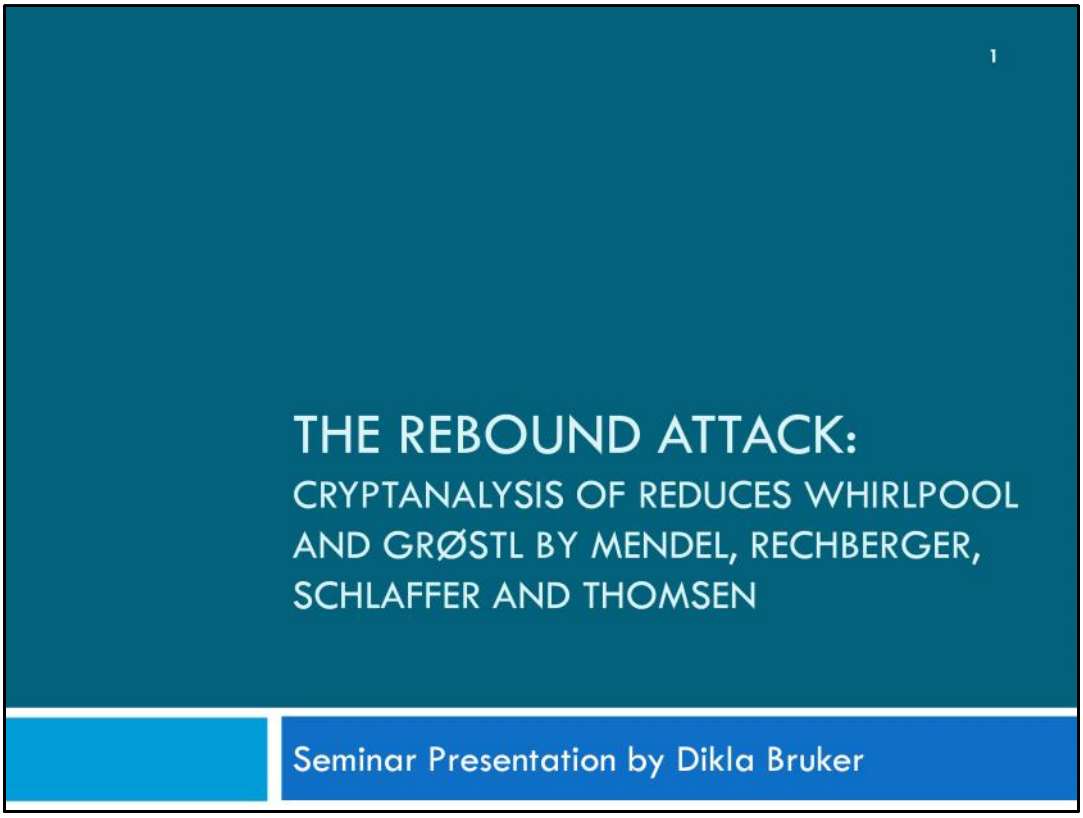

- Whirlpool מערבולת

What's On The Menu?

2

- Introduction
- Advanced Encryption Standard
- Whirlpool Block Cipher
- Rebound Attack on Whirlpool
- Grøstl Block Cipher
- Summery

What's On The Menu?

- Introduction
- Advanced Encryption Standard
- Whirlpool Block Cipher
- Rebound Attack on Whirlpool
- Grøstl Block Cipher
- Summery

Introduction

- The rebound attack is a technique for hash functions cryptanalysis.
- Divide the block cipher *E* into 3: $E = E_{fw} \circ E_{in} \circ E_{bw}$
- Inbound phase (The meet in the middle phase) and Outbound phase
- Our Goal: Find a differential trail that will cause the input and output differentials to cancel each other

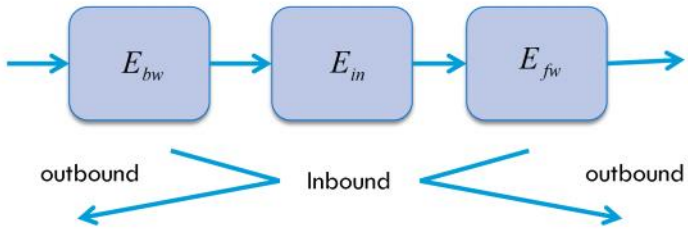

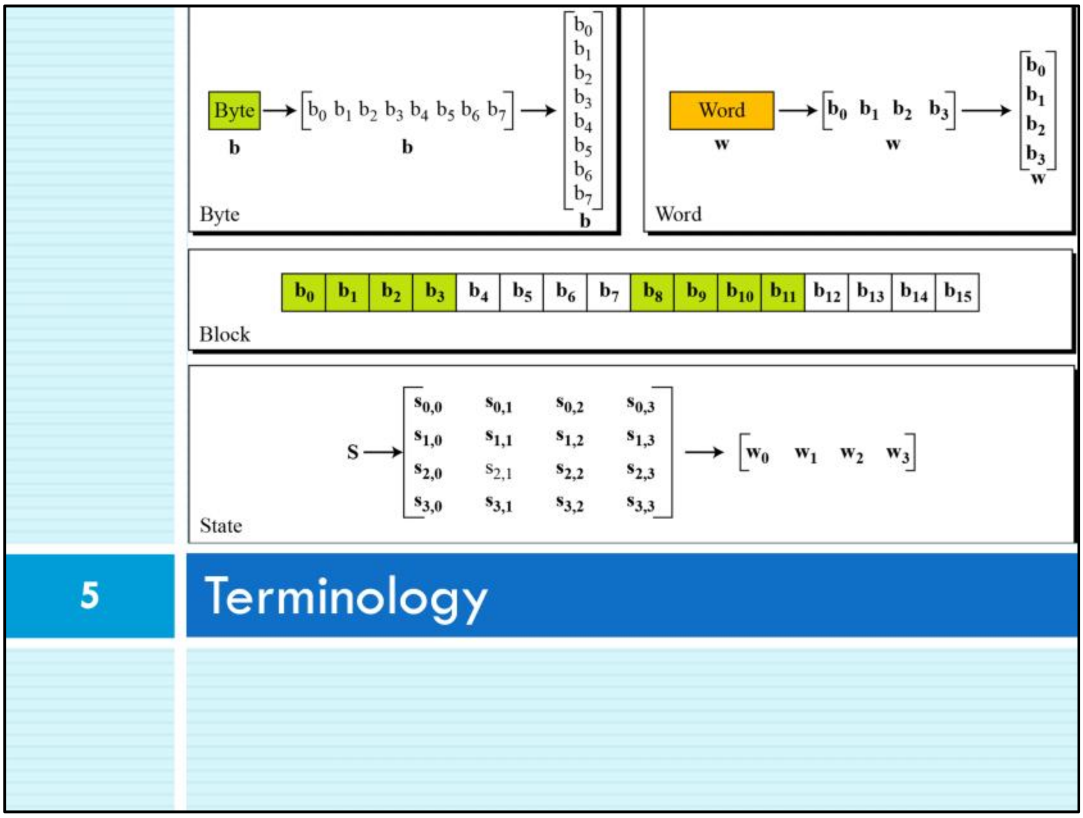

1 byte = 8 bit
Each byte is represented by 2 hexadecimal digits
1 word = 4 byte = 32 bit
1 block = 4 words = 128 bit

Advanced Encryption Standard (AES)

6

- A private key symmetric block.

Key Size (bits)	128	192	256
Plain Text Block Size (bits)	128	128	128
Number of Rounds N_r	10	12	14
Round Key Size (bits) $K_0, K_1, \ldots, K_{Nr}$	128	128	128

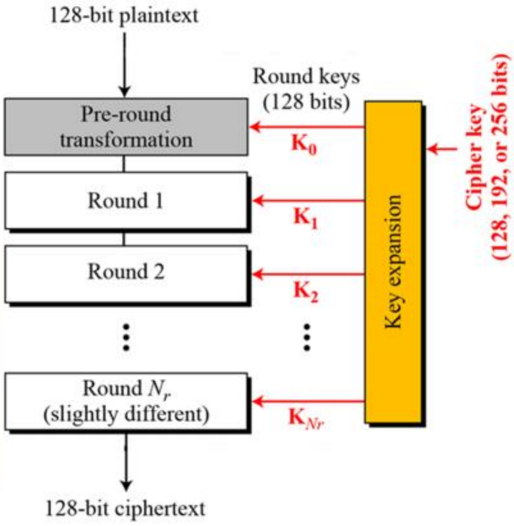

What's On The Menu?

7

- Introduction
- Advanced Encryption Standard
- Whirlpool Block Cipher
- Rebound Attack on Whirlpool
- Grøstl Block Cipher
- Summery

Advanced Encryption Standard (AES)

Structure Of Each Round

8

- Each round uses 4 transformations:
  - SubBytes
  - ShiftRows
  - MixColumns
  - Key Adding
- mixing transformation is missing from the last round.

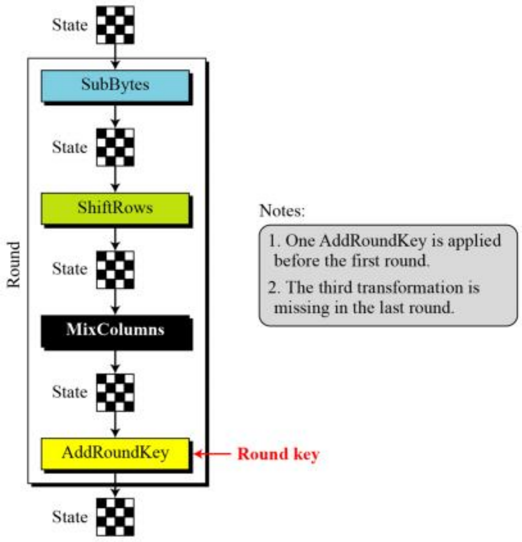

Advanced Encryption Standard (AES)

Structure Of Each Round (SubBytes)

- 16 independent byte-to-byte transformations.
- We interpret the byte as two hexadecimal digits
- Use a transformation table (as S-box) to transform 8 bit to 8 bit.

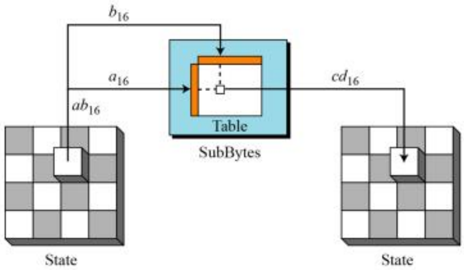

Table 7.1 *SubBytes transformation table*

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0	63	7C	77	7B	F2	6B	6F	C5	30	01	67	2B	FE	D7	AB	76
1	CA	82	C9	7D	FA	59	47	F0	AD	D4	A2	AF	9C	A4	72	C0
2	B7	FD	93	26	36	3F	F7	CC	34	A5	E5	F1	71	D8	31	15
3	04	C7	23	C3	18	96	05	9A	07	12	80	E2	EB	27	B2	75
4	09	83	2C	1A	1B	6E	5A	A0	52	3B	D6	B3	29	E3	2F	84
5	53	D1	00	ED	20	FC	B1	5B	6A	CB	BE	39	4A	4C	58	CF
6	D0	EF	AA	FB	43	4D	33	85	45	F9	02	7F	50	3C	9F	A8

We interpret the byte as two hexadecimal digits.

Advanced Encryption Standard (AES)
Structure Of Each Round (ShiftRows)

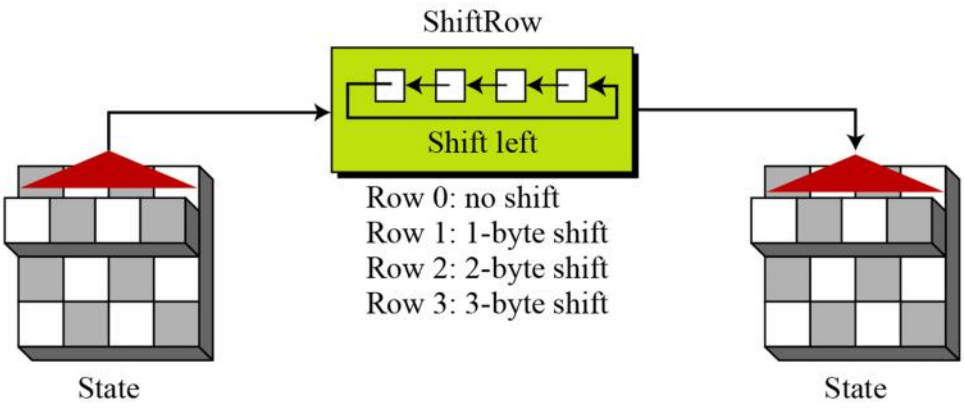

Advanced Encryption Standard (AES)

Structure Of Each Round (MixColumns)

- Change bits inside a byte, based on the bits inside the neighboring bytes
- Transforms each column of the state to a new column by multiplication with a constant matrix.

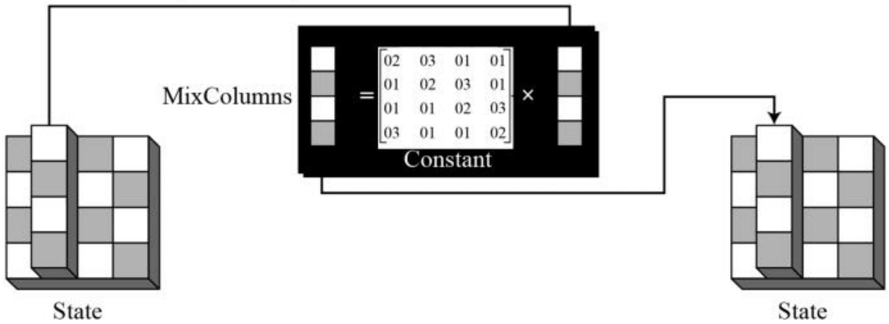

Advanced Encryption Standard (AES)

Structure Of Each Round (Add Round Key)

- The key is bitwise XORed to the state making the round function key dependent.

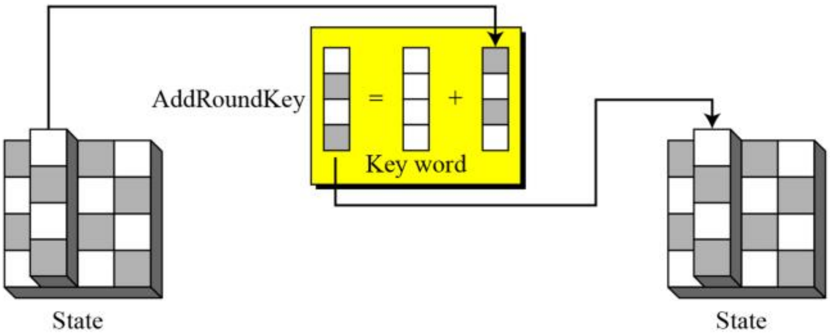

What's On The Menu?

- Introduction
- Advanced Encryption Standard
- Whirlpool Block Cipher
- Rebound Attack on Whirlpool
- Grøstl Block Cipher
- Summery

The Whirlpool Block Cipher

14

- first released in 2000 by Vincent Rijmen and Paulo S. L. M. Barreto. Since then a few revisions have taken place.
- Free, Whirlpool's designers have promised never to patent
- Named after the Whirlpool washing machine – Not!

The Whirlpool Block Cipher

15

- first released in 2000 by Vincent Rijmen and Paulo S. L. M. Barreto. Since then a few revisions have taken place.
- Free, Whirlpool's designers have promised never to patent
- Named after the Whirlpool galaxy, the first one recognized to have spiral structure by William Parsons, third Earl of Rosse

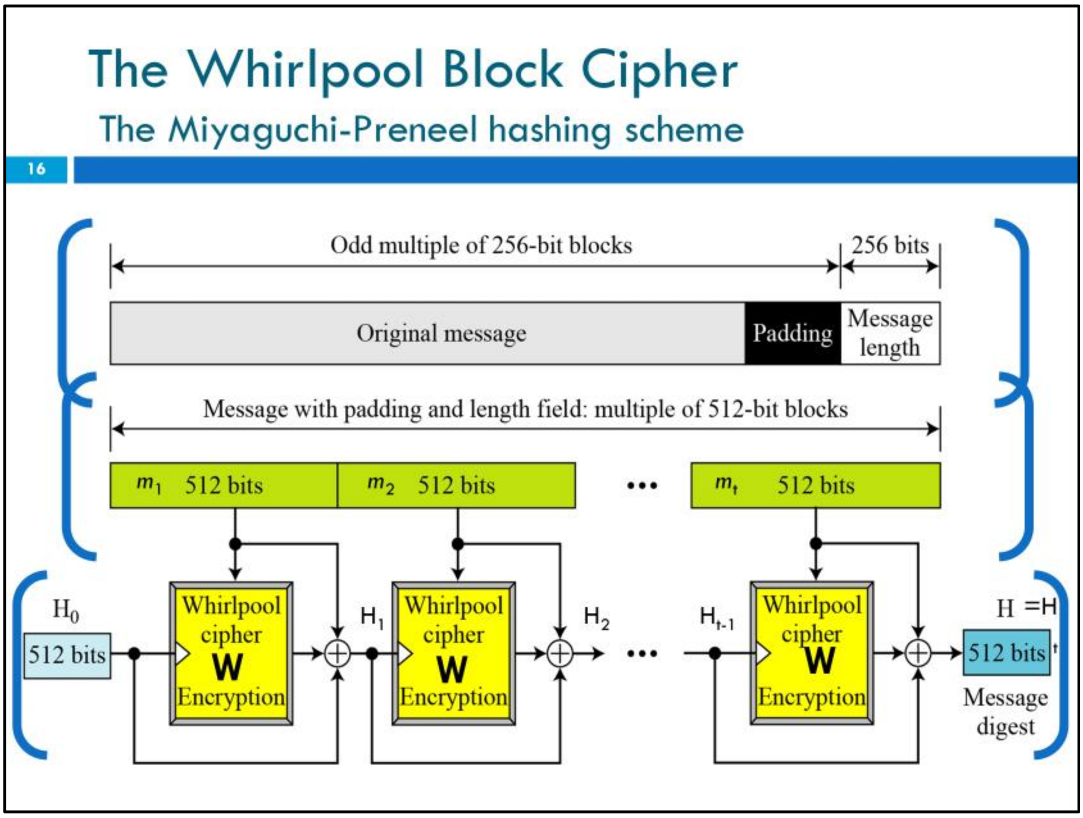

This hash function is based on the the Merkle-Damgård scheme we already seen in
class strengthening and the Miyaguchi-Preneel hashing scheme
The message is padded with a '1'-bit, then with a sequence of '0'-bits, and finally with the original length (in the form of a 256-bit integer value). The length after padding is a multiple of 512 bits.
The resulting message string is divided into a sequence of 512-bit blocks m_1, m_2, ... m_t which is then used to generate a sequence of intermediate hash values H_0, H_1, H_2, ... H_t. By definition, H_0 is a string of 512 '0'-bits.
To compute H_i, the block cipher
W encrypts m_i using H_{i-1} as key, and XORs the resulting ciphertext with both H_{i-1} and m_i. Finally, the **WHIRLPOOL** message digest is H_t.
The encryption algorithm is described in the next slide.

The Whirlpool Block Cipher

What goes on inside the blocks?

17

- 10 rounds.
- The encryption algorithm involves the use of 4 transformations:
  - Substitute Bytes (SB)
  - Shift Columns (SC)
  - Mix Rows (MR)
  - Add Key (AK)

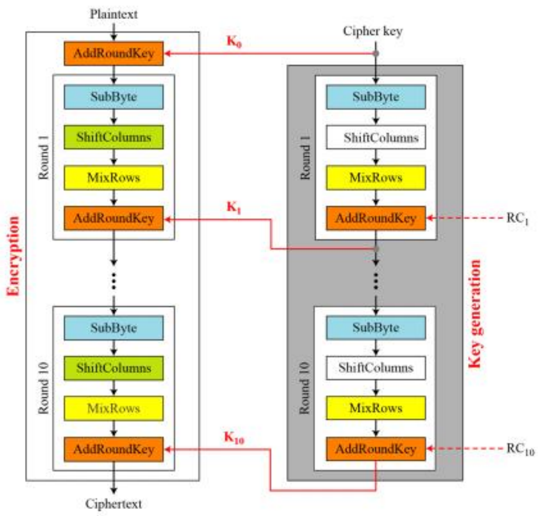

Whirlpool Cipher is very much like AES except minor differences:

1. The message length is 512 bit
2. Number of rounds is always 10
3. Key expansion is done in the round function and not in a dedicated algorithm
4. The S-Box in SybBytes is different
5. Instead of Shift Rows and Mix Columns we have Shift Columns and Mix Row

Table 1. Comparison of Whirlpool block cipher W and AES

	W	AES
Block size (bits)	512	128
Key size (bits)	512	128, 192, or 256
Matrix orientation	input is mapped row-wise	Input is mapped column-wise
Number of rounds	10	10, 12, or 14
Key expansion	W round function	dedicated expansion algorithm
$GF(2^8)$ polynomial	$x^8 + x^4 + x^3 + x^2 + 1$ (011D)	$x^8 + x^4 + x^3 + x + 1$ (011B)
Origin of S-box	recursive structure	multiplicative inverse in $GF(2^8)$ plus affine transformation
Origin of round constants	successive entries of the S-box	elements 2^i of $GF(2^8)$
Diffusion layer	right multiplication by 8×8 circulant MDS matrix (1, 1, 4, 1, 8, 5, 2, 9) - mix rows	left multiplication by 4×4 circulant MDS matrix (2, 3, 1, 1) - mix columns
Permutation	shift columns	shift rows

18

The Whirlpool Block Cipher

What goes on inside the blocks?

What's On The Menu?

The Whirlpool Block Cipher

Notations

- The transformation applied on a state in the i round of Whirlpool will be marked $r_i = AK \circ MR \circ SC \circ SB$
- The application of the S-Box on input *x* is marked *S(x)*
- The resulting stat of r_i will be marked as S_i.
  - The state after SubBytes – S'_i,
  - The state after Shift Columns– S''_i,
  - The state after Mix Rows– S'''_i,

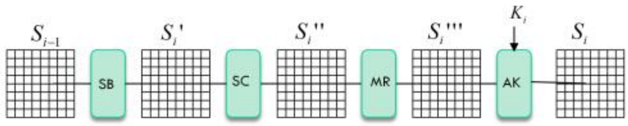

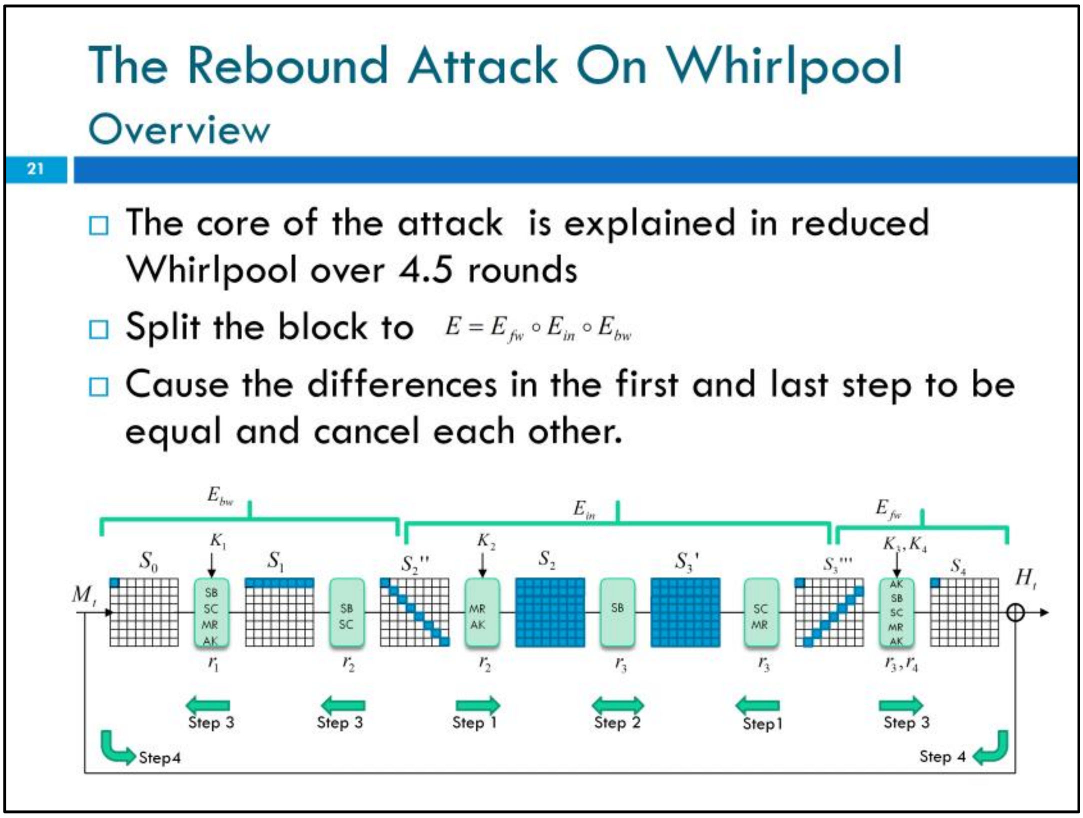

First, we will give an overview of the attack strategy which is the basis for the attacks on 4.5, 5.5 and 7.5 rounds.

The main idea of the attacks is to use a 4-round differential trail, which has the following sequence of active S-boxes: 1→8 →64 → 8 → 1

Using the Rebound Attack we can cover the most expensive middle part using an efficient match-in-the-middle approach (inbound phase).

If the differences in the first and last step are identical, they cancel each other through the feed-forward. The result is a collision of the round-reduced compression function of Whirlpool.

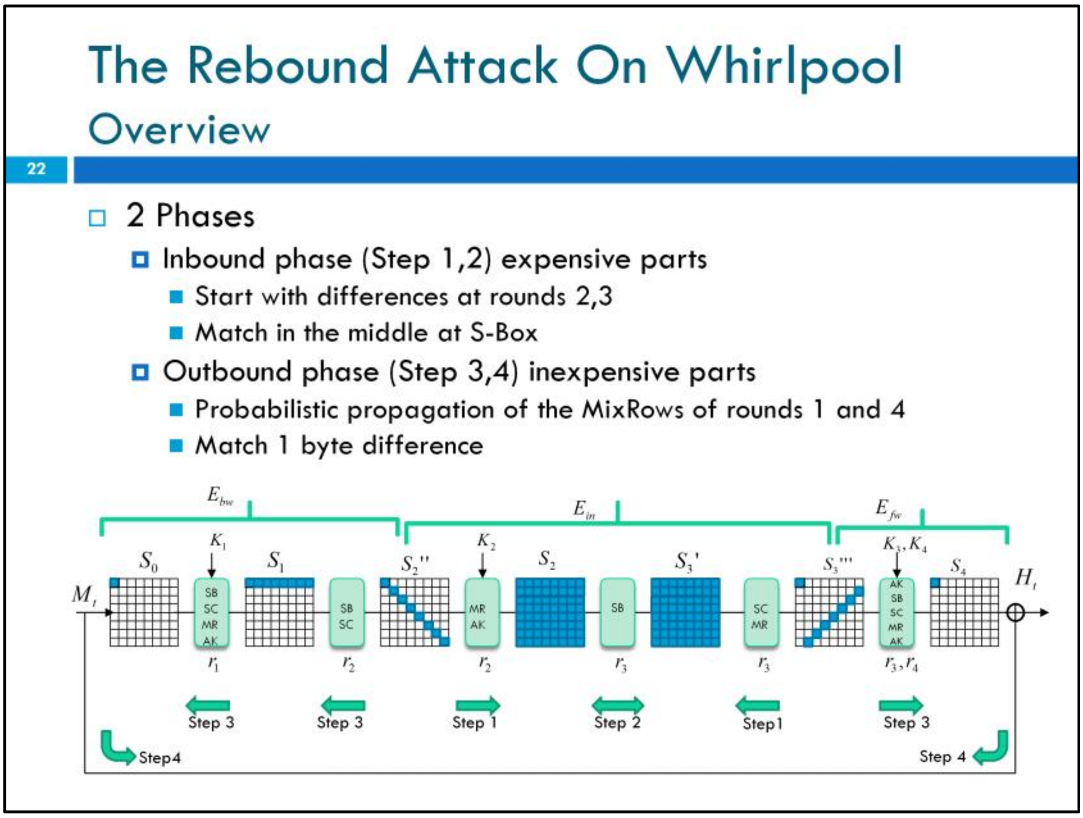

Inbound phase:

Step 1: start with 8-byte truncated differences at the MixRows layer of round r2 and r3, and propagate forward and backward to the S-box layer of round r3.

Step 2: connect the input and output of the S-boxes of round r3 to form the three middle states 8 →64 → 8 of the trail.

Outbound phase

Step 3: extend the trail both forward and backward to give the trail 1→8 →64 → 8 → 1 through MixRows in a probabilistic way.

Step 4: link the beginning and the end of the trail using the feed-forward of the hash function.

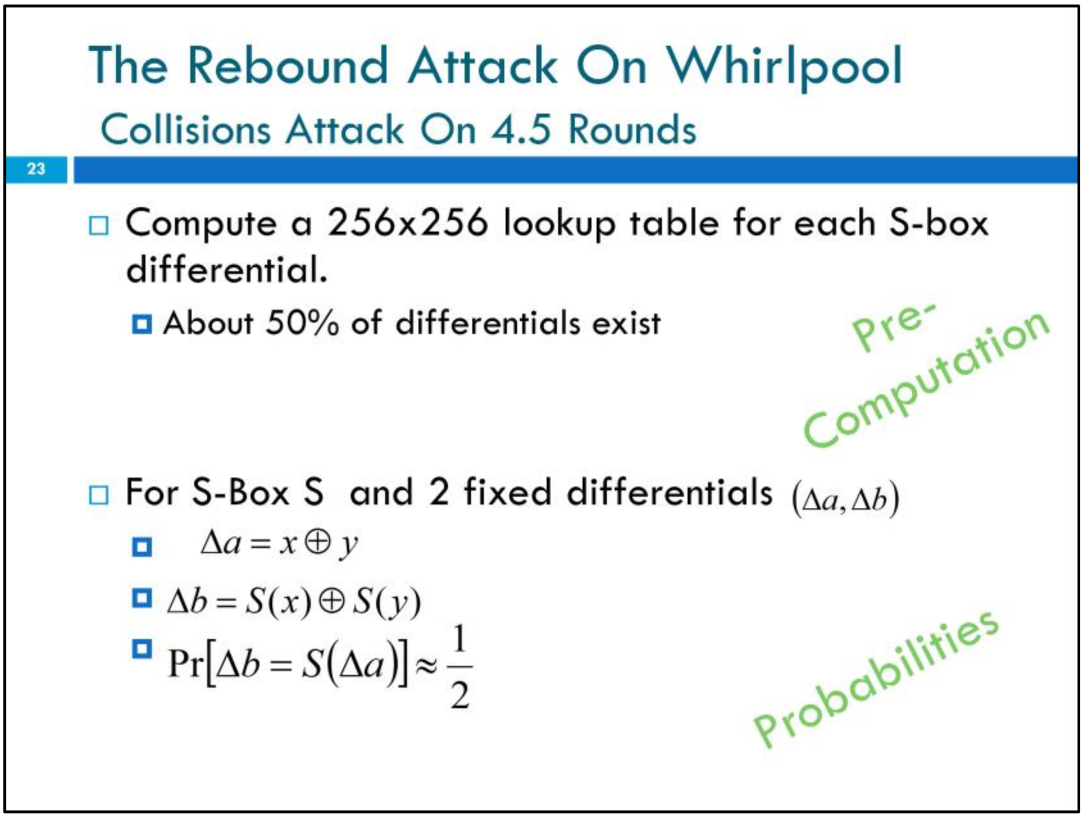

The provability can be verified by enumerating through all 256x256 input/output pairs (x; y) and (S(x); S(y)).
Note that for each possible S-box differential, we get at least the two symmetric values (x; y) and (y; x).
The table: The number of differentials and possible pairs (x; y) for the Whirlpool and AES S-boxes. The first row shows the number of impossible differentials and the last row corresponds to the zero differential.
In the case of Whirlpool, we get for a small fraction of differentials even 8 possible pairs. This corresponds to the maximum probability distribution of the Whirlpool S-box, which is 8*2^8= 2^5

.

The Rebound Attack On Whirlpool

Collisions Attack On 4.5 Rounds

Step 1

24

- Choose random difference at S_2''
- Propagate forward to S_2
- Choose random difference at S_3'''
- Propagate backward S_3'

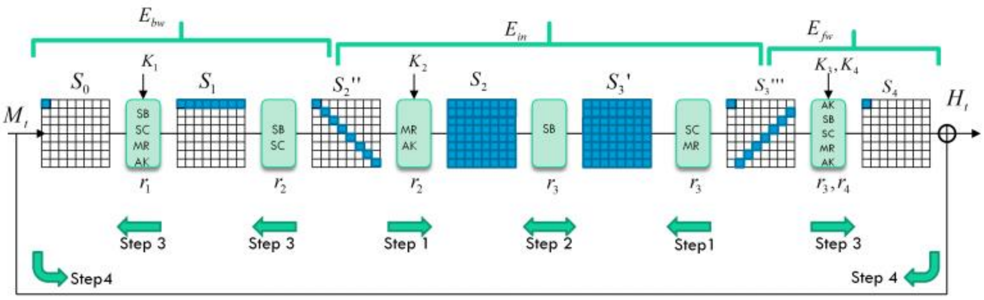

We start the attack by choosing a random difference with 8 active bytes of state S2’’ prior to the MixRows layer of round r2. Note that all active bytes have to be in the diagonal of state S2’’ . Then, the differences propagate forward to a full active state at the input of the next SubBytes layer (state) with a probability of 1. Next, we start with another difference and 8 active bytes in state S3’’’ after the MixRows transformation of round r3 and propagate backwards. Again, the diagonal shape ensures that we get a full active state at the output of SubBytes of round r3.

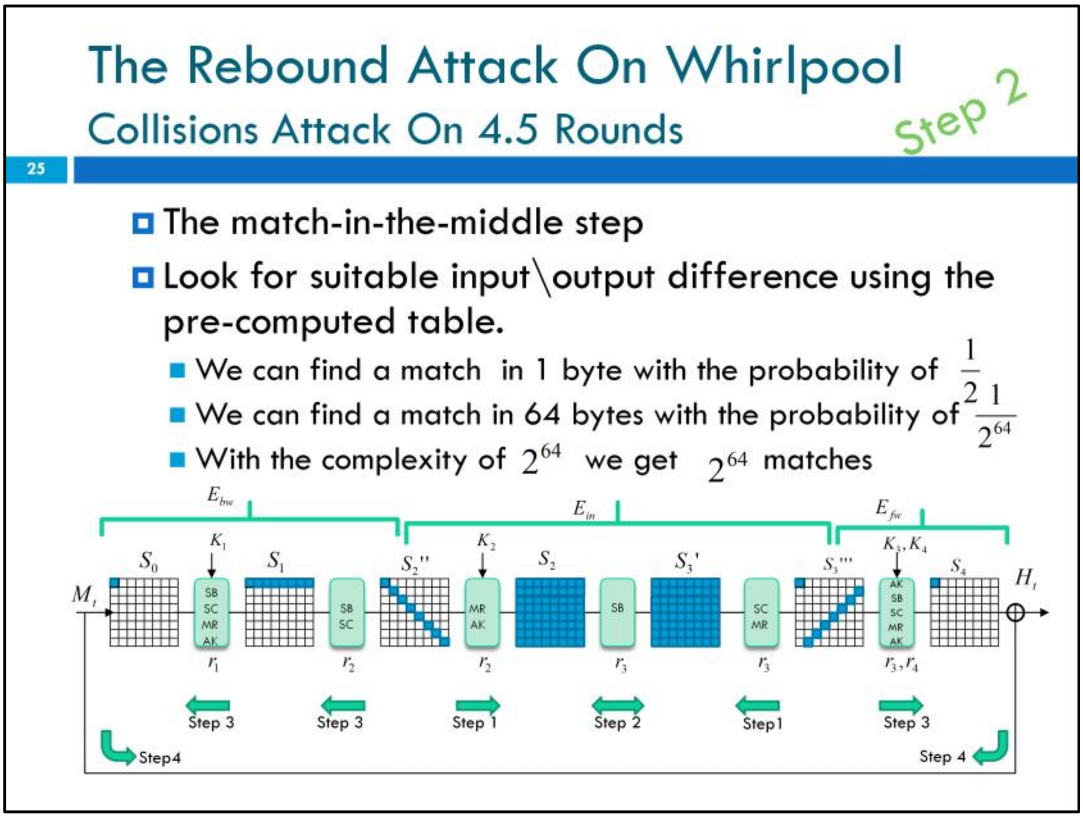

We look for a matching input/output difference of the SubBytes layer of round r3 using the pre-computed S-box differential table.
Since we can find a match with a probability of 0.5 for each byte, we can find a differential for the whole active SubBytes layer with a probability of about 2^-64 .
Hence, after repeating Step 1 of the attack about 2^64 times, we expect to find a SubBytes differential for the whole state.
Each match gets 2-8 posibilities. Since we get at least two state values for each S-box match, we get about 2^64 starting points for the outbound phase

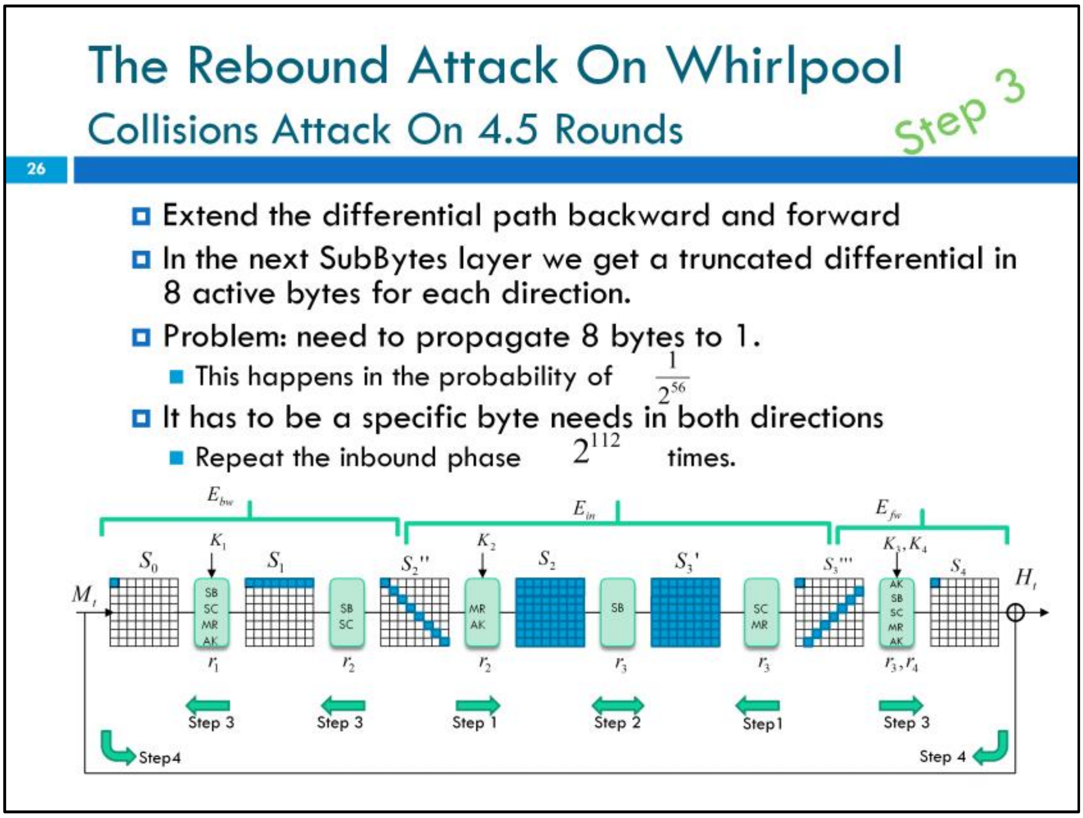

In the outbound phase, we further extend the differential path backward and forward. By propagating the matching differences and state values through the next SubBytes layer, we get a truncated differential in 8 active bytes for each direction. Next, the truncated differentials need to follow a specific active byte pattern. In the case of the 4 round Whirlpool attack, the truncated differentials need to propagate from 8 to one active byte through the MixRows transformation, both in the backward and forward direction. The propagation of truncated differentials through the MixRows transformation is modeled in a probabilistic way. The transition from 8 active bytes to one active byte through the MixRows transformation has a probability of about 2^-56 (7 bytes * 8 bits).
Note that we require a specific position of the single active byte to find a match in the feed-forward (Step 4). Since we need to fulfill one 8 -> 1 transitions in the backward and forward direction, the probability of the outbound phase is 2^-2*56 = 2^-112. In other words, we have to repeat the inbound phase about 2^112 times to generate 2^112 starting points for the outbound phase of the attack.

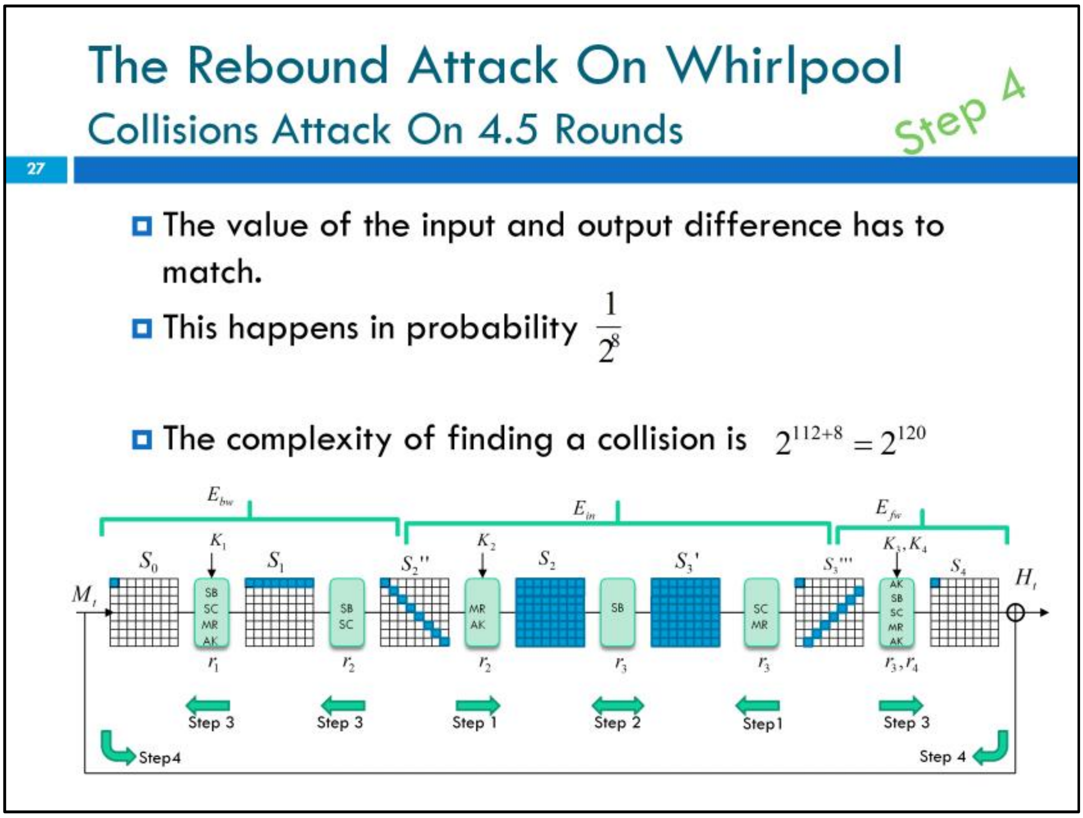

To construct a collision at the output of this 4 round compression function, the exact value of the input and output difference has to match.
Since only one byte is active, this can be fulfilled with a probability of 2^-8. Hence, the complexity to find a collision for 4 rounds of Whirlpool is 2^(112+8)=2^120. Note that we can add half of a round (SB,SC) at the end for free, since we are only interested in the number of active bytes. Remember that we can construct up to 2^128 starting points in the inbound phase of the attack, hence we have enough degrees of freedom for the attack. Note that the values of the key schedule are not influenced.

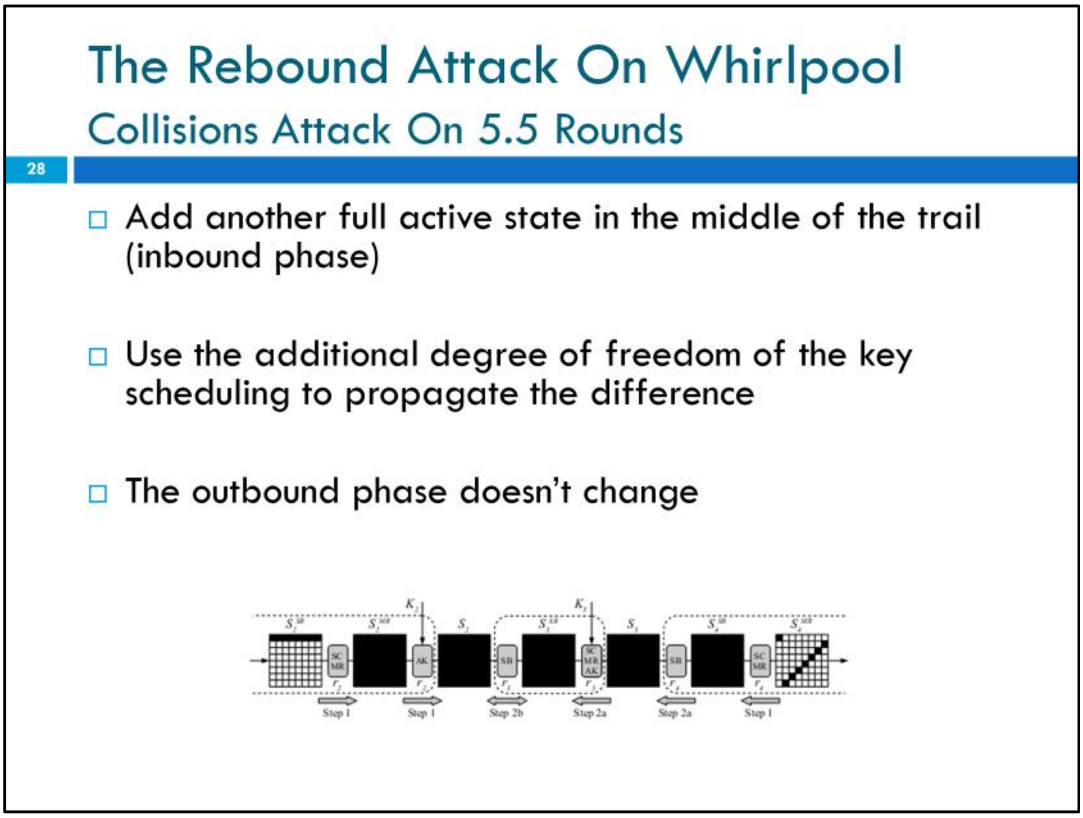

We can extend the collision attack on 4.5 rounds to a semi-free-start collision attack on 5.5 rounds of Whirlpool. The idea is to add another full active state in the middle of the trail. We use the additional degrees of freedom of the key schedule to fulfill the difference propagation through two full active S-box trans-formations.

Note that the outbound part of the attack stays the same and the new sequence of active S-boxes is: 1->8->64->64->8->1->1

The Rebound Attack On Whirlpool

Collisions Attack On 5.5 Rounds

Step 1

29

- Choose initial differences with 8 active bytes at S_2'' and S_4'''
- Propagate differences forwards to S_2 and backwards to S_4'
- Find a matching SubBytes differential of 2 consecutive S-boxs

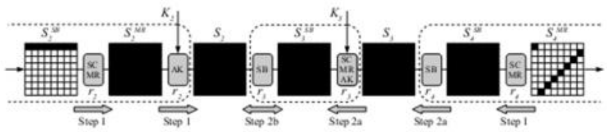

Again, we can choose from up to 2^64 initial differences with 8 active bytes at state S2’’ and S4’’ and linearly propagate forward toS2 and backward to S4 until we hit the first S-box layer. Then, we need to find a matching SubBytes differential of two consecutive S-box layers in the match-in-the-middle phase

The Rebound Attack On Whirlpool

Collisions Attack On 5.5 Rounds

Step 2

30

- Select a value for S_4' (out of 2^{512} possible ones)
- Propagate towards S_3
- Propagate further back to state S_3' with 512 degrees of freedom of the key.

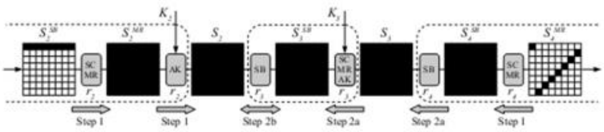

To pass the S-box of round r4 in the backward direction, we choose one of 2^512 possible values for state S4’.
This also determines the input values and differences of the SubBytes layer (state S3).
Then, we propagate the difference further back to state S3’ with 512 degrees of freedom of the key.
That allows us to still assign arbitrary values to the state S3’ .
Hence, the correct difference propagation of the S-box in round r3 can be fulfilled by using these additional degrees of freedom to choose the state S3’ .
The complexity of the attack does not change and is determined by the 2^120 trials of the outbound phase

The Rebound Attack On Whirlpool

Collisions Attack On 5.5 Rounds

31

- Since the outbound phase remains the same the complexity of the attack is still 2^{120}
- The main difference is that the round keys are determined by the inbound phase

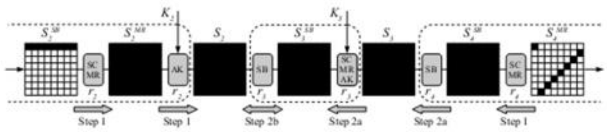

The complexity of the attack does not change and is determined by the 2^120 trials of the outbound phase
The outbound phase (Step 3 and Step 4) of the 5.5 round attack is equivalent
to the 4.5 round case.
However, we cannot choose the round keys, and hence the chaining values, anymore since they are determined by the difference propagation of the S-box of round r3.
Therefore, this 5.5 round attack is only a semi-free-start collision attack on the hash function of Whirlpool.

What's On The Menu?

- Introduction
- Advanced Encryption Standard
- Whirlpool Block Cipher
- Rebound Attack on Whirlpool
- Grøstl Block Cipher
- Rebound Attack on Grøstl
- Summery

What's On The Menu?

- Introduction
- Advanced Encryption Standard
- Whirlpool Block Cipher
- Rebound Attack on Whirlpool
- Grøstl Block Cipher
- Summery

Grøstl Block Cipher

Overview

- An iterated hash function designed by a team of cryp
  Denn

- Grö ver
  pota pan
  toge
  some ash.

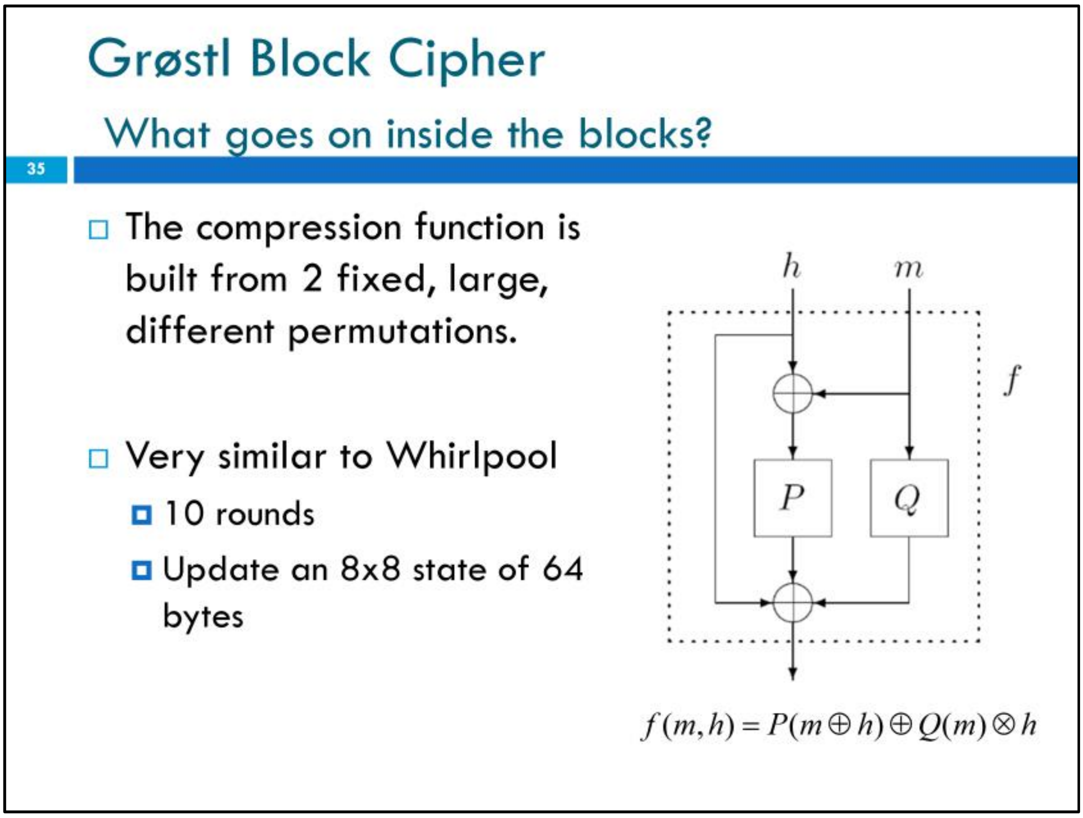

The two permutations P and Q are constructed using the wide trail design strategy and borrow components from the AES. The design of the two permutations
is very similar to the block cipher W used in Whirlpool instantiated with a fixed key input. Both permutations update an 8x8 state of 64 bytes in
10 rounds each.

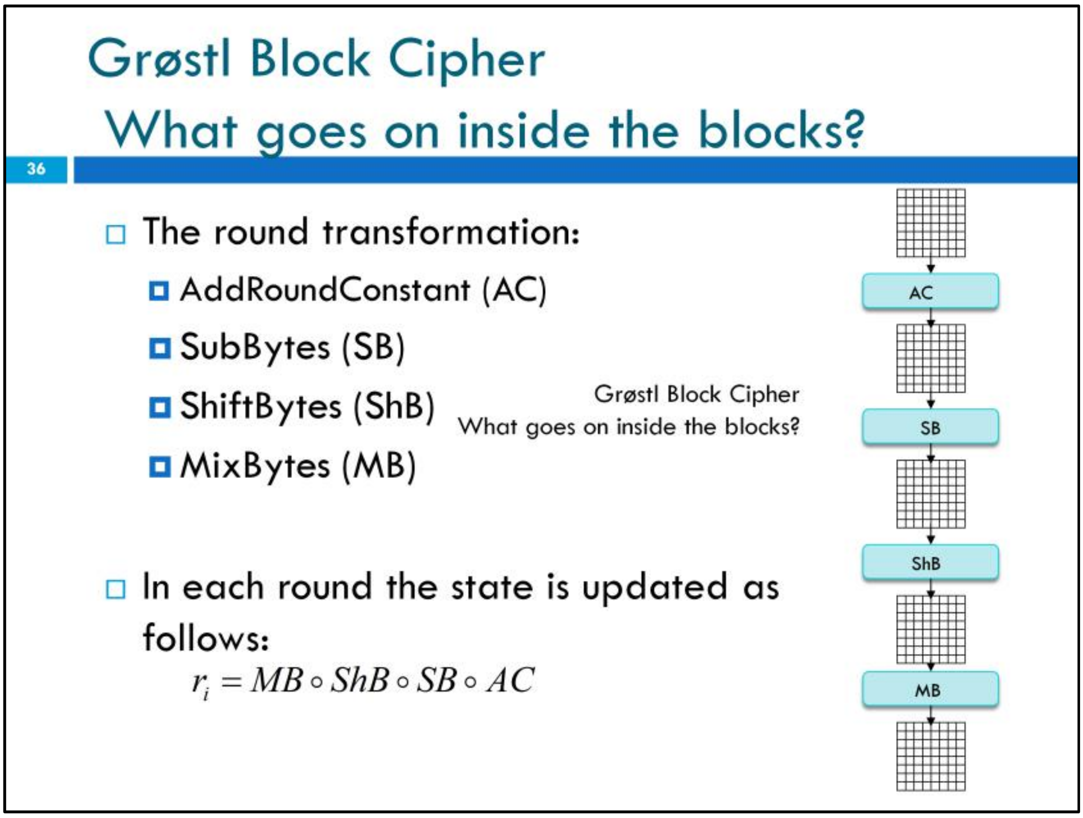

AddRoundConstant (AC) adds different one-byte round constants to the 8x8 states of P and Q. (P & Q have different constants)
the non-linear layer SubBytes (SB) applies the AES S-Box to each byte of the state independently
The cyclical permutation ShiftBytes (ShB) rotates the bytes of row j left by j positions in P, Q
the linear diffusion layer MixBytes (MB) multiplies the state by a constant matrix
In the MixBytes transformation, each column in the matrix is transformed independently by multiplying each column in a constant 8x8 matrix,

Rebound Attack on Grøstl

Rebound Attack On 5 Rounds

- Use the following differential trail

$$8 \xrightarrow{1} 8 \xrightarrow{2} 64 \xrightarrow{3} 8 \xrightarrow{4} 8 \xrightarrow{5} 64$$

- Semi free collisions 2^{120}
- Complexity of the attack

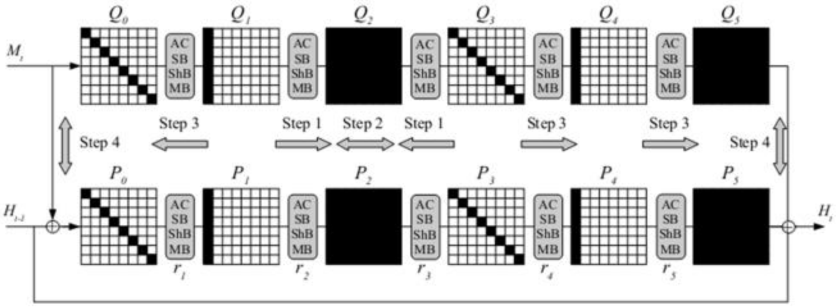

In the attack on 5 rounds, we use the following differential trail for both permutations:
8->8->64->8->8->64
Do not allow diffs in H. all diffs are in m.
By using an equivalent differential trail in the second permutation one can find a collision for the compression function ofGrstl-256 reduced to 5 rounds with a complexity of 2^64
Step 1 and 2 the same as whirlpool only on 2 permutations
Require that the differential output of round 5 are equal
To prevent feed forwards to destroy the collision do not allow differences in H

What's On The Menu?

38

- Introduction
- Advanced Encryption Standard
- Whirlpool Block Cipher
- Rebound Attack on Whirlpool
- Grøstl Block Cipher
- Summery

Summery

- The idea of the rebound attack is to bypass the low probability parts of a hash function differential trail
- AES-based hash functions seem like natural candidates for such an attack.
- Can this attack be applied to a wider range of hash functions?

hash function	rounds	computational complexity	memory requirements	type
Whirlpool	4.5/10	2^{120}	2^{16}	collision
	5.5/10	2^{120}	2^{16}	semi-free-start collision
	7.5/10	2^{128}	2^{16}	semi-free-start near-collision
Grøstl-256	6/10	2^{120}	2^{70}	semi-free-start collision

References

- The Rebound Attack: Cryptanalysis of Reduced Whirlpool and Grøstl \ Mendel, Rechberger, Schlaffer and Thomsen
- Specifications for the Advanced Encryption Standard\ Federal Information Processing Standarf Publication 197
- The Whirlpool Hashing Function \ Paulo S. L. M. Barreto and Vincent Rijmen.
- The Design of Rijndael \ Daemen and Rijmen
- Grøstl – a SHA-3 candidate \ Gauravaram, Knudsen, Matusiewicz, Mendel, Rechberger, Schlaffer and. Thomsen