain with a 3-round EM scheme with 1 key
- We obtain $T \approx 2^{124}$, $M \approx 2^{60}$ and $D = 2^{49}$
- About the **same** time and memory complexities as the previous 6-round attack, and the data is **reduced** by a factor of about 1000!

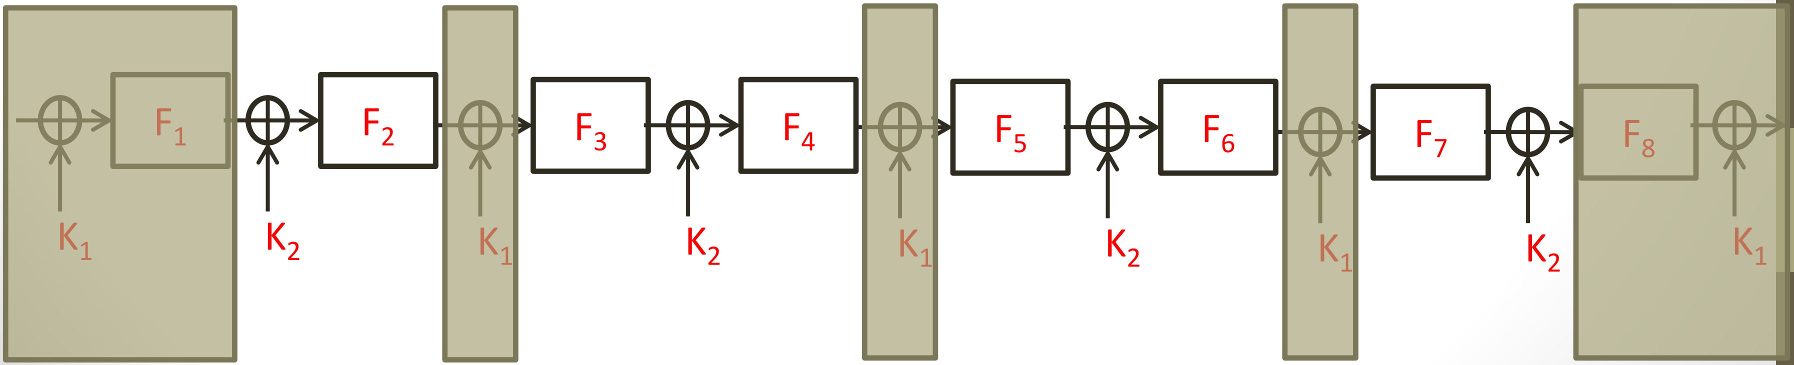

2-Round EM with Independent Keys

- A simple meet-in-the-middle attack has time and memory complexity of 2^n

- t-way collisions on X_i+Y_i do not seem to help

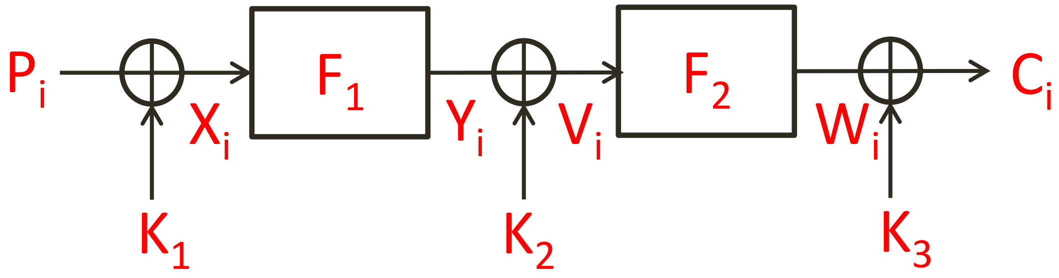

Our Attack on 2-Round EM with Independent Keys: The Main Idea

- Use the **differential** algorithm of Mendel et al. from ASIACRYPT 2012
- However, we apply attack even when F_1 and F_2 do not have any **statistical weakness**!
- The attack uses **additional** techniques…

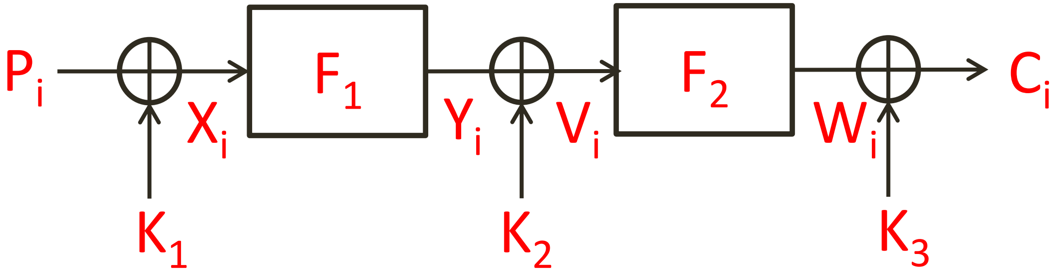

Application to AES^2

- AES^2 is 128-bit block cipher presented at EUROCRYPT 2012 by Bogdanov et al.

- A 2-round EM with independent 128-bit keys

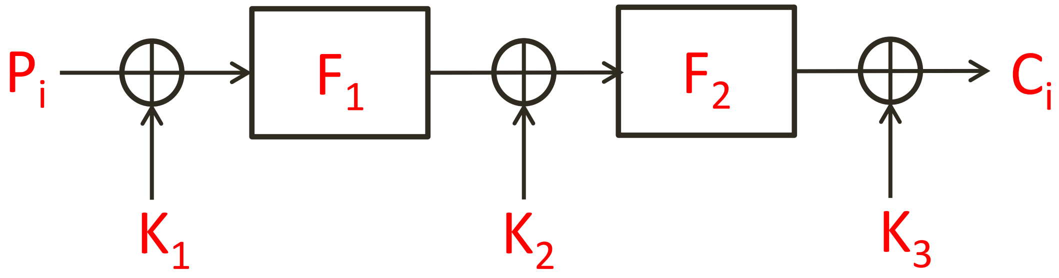

Application to AES^2

- Each public permutations is a **complete** AES-128 fixed-key encryption and is thus **very strong**
- The designers conjecture that the most efficient attack on AES^2 is a **basic meet-in-the-middle**

- Our attack is about 7 times **faster**
  - uses 7 times less memory (but requires much more data)

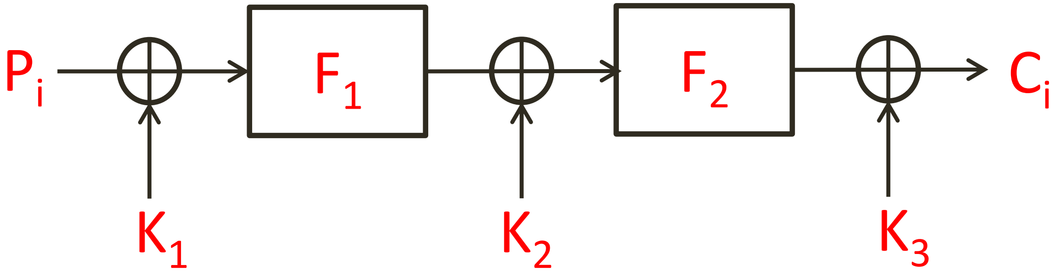

Conclusions

- We presented **improved** attacks on several schemes based on iterated Even-Mansour
- We described the **first** attack on full AES^2
- We **increased** the number of steps that can be attacked for LED-128 from 6 to 8
- The attacks are **unlikely** to be practically significant
- They show that a 1-key EM scheme needs to have **at least** 4 rounds to provide n-bit security

Thank you for your attention!