Skein

John Kevin Hicks

Outline

- Introduction
- Skein Overview
  - Threefish Block Cipher
  - Unique Block Iteration
  - Optional Argument System
- Skein Performance
- Security Claims and Current Cryptanalysis
- Conclusions

Introduction

- SHA-3 candidate submitted by a team headed by Niels Ferguson and Bruce Schneier
- One of the fourteen candidates to move onto the second round of the SHA-3 competition
- Combines “speed, security, simplicity, and a great deal of flexibility in a modular package”
- Skein: a loosely coiled length of yarn or thread

Skein Overview

- Family of hash functions with various state sizes
  - Internal state size of 256, 512, or 1024 bits
  - Output size up to 2^{64} bits
- Naming Convention: Skein-512-160

Internal State Size | Output Size

Replace	With	State Size	Output Size
MD5	Skein-256-128	256	128
	Skein-512-128	512	128
SHA-1	Skein-256-160	256	160
	Skein-512-160	512	160
SHA-224	Skein-256-224	256	224
	Skein-512-224	512	224
SHA-256	Skein-256-256	256	256
	Skein-512-256	512	256
SHA-384	Skein-512-384	512	384
	Skein-1024-384	1024	384
SHA-512	Skein-512-512	512	512
	Skein-1024-512	1024	512

Image From [1]

Skein Overview

- Builds hash function out of a tweakable block cipher
  - Hash configuration data along with input text of every block
  - Makes every instance of the compression function unique

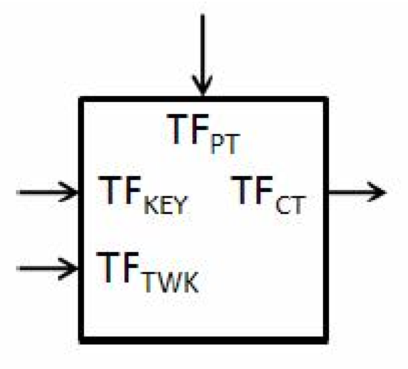

- Unique Block Iteration is the Matyas-Meyer-Oseas (MMO) construction with the tweakable block cipher

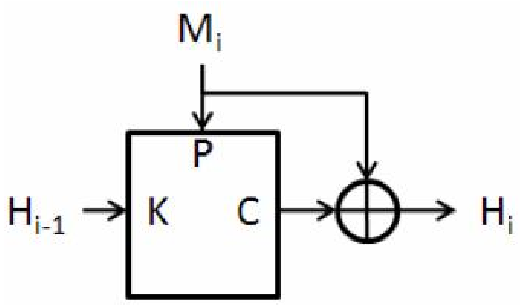

Image From [4]

Threefish Block Cipher

- Design philosophy: Large number of simple rounds provides better security than a small number of complex rounds
- Simple rounds using only addition, exclusive-or, and constant rotations on 64-bit words

Internal State Size	Rounds
256 bits	72
512 bits	72
1024 bits	80

Threefish Block Cipher

- Four Rounds of Skein-512

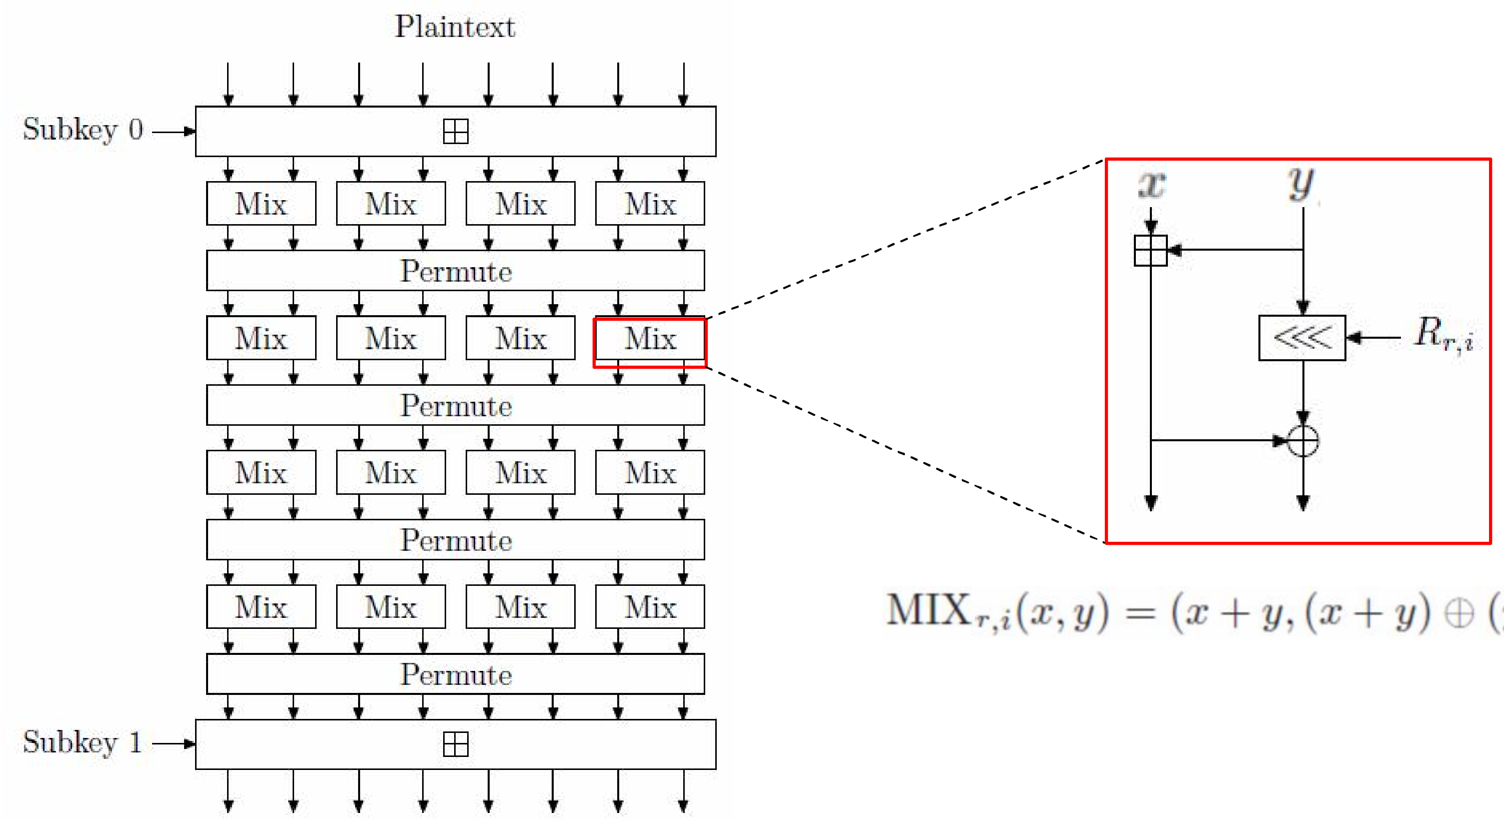

$$\mathrm{MIX}_{r,i}(x,y) = (x+y, (x+y) \oplus (y \lll R_{r,i}))$$

Threefish Block Cipher

- Subkey Generation

Key Scheduling Algorithm		
Inputs:	$K = k_0, \ldots, k_{N_w-1}$	block cipher key split into 64-bit words
	$T = t_0, t_1$	tweak split in two 64-bit words
Outputs:	$sK_s = sk_{s,0}, \ldots, sk_{s,N_w-1}$	subkey s
Algorithm:		
	$k_{N_w} = \lfloor 2^{64}/3 \rfloor \oplus \bigoplus_{i=0}^{N_w-1} k_i$	
	$t_2 = t_0 \oplus t_1$	
	$sk_{s,i} = k_{(s+i) \bmod (N_w+1)}$	for i = $0, \ldots, N_w - 4$
	$sk_{s,i} = k_{(s+i) \bmod (N_w+1)} + t_{s mod 3}$	for i = $N_w - 3$
	$sk_{s,i} = k_{(s+i) \bmod (N_w+1)} + t_{(s+1) mod 3}$	for i = $N_w - 2$
	$sk_{s,i} = k_{(s+i) \bmod (N_w+1)} + s$	for i = $N_w - 1$

Image From [4]

UBI Chaining Mode

- Requires three inputs
  - Chaining variable (G)
  - Message or data portion (M)
  - Tweak value (T_s)
- Tweak value (128 bits) includes the number of processed bytes, flags for the first and/or last block of UBI, type of argument (message, configuration, output)

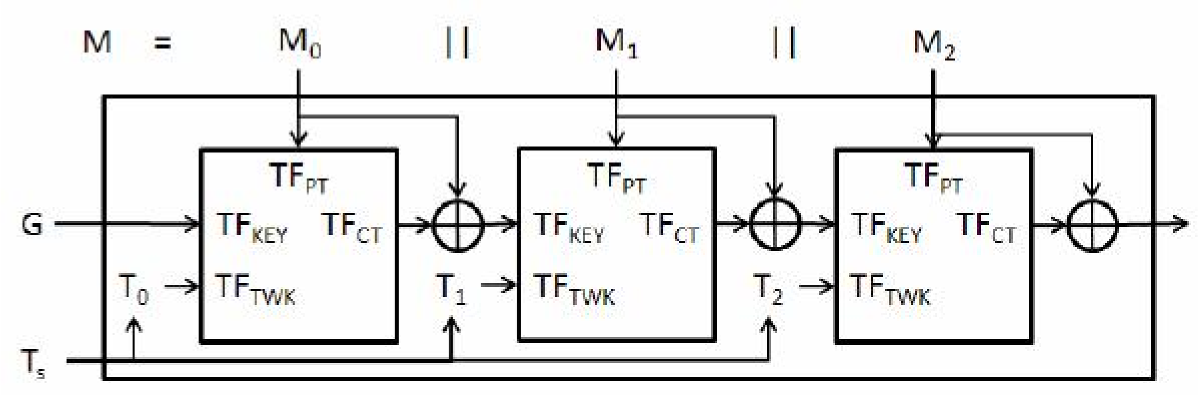

Image From [4]

Skein Hashing

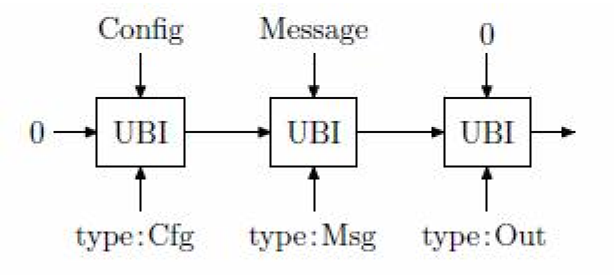

- Config encodes output length and other parameters
- Config output is used as IV for hashing message
- Output transform used to produce desired output size

Image From [1]

Skein Hashing

- Output Size ≤ 3 * Internal State Size

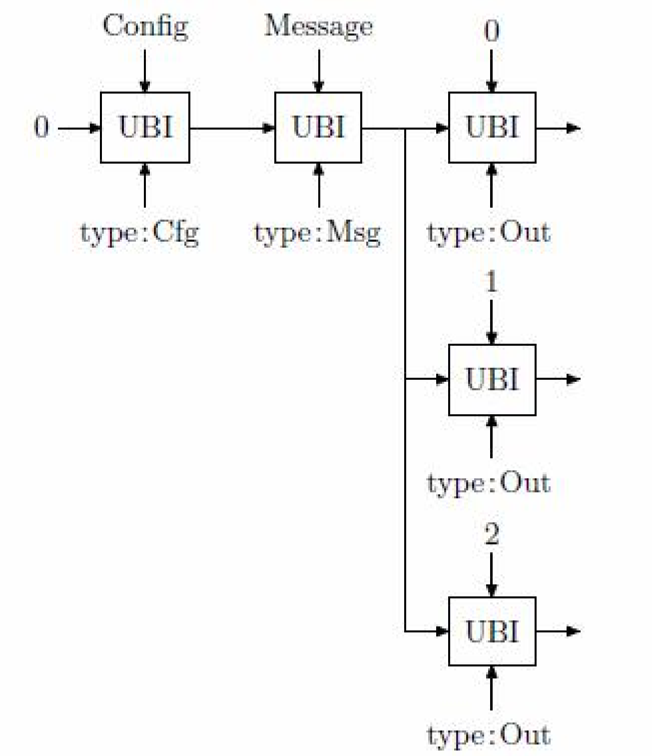

Image From [1]

Optional Argument System

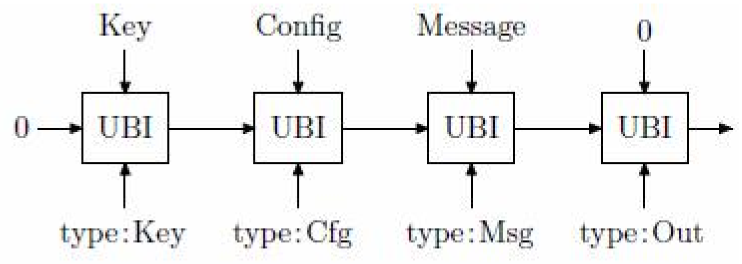

- Key argument can be processed for MAC functionality
- Additional arguments include
  - Personalization
  - Public Key
  - Key Derivation Identifier
  - Nonce

Image From [1]

Performance of Skein

- Recall:

Replace	With	State Size	Output Size
SHA-1	Skein-256-160	256	160
	Skein-512-160	512	160

- Clocks per Byte in C on a 64-bit machine

	Message Length (bytes)					
	1	10	100	1000	10,000	100,000
SHA-1	677	74.2	14.0	10.4	10.0	10.0
SHA-224	1379	143.1	27.4	20.7	20.1	20.0
SHA-256	1405	145.7	27.6	20.7	20.1	20.0
SHA-384	1821	187.3	19.6	13.7	13.4	13.3
SHA-512	1899	192.5	20.6	13.8	13.4	13.3

	Message Length (bytes)					
	1	10	100	1000	10,000	100,000
Skein-256	774	77	16.6	9.8	9.2	9.2
Skein-512	1086	110	15.6	7.3	6.6	6.5
Skein-1024	3295	330	33.2	14.2	12.3	12.3

Image From [1]

Skein Security Claims

- First preimage resistance
  - **Given**: Hash h
  - **Find**: Message M such that $hash(M) = h$
  - **Skein Complexity:** $2^{\min\{\text{ internal state size, output size }\}}$

- Second preimage resistance
  - **Given**: Fixed message M_1
  - **Find**: Different message M_2 such that $hash(M_2) = hash(M_1)$
  - **Skein Complexity:** $2^{\min\{\text{ internal state size, output size }\}}$

Skein Security Claims

- Collision resistance
  - **Find**: Two messages $M_1 \neq M_2$ such that $hash(M_2) = hash(M_1)$
  - **Skein Complexity**: $2^{\min\{\text{ internal state size, output size }\} / 2}$
- R-collision resistance
  - **Find**: R different messages $M_1, \ldots, M_r$ with $hash(M_1) = \ldots = hash(M_r)$
  - **Skein Complexity**: $2^{\min\{ x / 2, (r-1) * y / r\}}$
    - x = internal state size
    - y = min{internal state size, output size}

Skein Security Claims

- Near-Collision
  - **Given**: Hamming weight h
  - **Find**: Two messages $M_1 \neq M_2$ with $hash(M_1) \neq hash(M_2)$, where $n - h$ bits in $hash(M_1)$ and $hash(M_2)$ are the same and h bits differ
  - **Skein Complexity**: No more than

$$\binom{n}{h} = \frac{n!}{h! \cdot (n-h)!}$$

times faster than the corresponding full collision

- Also applies to near-first/second-preimage

Skein Cryptanalysis

- Security proofs used to demonstrate security of Skein based on the assumption that Threefish and the compression function are secure
- Focus should be on these underlying base primitives
- Current cryptanalysis has targeted Threefish-512 with a reduced number of rounds

Skein Cryptanalysis

- Threefish-512 Cryptanalysis

Rounds	Time	Memory	Type
8	1	–	511-bit near-collision
16	2^{6}	–	459-bit near-collision
17	2^{24}	–	434-bit near-collision
17	$2^{8.6}$	–	related-key distinguisher*
21	$2^{3.4}$	–	related-key distinguisher
25	$2^{416.6}$	–	related-key key recovery
26	$2^{507.8}$	–	related-key key recovery
32	2^{312}	2^{71}	related-key boomerang key recovery
34	2^{398}	–	related-key boomerang distinguisher
35	2^{478}	–	known-related-key boomerang distinguisher

- *Recall: Skein-256 and Skein-512 – 72 rounds; Skein-1024 – 80 rounds*

Image From [3]

Conclusion

- Skein provides a fast, simple, secure, and flexible hash algorithm
- Current cryptanalysis suggests 36 or more rounds of Threefish to provide optimal security, well under the 72 or 80 rounds of the current submission
- Skein is a serious contender for selection in the SHA-3 competition

Questions

References

- [1] N. Ferguson, S. Lucks, B. Schneier, D. Whiting, M. Bellare, T. Kohno, J. Callas, J. Walker. The Skein Hash Function Family v1.2. Available online at http://www.skein-hash.info/sites/default/files/skein1.2.pdf, September 2009.

- [2] M. Bellare, T. Kohno, S. Lucks, N. Ferguson, B. Schneier, D. Whiting, J. Callas, J. Walker. Provable Secuirty Support for the Skein Hash Family. Available online at http://www.skein-hash.info/sites/default/files/skein-proofs.pdf, April 2009.

- [3] J. Aumasson, C. Calik, W. Meier, O. Ozen, R. Phan, K. Varici. Improved Cryptanalysis of Skein. In Advances in Cryptology - ASIACRYPT 2009, volume 5912 of Lecture Notes in Computer Science. Springer, Berlin, Germany, 2009.

- [4] A. Schorr. Performance Analysis of a Scalable Hardware FPGA Skein Implementation. Master's thesis, Rochester Institute of Technology, Rochester, New York, February 2010.

- [5] P. Hoffman, B. Schneier. Attacks on Cryptographic Hashes in Internet Protocols. RFC 4270: Network Working Group. November 2005. Available online at http://tools.ietf.org/html/rfc4270