PFLASH - Secure Asymmetric Signatures on Smartcards

Ming-Shing Chen [1] **Daniel Smith-Tone** [2] Bo-Yin Yang [1]

[1]Academia Sinica, Taipei, Taiwan

[2]National Institute of Standards and Technology, Gaithersburg, MD, USA

20th July, 2015

Minimalism

- Low power

Minimalism

- Low power
- No arithmetic coprocessor

Minimalism

- Low power
- No arithmetic coprocessor
- Minimalist Architecture

Minimalism

- Low power
- No arithmetic coprocessor
- Minimalist Architecture
- Little Cost

SFLASH

In 2003, the NESSIE consortium considered SFLASH the most attractive digital signature algorithm for use in low cost smart cards.

SFLASH

In 2003, the NESSIE consortium considered SFLASH the most attractive digital signature algorithm for use in low cost smart cards.

- Very Fast

SFLASH

In 2003, the NESSIE consortium considered SFLASH the most attractive digital signature algorithm for use in low cost smart cards.

- Very Fast
- Naturally resistant to Timing attacks

SFLASH

In 2003, the NESSIE consortium considered SFLASH the most attractive digital signature algorithm for use in low cost smart cards.

- Very Fast
- Naturally resistant to Timing attacks
- Easy algebraic method for resisting SPA and DPA attacks

SFLASH

In 2003, the NESSIE consortium considered SFLASH the most attractive digital signature algorithm for use in low cost smart cards.

- Very Fast
- Naturally resistant to Timing attacks
- Easy algebraic method for resisting SPA and DPA attacks
- Small footprint

PFLASH

SFLASH was completely broken in 2007.

PFLASH

SFLASH was completely broken in 2007.

Multivariate Security

- Differential Structure
- Rank Structure
- Q-rank Structure

PFLASH

SFLASH was completely broken in 2007.

Multivariate Security

- Differential Structure
- Rank Structure
- Q-rank Structure

Speed

PFLASH trades a little speed for "provable" security.
PFLASH is as little as 4 times slower than SFLASH.

PFLASH

SFLASH was completely broken in 2007.

Multivariate Security

- Differential Structure
- Rank Structure
- Q-rank Structure

Speed

PFLASH trades a little speed for "provable" security.
PFLASH is as little as 4 times slower than SFLASH.

PFLASH retains the desirable traits from SFLASH.

The C^* Scheme

The C^* cryptosystem is the simplest example of a "big field" scheme.

The C^* Scheme

The C^* cryptosystem is the simplest example of a "big field" scheme.

Construction

$$\begin{matrix} k \\ | \\ \mathbb{F}_q \end{matrix} \Bigg] n$$

The C^* Scheme

The C^* cryptosystem is the simplest example of a "big field" scheme.

Construction

$$\begin{matrix} k \\ | \\ \mathbb{F}_q \end{matrix} \Bigg] n$$

We can identify $\mathbf{x} \in k$ with $x \in \mathbb{F}_q^n$.

The C^* Scheme

The C^* cryptosystem is the simplest example of a "big field" scheme.

Construction

$$\begin{matrix} k \\ | \\ \mathbb{F}_q \end{matrix} \Bigg] n$$

We can identify $\mathbf{x} \in k$ with $x \in \mathbb{F}_q^n$.

Encryption Scheme

$$y = P(x) = (T \circ f \circ U)x \text{ where } f(\mathbf{x}) = \mathbf{x}^{q^\theta + 1}.$$

Variations on a Theme by Matsumoto-Imai

Several modifications of the C^* scheme:

Variations on a Theme by Matsumoto-Imai

Several modifications of the C^* scheme:

1. HFE - Replace the monomial map with a polynomial.

Variations on a Theme by Matsumoto-Imai

Several modifications of the C^* scheme:

1. HFE - Replace the monomial map with a polynomial.
2. $(-)$ - Remove r of the public equations.

Variations on a Theme by Matsumoto-Imai

Several modifications of the C^* scheme:

1. HFE - Replace the monomial map with a polynomial.
2. $(-)$ - Remove r of the public equations.
3. (p) - Fix d input variables to constant values.

PFLASH Description

$$\begin{array}{ccccccc} & & k & \overset{f}{\longrightarrow} & k & & \\ & & \uparrow & & \downarrow & & \\ & & \phi & & \phi^{-1} & & \\ \mathbb{F}_q^n & \xrightarrow{S} & \mathbb{F}_q^n & \xrightarrow{\overline{f}} & \mathbb{F}_q^n & \xrightarrow{T} & \mathbb{F}_q^n \end{array}$$

Here ϕ is a vector space isomorphism, $f(x) = x^{q^\theta+1}$ is a C^* monomial, S is a codimension d projection (p), T is a codimension r projection $(-)$ and $\overline{f}$ is the quadratic map which makes the diagram commutative.

High Power Performance

Scheme	PK	SK	Sig	Sign	Ver.
PFLASH-62	39,040B	3,937B	244b	288,093c	17,007c
PFLASH-74	72,124B	5,587B	292b	509,355c	23,829c
PFLASH-94	142,848B	8,977B	372b	634,051c	38,044c
ed25519	32B	64B	512b	61,976c	184,992c
ec p256	64B	96B	512b	381,696c	913,848c
RSA 1024	128B	1024B	344b	1,186,912c	33,676c
RSA 2048	256B	2048B	344b	5,134,876c	67,916c

Constant time implementation data for PFLASH with SSE instructions on Intel Xeon E3-1245 v3 3.40 GHz, avg. for 1000 trials. Also listed are comparable data from eBATS http://bench.cr.yp.to/results-sign.html on an Intel Xeon E3-1275 v3 3.50 GHz (same architecture).

Performance Across Platforms

Xeon (Haswell)	Sign (c)	Ver (c)
PFLASH-74	1,253,068	201,598
RSA1024	1,186,912	33,676
Ed25519	61,976	184,992
ECDSAp256	381,696	913,848
ARM Cortex-A8		
PFLASH-74	4,628,701	740,429
RSA1024	7,878,747	3,860,809
Ed25519	819,157	2,594,303
ECDSAp256	5,378,137	6,317,331
MIPS o32		
PFLASH-74	5,710,020	1,105,242
RSA1024	17,756,132	385,956
Ed25519	2,612,848	8,762,140
ECDSAp256	14,586,352	17,535,264

Implementations without vector instructions. Cycles are listed for the instruction sets of Xeon (Haswell), ARM Cortex-A8, and MIPS o32. The cycle-counts of these number theoretic schemes change more dramatically than PFLASH based on the width of the available multiplication instructions.

Security

1. Differential Symmetric/Invariant Structure

Security

1. Differential Symmetric/Invariant Structure
2. Rank

Security

1. Differential Symmetric/Invariant Structure
2. Rank
3. Q-Rank

Discrete Differential

Definition

The *Discrete Differential* of a map $f : k \to k$ is given by:

$$Df(a, x) = f(a + x) - f(x) - f(a) + f(0).$$

Discrete Differential

Definition

The *Discrete Differential* of a map $f : k \to k$ is given by:

$$Df(a, x) = f(a + x) - f(x) - f(a) + f(0).$$

Elementary Properties

1. Linear operator.

Discrete Differential

Definition

The *Discrete Differential* of a map $f : k \to k$ is given by:

$$Df(a, x) = f(a + x) - f(x) - f(a) + f(0).$$

Elementary Properties

1. Linear operator.
2. Reduces complexity of a function: If f is quadratic, Df is bilinear.

Discrete Differential

Definition

The *Discrete Differential* of a map $f : k \to k$ is given by:

$$Df(a, x) = f(a + x) - f(x) - f(a) + f(0).$$

Elementary Properties

1. Linear operator.
2. Reduces complexity of a function: If f is quadratic, Df is bilinear.
3. If f is quadratic, $D(Tf(Ux + c) + d) = D(Tf(Ux))$.

General Linear Symmetries

Definition [*based on* Dubois et al. (2007)]

We say that f satisfies a *general linear symmetry* if the following equation holds:

$$Df(La, x) + Df(a, Lx) = \Lambda_L Df(a, x).$$

General Linear Symmetries

Definition [*based on* Dubois et al. (2007)]

We say that f satisfies a *general linear symmetry* if the following equation holds:

$$Df(La, x) + Df(a, Lx) = \Lambda_L Df(a, x).$$

Definition

We denote the space of linear maps inducing the above symmetry, S_G, and call this the *space of symmetries*.

Classification of the Space of Symmetries for C^*

Multiplicative Symmetry

Since multiplication by σ, $M_\sigma \in S_G$, $k < S_G$.

Classification of the Space of Symmetries for C^*

Multiplicative Symmetry

Since multiplication by σ, $M_\sigma \in S_G$, $k < S_G$.

Theorem

If f is a C^* monomial then S_G equipped with standard multiplication is a k-algebra. Furthermore, if $3\theta \neq n$ then $k \cong S_G$.

Classification of the Space of Symmetries for C^*

Multiplicative Symmetry

Since multiplication by σ, $M_\sigma \in S_G$, $k < S_G$.

Theorem

If f is a C^* monomial then S_G equipped with standard multiplication is a k-algebra. Furthermore, if $3\theta \neq n$ then $k \cong S_G$.

In the case of the C^{*-} scheme, the "large" size of this space allows a portion to "survive" when restrictions are made to the maps inducing symmetry while most arbitrary linear maps are eliminated.

Classification of Maps Inducing Symmetry

Theorem

Let $f(x) = x^{q^\theta+1}$, be a C^* monomial map. Let $\pi x = \sum_{i=0}^{d} x^{q^i}$ and $Mx = \sum_{i=0}^{n-1} m_i x^{q^i}$ be linear. Suppose $Df(Ma, \pi x) + Df(\pi a, Mx) = \Lambda_M Df(\pi a, \pi x)$. If $\theta + d < \frac{n}{2}$, $|n - 3\theta| > d$, and $0 < d < \theta - 1$, then $M = M_\sigma \pi$ for some $\sigma \in k$.

Classification of Maps Inducing Symmetry

Theorem

Let $f(x) = x^{q^\theta+1}$, be a C^* monomial map. Let $\pi x = \sum_{i=0}^{d} x^{q^i}$ and $Mx = \sum_{i=0}^{n-1} m_i x^{q^i}$ be linear. Suppose $Df(Ma, \pi x) + Df(\pi a, Mx) = \Lambda_M Df(\pi a, \pi x)$. If $\theta + d < \frac{n}{2}$, $|n - 3\theta| > d$, and $0 < d < \theta - 1$, then $M = M_\sigma \pi$ for some $\sigma \in k$.

In the case of a codimension 1 projection, this result proves that the only linear symmetries of a pC^{*-} scheme are the trivial scalar symmetries.

First-order Differential Invariant

Definition

Let $f : k \to k$. A first-order differential invariant of f is an $\mathbb{F}_q$-subspace $V \subseteq k$ such that there exists a subspace $W \subseteq k$ of dimension at most $dim(V)$ for which simultaneously $AV \subseteq W$ for all $A \in Span(Df_i)$.

C^* Invariants

Theorem

Let f be a C^* monomial. Then f has no nontrivial first-order differential invariant.

C^* Invariants

Theorem

Let f be a C^* monomial. Then f has no nontrivial first-order differential invariant.

Theorem

Let $f : k \to k$ be a C^* monomial and let $\pi : k \to k$ be a projection. There exist no nontrivial first-order differential invariants of $f \circ \pi$ beyond $ker(\pi)$.

Rank

The corank of the matrix representations of the quadratic forms is low with very high probability. We can show that nonzero matrices with high corank do not occur in the span of the quadratic forms, and thus there is no rank structure to exploit.

Q-Rank

The Q-rank of a quadratic map is the rank of the public key considered as a quadratic form from an n-dimensional representation of k over itself.

Q-Rank

The Q-rank of a quadratic map is the rank of the public key considered as a quadratic form from an n-dimensional representation of k over itself.

Q-rank of PFLASH

Since the C^* monomial has Q-rank 1, we can find a basis for which the composition $f \circ S$ has Q-rank no greater than d.

Q-Rank

The Q-rank of a quadratic map is the rank of the public key considered as a quadratic form from an n-dimensional representation of k over itself.

Q-rank of PFLASH

Since the C^* monomial has Q-rank 1, we can find a basis for which the composition $f \circ S$ has Q-rank no greater than d.

This Q-rank is low, but can only be exploited when the corank of T is 1, and so PFLASH is secure against Q-rank attacks.
From this perspective PFLASH is equivalent to an HFE^- scheme with a far more efficient signing process.

Summary

1. Sig. Sizes: 244, 292, 372 bits.

Summary

1. Sig. Sizes: 244, 292, 372 bits.
2. Public Key: few dozen Kbytes.

Summary

1. Sig. Sizes: 244, 292, 372 bits.
2. Public Key: few dozen Kbytes.
3. Private Key: few Kbytes.

Summary

1. Sig. Sizes: 244, 292, 372 bits.
2. Public Key: few dozen Kbytes.
3. Private Key: few Kbytes.
4. Best known attack: brute force.

Summary

1. Sig. Sizes: 244, 292, 372 bits.
2. Public Key: few dozen Kbytes.
3. Private Key: few Kbytes.
4. Best known attack: brute force.
5. Side-Channel Resistant

Summary

1. Sig. Sizes: 244, 292, 372 bits.
2. Public Key: few dozen Kbytes.
3. Private Key: few Kbytes.
4. Best known attack: brute force.
5. Side-Channel Resistant
6. Time Constant

Summary

1. Sig. Sizes: 244, 292, 372 bits.
2. Public Key: few dozen Kbytes.
3. Private Key: few Kbytes.
4. Best known attack: brute force.
5. Side-Channel Resistant
6. Time Constant
7. Scalable under various criteria.

Summary

1. Sig. Sizes: 244, 292, 372 bits.
2. Public Key: few dozen Kbytes.
3. Private Key: few Kbytes.
4. Best known attack: brute force.
5. Side-Channel Resistant
6. Time Constant
7. Scalable under various criteria.
8. Appropriate for low-power.

Done

Thanks!

Please see the references in the paper.