

# FastFlex (*efflex*)

(A New Fast and Flexible Cryptographic Function)

**Ashish Sharma,**

Army Institute of Technology, Pune, India 411015

**Email:** [ashish.sharma.ait@gmail.com](mailto:ashish.sharma.ait@gmail.com)

**Web site:** <http://fastflex.sourceforge.net>

## Abstract

FastFlex (*pronounced efflex*) is a New Cryptographic Function which can be used to construct Stream Ciphers, Hash Functions, Message Authentication Codes and Block Ciphers which perform considerably well in both hardware and software environments on Processors of varying word lengths. FastFlex uses a modified version of the Salsa20 [1] Hash Function as its core, along with a 256:256 bit mapping generated using a 256\*256 virtual s-box. FastFlex only uses 4 primitive operations, Word Additions, Word Rotations, Word Multiplications and Word XORs, which gives it speed and flexibility. Furthermore Word Multiplications are used only in moderation where the advantages of using them outweigh the slight speed penalty induced by using multiplications over additions and xor. Finally, the design of FastFlex allows significant tradeoffs in Speed, Size and Security in both Hardware and Software. The cryptanalysis of FastFlex is in progress and initial results have been very encouraging. The Soucre Code for the same may be downloaded at <http://fastflex.sourceforge.net>

**Keywords:** Stream Cipher, Deterministic Random Number Generator, Hash Function, Keyed Hash, 192-bit key security

## Contents

<b>1</b>	<b>Introduction</b>
<b>2</b>	<b>FastFlex Design Goals</b>
<b>3</b>	<b>FastFlex Building Blocks</b>
3.1	HashCore
3.2	SAN, 256:256 Bit Mapping
<b>4</b>	<b>FastFlex Design Justifications</b>
<b>5</b>	<b>FastFlex Applications</b>
5.1	FastStream192
5.2	FastHash256
<b>6</b>	<b>WeakFlex</b>
6.1	WeakFlex8
6.2	WeakStream48
<b>7</b>	<b>Comparison with AES</b>
<b>8</b>	<b>Conclusions</b>
<b>9</b>	<b>Appendices</b>

### Definitions:

**Constructions:** Ciphers / Other Networks created using the FastFlex Function at their core.

**Word:** Unless otherwise stated, a Word denotes a 32-bit Word.

## 1 INTRODUCTION

Presented in this paper is FastFlex, a New Fast and Flexible Cryptographic Function. FastFlex uses a Weak Hash Function combined with an irreversible 256:256 bit mapping to produce a strong cryptographic function which can then be used for designing other Stream Ciphers, Hash Functions and Block Ciphers. The advantage of using such an approach is that if the security properties of the underlying cryptographic function are well established, it is much easier to establish the security of the constructions using the function. Furthermore, it allows a large part of the development effort to be shared and reused between various different constructions serving different purposes and thus reduces the overall development and implementation cost of the final construction.

## 2 FASTFLEX DESIGN GOALS

The primary design goals of FastFlex were as follows:

- The cryptographic function must be considerably faster than existing functions offering the same level of security
- It must be possible to implement the cryptographic function on different platforms without any significant performance penalties
- It must be possible to implement the cryptographic function on low end smart cards and other constrained environments
- The cryptographic function must accept a key size of 192 bits.
- The cryptographic function must be suitable for use in Stream Ciphers, Hash Functions and Block Ciphers
- The cryptographic function must have small key setup times
- The cryptographic function should not use any operations which make it inefficient on 32-bit microprocessors
- It should be possible to implement the cryptographic function on 8-bit microcontrollers with only 128 bytes of RAM
- The cryptographic function should have variants with different number of rounds to offer tradeoffs between security and performance

Now that the primary design goals have been specified, the structure of the cryptographic function can be described.

## 3 FASTFLEX BUILDING BLOCKS

The FastFlex Cryptographic Function is made up of two subroutines:

**The HashCore** which iterates a specified number of times during FastFlex depending on the security requirements, and a **Substitution Addition Network (SAN)** which is a highly nonlinear 256:256 bit mapping executing only once during FastFlex, irrespective of the security requirements.

The input to FastFlex is the input to the HashCore. The output of the HashCore is the input to the SAN. The output of the SAN is the output of FastFlex.




Fig. 1. Structure of FastFlex

### 3.1 HASHCORE

The HashCore is central to the design of FastFlex. The requirements for the HashCore are as follows:

- **The HashCore should be very fast** and should only use operations easily implementable on all target Hardware and Software platforms. To this end only Word Additions, Word XOR, Word Rotations and Word Multiplications are used in the HashCore. Furthermore Word Multiplications are used only in moderation where Word Additions and Word XOR operations would either require more time for the same level of diffusion or would not provide the desired diffusion at all.
- The HashCore should not have a very large state. Nor should it have large memory requirements for storage of temporary data.
- The HashCore should have high collision resistance.
- The HashCore need not be cryptographically secure or one-way as the irreversible 256:256 bit mapping SAN prevents direct attacks on it.

Now that the requirements of the HashCore have been specified, it can be defined.

The HashCore has a state size of 8 Words arranged as a 2\*4 array.

Let the Initial state of the HashCore be

$$\begin{Bmatrix} W_0 & W_1 & W_2 & W_3 \\ W_4 & W_5 & W_6 & W_7 \end{Bmatrix}$$

The transformation which works on words in the same row can then be defined as

$$\begin{Bmatrix} Y_1 = W_1 \oplus ((W_0 + W_3) \lll 7), \\ Y_2 = W_2 \oplus ((Y_1 + W_0) \lll 13), \\ Y_3 = W_3 \oplus ((Y_2 + Y_1) \lll 17), \\ Y_0 = W_0 \oplus ((Y_3 + Y_2) \lll 23) \end{Bmatrix}$$

which can be written as

$$\{Y_0, Y_1, Y_2, Y_3\} = \text{Row}(W_0, W_1, W_2, W_3)$$

The transformation working on the second row is

$$\{Y_6, Y_7, Y_4, Y_5\} = \text{Row}(W_6, W_7, W_4, W_5)$$

The second Row thus becomes

$$\begin{Bmatrix} Y_7 = W_7 \oplus ((W_6 + W_5) \lll 7), \\ Y_4 = W_4 \oplus ((Y_7 + W_6) \lll 13), \\ Y_5 = W_5 \oplus ((Y_4 + Y_7) \lll 17), \\ Y_6 = W_6 \oplus ((Y_5 + Y_4) \lll 23) \end{Bmatrix}$$

Where  $\lll$  denotes Left Rotation

And  $\oplus$  denotes bitwise XOR

It can be seen that the second row transformation is the same as the first row transformation with inputs  $\{W_6, W_7, W_4, W_5\}$  instead of  $\{W_0, W_1, W_2, W_3\}$ .

The transformation which works on the words of the same column is defined as

$$\begin{Bmatrix} Y_1 = (Y_1 \otimes Y_{1+4})_L \\ Y_{1+4} = (Y_1 \otimes Y_{1+4})_H \oplus 8b9a7465_{16} \end{Bmatrix}$$

which can be written as

$$\{Y_1, Y_{1+4}\} = \text{Col}(Y_1, Y_{1+4})$$

where  $\otimes$  denotes 32-bit Word Multiplication,

H and L denote the higher and lower words of the product.

That is  $8b9a7465_{16}$  is XORed with the higher word of the result of the multiplication. The result so obtained is stored in  $Y_{1+4}$ . The lower word of the result of the multiplication is stored in  $Y_1$ .

The result of the row transformation, followed by the column transformation is the final state of one round of HashCore. The final state is thus given by

$$\begin{Bmatrix} Y_0 & Y_1 & Y_2 & Y_3 \\ Y_4 & Y_5 & Y_6 & Y_7 \end{Bmatrix}$$

The output of the present round becomes the input to the next round till there are no more rounds of HashCore remaining.

**Please see Appendix 9.4 for an example.**

**Unless otherwise stated, the HashCore iterates 8 times during one round of FastFlex.**




Fig. 2. HashCore

### 3.2 SUBSTITUTION-ADDITION NETWORK (SAN)

The Substitution-Addition Network (SAN) is a one way substitution which serves to hide the internal state of the FastFlex function from an attacker. It must be noted that even though the HashCore is quite difficult to reverse, it is the SAN which guarantees the claimed security properties of FastFlex.

The SAN is constructed using a 8\*32 bit S-Box, which is in turn derived from the Skipjack F-table. Designed by the Information Security Research Centre (ISRC) at the Queensland University of Technology, Australia, the 8\*32 S-box is the same as is used in the Stream Cipher Proposal NLS (Non-Linear Sober) [2], submitted to the 'estream –Stream Cipher Project' [3].

The S-Box has a high degree on Non-Linearity and is not known to possess any algebraic structure. Such an S-Box is known to be resistant to differential and linear cryptanalysis.

Some key features of the 8\*32 S-Box, hereafter referred to as the Qbox are:

- High Non-Linearity of 114
- Low Imbalance
- Very Low Autocorrelation of less than 24
- 32:32 bit mapping
- No known algebraic structure

The Qbox [ ] is as follows:

```

0XA3AA1887, 0XD65E435C, 0X0B65C042, 0X800E6EF4,
0XFC57EE20, 0X4D84FED3, 0XF066C502, 0XF354E8AE,
0XBB2EE9D9, 0X281F38D4, 0X1F829B5D, 0X735CDF3C,
0X95864249, 0XBC2E3963, 0XA1F4429F, 0XF6432C35,
0XF7F40325, 0X3CC0DD70, 0X5F973DED, 0X9902DC5E,
0XDA175B42, 0X590012BF, 0XDC94D78C, 0X39AAB26B,
0X4AC11B9A, 0X8C168146, 0XC3EA8EC5, 0X058AC28F,
0X52ED5C0F, 0X25B4101C, 0X5A2DB082, 0X370929E1,
0X2A1843DE, 0XFE8299FC, 0X202FBC4B, 0X833915DD,
0X33A803FA, 0XD446B2DE, 0X46233342, 0X4FCEE7C3,
0X3AD607EF, 0X9E97EBAB, 0X507F859B, 0XE81F2E2F,
0XC55B71DA, 0XD7E2269A, 0X1339C3D1, 0X7CA56B36,
0XA6C9DEF2, 0XB5C9FC5F, 0X5927B3A3, 0X89A56DDF,
0XC625B510, 0X560F85A7, 0XACE82E71, 0X2ECB8816,
0X44951E2A, 0X97F5F6AF, 0XD9C2B2B3, 0XCE4FF55D,
0XC6B6B214, 0X2B0B83E3, 0X549EA6F5, 0X9DE041AF,
0X792F1F17, 0XF73B99EE, 0X39A65EC0, 0X4C7016C6,
0X857709A4, 0XD6326E01, 0XC7B280D9, 0X5CFB1418,
0XA6AFF227, 0XFD548203, 0X506B9D96, 0XA117A8C0,
0X9CD5BF6E, 0XDCEE7888, 0X61FCFE64, 0XF7A193CD,
0X050D0184, 0XE8AE4930, 0X88014F36, 0XD6A87088,
0X6BAD6C2A, 0X1422C678, 0XE9204DE7, 0XB7C2E759,
0X0200248E, 0X013B446B, 0XDA0D9FC2, 0X0414A895,
0X3A6CC3A1, 0X56FEF170, 0X86C19155, 0XCF7B8A66,
0X551B5E69, 0XB4A8623E, 0XA2BDF3A5, 0XC4F068CC,
0X573A6ACD, 0X6355E936, 0X03602DB9, 0X0EDF13C1,
0X2D0BB16D, 0X6980B83C, 0XFEB23763, 0X3DD8A911,
0X01B6BC13, 0XF55579D7, 0XF55C2FA8, 0X19F4196E,
0XE7DB5476, 0X8D64A866, 0XC06E16AD, 0XB17FC515,
0XC46FEB3C, 0X8BC8A306, 0XAD6799D9, 0X571A9133,
0X992466DD, 0X92EB5DCD, 0XAC118F50, 0X9FAFB226,
0XA1B9CEF3, 0X3AB36189, 0X347A19B1, 0X62C73084,
0XC27DED5C, 0X6C8BC58F, 0X1CDDE421, 0XED1E47FB,
0XCDCD715E, 0XB9C0FF99, 0X4B122F0F, 0XC4D25184,
0XAF7A5E6C, 0X5BBF18BC, 0X8DD7C6E0, 0X5FB7E420,
0X521F523F, 0X4AD9B8A2, 0XE9DA1A6B, 0X97888C02,
    
```

```


0X19D1E354, 0X5ABA7D79, 0XA2CC7753, 0X8C2D9655,
0X19829DA1, 0X531590A7, 0X19C1C149, 0X3D537F1C,
0X50779B69, 0XED71F2B7, 0X463C58FA, 0X52DC4418,
0XC18C8C76, 0XC120D9F0, 0XAF80D4D, 0X3B74C473,
0XD09410E9, 0X290E4211, 0XC3C8082B, 0X8F6B334A,
0X3BF68ED2, 0XA843CC1B, 0X8D3C0FF3, 0X20E564A0,
0XF8F55A4F, 0X2B40F8E7, 0XFEA7F15F, 0XCF00FE21,
0X8A6D37D6, 0XD0D506F1, 0XADE00973, 0XEFBBDE36,
0X84670FA8, 0XFA31AB9E, 0XAEDAB618, 0XC01F52F5,
0X6558EB4F, 0X71B9E343, 0X4B8D77DD, 0X8CB93DA6,
0X740FD52D, 0X425412F8, 0XC5A63360, 0X10E53AD0,
0X5A700F1C, 0X8324ED0B, 0XE53DC1EC, 0X1A366795,
0X6D549D15, 0XC5CE46D7, 0XE17ABE76, 0X5F48E0A0,
0XD0F07C02, 0X941249B7, 0XE49ED6BA, 0X37A47F78,
0XE1CFFFBFD, 0XB007CA84, 0XBB65F4DA, 0XB59F35DA,
0X33D2AA44, 0X417452AC, 0XC0D674A7, 0X2D61A46A,
0XD6C3152A, 0X3E12B7AA, 0X6E615927, 0XA14FB118,
0XA151758D, 0XBA81687B, 0XE152F0B3, 0X764254ED,
0X34C77271, 0X0A31ACAB, 0X54F94AEC, 0XB9E994CD,
0X574D9E81, 0X5B623730, 0XCE8A21EC, 0X37917F0B,
0X8A9B5D6, 0X9697ADF8, 0XF3D30431, 0X5DCAC921,
0X76B35D46, 0XAA430A36, 0XC2194022, 0X22BCA65E,
0XDAEC70BA, 0XDFAE88CC, 0X777BAE8B, 0X242924D5,
0X1F098A5A, 0X4B396B81, 0X55DE2522, 0X435C1CB8,
0XAE8FE1D, 0X9DB3C697, 0X5B164F83, 0XE0C16376,
0XA319224C, 0XD0203B35, 0X433AC0FE, 0X1466A19A,
0X45F0B24F, 0X51FDA998, 0XC0D52D71, 0XFA0896A8,
0XF9E6053F, 0XA4B0D300, 0XD499CBCC, 0XB95E3D40
    
```

The SAN function uses a virtual S-Box constructed using the Qbox. The Virtual S-box, Vbox, provides even higher resistance against differential and linear cryptanalysis.

It is made up of fundamental operations which do not commute and should thus provide resistance to a wide spectrum of attacks.

Let us define the Vbox.

- The Vbox is a 256:256, one way bit mapping
- The Vbox uses only fundamental operations which do not commute
- The Vbox is designed for maximum parallelization while maintaining a high level of security



The Vbox consists of 4 layers of operations and is essentially a one way mapping.

Layer 1:

$$\begin{aligned}
 \mathbf{A} &= [\mathbf{A} + \mathbf{S}(\mathbf{B}) + \mathbf{F}] \\
 \mathbf{B} &= [\mathbf{B} + \mathbf{S}(\mathbf{C}) + \mathbf{G}] \\
 \mathbf{C} &= [\mathbf{C} + \mathbf{S}(\mathbf{D}) + \mathbf{H}] \\
 \mathbf{D} &= [\mathbf{D} + \mathbf{S}(\mathbf{A}) + \mathbf{E}]
 \end{aligned}$$

Layer 2:

$$\begin{aligned}
 \mathbf{E} &= [\mathbf{E} + \mathbf{S}(\mathbf{A})] \\
 \mathbf{F} &= [\mathbf{F} + \mathbf{S}(\mathbf{B})]
 \end{aligned}$$

$$G = [G + S(C)]$$

$$H = [H + S(D)]$$

Layer 3:

$$A' = [A \oplus S(G)]$$

$$B' = [B \oplus S(H)]$$

$$C' = [C \oplus S(E)]$$

$$D' = [D \oplus S(F)]$$

Layer 4:

$$E' = [E + A'] \oplus F$$

$$F' = [F + B'] \oplus G$$

$$G' = [G + C'] \oplus H$$

$$H' = [H + D] \oplus E$$

Where  $S()$  is defined as

$$S(X) = [Qbox(X_H) \oplus X]$$

and  $X_H$  is the most significant byte of  $X$ .




Fig. 3. The Virtual S-box (V-box)

Unless otherwise stated, **SAN executes only once during FastFlex.**

Now that both the HashCore and the SAN have been specified, let us summarize FastFlex:

- FastFlex takes 8 words as input and produces 8 words as output.
- A round of FastFlex translates to 8 rounds of HashCore followed by 1 round of SAN.
- The HashCore takes 8 words as input and produces 8 words as output.
- The SAN acts as a virtual S-Box It takes 8 words as inputs and produces 8 words as output.
- FastFlex is guaranteed to be a one to one mapping of 256:256 bits. No collisions have been found during the initial analysis of FastFlex.

- We denote the FastFlex function as

$$\{A', B', C', D', E', F', G', H'\} = \text{FastFlex}(W_0, W_2, W_5, W_7, W_1, W_3, W_4, W_6)$$

FastFlex modifies 8 words in place.

#### 4 FASTFLEX DESIGN JUSTIFICATIONS

- **The Use of Primitive Operations:** Only primitive operations such as Word Addition, Word XOR, Word Rotations and Word Multiplications are used in FastFlex. A long chain of simple operations can approximate any complex operation [4]. On the other hand, implementing complex functions in constrained environments is difficult. Most constrained environments, more so than not, already have architectures for implementing simple operations such as Addition, Xoring, Rotation and Multiplication The chain of simple operations as used in FastFlex can thus reach the same desired security level as of any other more complex operation, and at the same time be implementable in constrained environments.
- **The use of Word Multiplications:** Integer multiplications, it can be argued are slow and difficult to implement in Hardware, although on most modern general purpose processors, they are quite fast. Some new hardware chips such as the Xilinx Spartan 3 [5] and Altera Cyclone have efficient multiplier architectures pre-fabricated on chip, and this is a trend which seems to be continuing. Thus it is a reasonable assumption that most future hardware chips will also boast such efficient multiplication architectures and thus the performance penalty mentioned above will soon disappear. At present, the advantage of using word multiplications for diffusion on general purpose processors and the fact that future hardware chips will implement word multiplications efficiently leaves no reason not to use them.
- **The use of Word Rotations:** It may be argued that some microprocessors do not have efficient Word Rotations or do not have Word Rotations at all. To overcome this, rotations may be achieved using bit-shifts on such processors.  

$$X \lll Y = [(X \lll Y) | (X \gg 32 - Y)]$$
- **Rotation Distances:** The exact rotation distances do not matter much. Rotation by any different prime amounts should suffice.
- **The use of S-boxes:** An S-box provides thorough diffusion and confusion in a single operation. S-boxes however take up considerable memory which may increase the cache pressure on some CPUs. Furthermore, constrained environments may not have enough space to implement large S-boxes. A balance is thus needed between the number of S-boxes, the size of S-boxes and the frequency of use of S-boxes in code. The frequency of use of S-boxes is a fact

often neglected in many implementations. S-boxes are implemented as array lookups and using an S-box is equivalent to fetching an element from memory. If S-boxes are not frequently used, and they are frequently overwritten in cache, then there can be significant performance degradation as they will have to be fetched from memory again. FastFlex overcomes this problem by using only a single 8\*32 S-box which requires only 1KB of storage, and using the S-box thoroughly when it is brought into cache.

- **Size of the Internal State of FastFlex:** FastFlex has an internal state of 256 bits. It must be noted that a state size of 256 bits is more than overwhelming considering the amount of processing power available today and the present state of cryptanalysis [6]. Only a major breakthrough in electronics technology or cryptanalytic techniques, if not both, will change this fact. It seems more reasonable to use better design principles than to increase the state size as of now.
- **The Number of Rounds:** 8 rounds of HashCore and 1 round of SAN seems to be quite conservative and as peer confidence in FastFlex increases, the number of rounds can be decreased. Meanwhile, decreasing the number of rounds is not advisable.
- **The SAN:** SAN hides the internal state of the HashCore. This makes direct guess and determine attacks very difficult against FastFlex.
- **Why not more S-boxes:** A single S-box cannot provide the same security as provided by two or more S-boxes. However, using additional S-boxes would have nearly doubled the memory requirements of the FastFlex object code. Furthermore, S-box lookups are vulnerable to timing attacks on most platforms and there is no efficient workaround to this [7]. S-box lookups must thus be kept to a minimum whenever possible.

## 5 FASTFLEX APPLICATIONS

**In this section are presented a few constructions using the FastFlex Cryptographic Function.**

The following are constructed using the FastFlex cryptographic function:

- **FastStream192:** The construction of a synchronous stream cipher providing 192 bit key security.
- **FastHash256:** A hash function taking input in 256 bit blocks and producing a 256 bit message digest.
- **KeyHash256:** A keyed hash function producing a 256 bit message digest.

### 5.1 FASTSTREAM192

**FastStream192** is a synchronous keystream generator. In a synchronous keystream generator, the sender and receiver generate the same keystream using identical copies of the generator and using the same key. The requirements of a good keystream generator are as follows:

- The generated keystream should be unique for every different key.
- The generated keystream should pass all known randomness tests and should thus be indistinguishable from a random message source.
- It should not be practical to recover the secret key by observing the keystream corresponding to that key.

It must be noted that in theory it is not possible for any generator to output a truly random message stream using any fixed series of operations, no matter how complex each operation is taken to be, if the input to the generator is not random. We suffice with what is sufficient randomness, or pseudo randomness, as required by the application.

For a stream cipher using a single key, a keystream length of period  $2^{64}$  bits is considered sufficient.

Let us assume such a keystream generator G generating keystream  $G_T$ .




Fig. 4. Keystream Generator

Encryption then takes place at the transmitter by XORing the message to be encrypted with the corresponding Keystream Word  $G_T$ .

$$C_T = M_T \oplus G_T$$

where  $C_T$  is the transmitted ciphertext corresponding to the Message  $M_T$

Decryption is done by XORing the ciphertext  $C_T$  with a copy of the corresponding keystream  $G'_T$ , generated locally by the receiver using the same key.

$$M'_T = G'_T \oplus C_T$$

If  $G'_T$  and  $G_T$  are identical, then it can be seen that  $M'_T$  and  $M_T$  are also identical.

At first, XOR might not seem a formidable operation for encryption, but the only theoretically secure stream cipher(or for that matter any cipher) the **One-Time Pad** [8] uses XOR as the only operation.

Now that synchronous stream ciphers have been introduced, FastStream192 can be described.

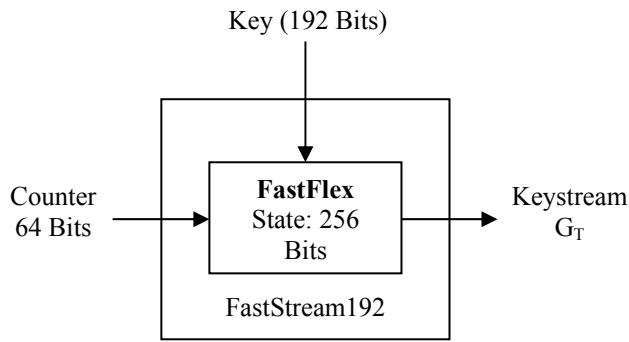


Fig. 5. FastStream192 Keystream Generator

**At the Transmitter (Encryption):**

$$C_T = G_T \oplus M_T$$

**At the Receiver (Decryption):**

$$M_T = G_T \oplus C_T$$

Synchronized copies of the keystream generator run at both the transmitter and the receiver.

5.1.A KEYSTREAM GENERATION IN FASTSTREAM192

Keystream generation in FastStream192 is done by one round of FastFlex. The keystream generator takes two inputs, a 192 bit secret key  $K \{K_0, K_1, K_2, K_3, K_4, K_5\}$  and a 64 bit nonce  $N \{N_0, N_1\}$ , where  $K_0, N_0$  are the least significant words of  $K$  and  $N$  respectively.

The process of keystream generation is as follows:

- The inner state of the FastFlex is loaded with  $K$  and  $N$  as shown below

$$\left\{ \begin{array}{cccc} N_0 & K_0 & K_1 & K_2 \\ K_3 & K_4 & K_5 & N_1 \end{array} \right\}$$

- FastFlex is allowed to run four times with the output of the present round becoming the inner state of FastFlex for the next round.

$$\left\{ \begin{array}{cccc} N_0 & K_0 & K_1 & K_2 \\ K_3 & K_4 & K_5 & N_1 \end{array} \right\}$$

↓ 4 Rounds of FastFlex

$$\left\{ \begin{array}{cccc} A & B & C & D \\ E & F & G & H \end{array} \right\}$$

- $\{A, B, C, D, E, F, G, H\}$  becomes the internal state of FastStream192.  $\{H, A\}$  form a 64 bit counter, with a  $\{H\}$

as the higher word and  $\{A\}$  as the lower word of the counter.

- The keystream generator is now ready to produce keystream.
- A single round of FastFlex produces 8 words of keystream.

$$\left\{ \begin{array}{cccc} A & B & C & D \\ E & F & G & H \end{array} \right\}$$

↓ 1 Round of FastFlex

$$\left\{ \begin{array}{cccc} G_0 & G_1 & G_2 & G_3 \\ G_4 & G_5 & G_6 & G_7 \end{array} \right\}$$

or in other words

$$\{G\} = \text{FastFlex}(ABCDEFGHIH)$$

- The counter is updated for generating the next 8 words of keystream by incrementing  $ah$  by 1. The value of the counter to be loaded for the next round of encryption thus becomes  $HA + 1$ .
- On generating every 8 words of keystream, the counter  $HA$  is again incremented by 1 before generating the next words of the keystream.

5.1.B FASTSTREAM192 DESIGN JUSTIFICATIONS

- Key Size of 192 bits:** Although it is possible to implement a larger key size by increasing the state size of FastFlex (FastFlex scales well to increase in size of its state), a 192 bit key seems to be quite secure for quite some time to come. It is still impossible to brute force even a 96 bit key being used in a secure implementation. What is required is more resistance to cryptanalytic tools.
- Stream Independent of Plaintext and Ciphertext:** The success of many cryptanalytic attacks depends on the ability of the cryptanalyst to induce/toggle some bits in the internal state of the cipher. A synchronous cipher using a counter is thus more secure against such attacks as the cryptanalyst cannot induce bits into its internal state. Furthermore, making the stream dependent on ciphertext forces serialization of the cipher. A serial cipher cannot enjoy the benefits of parallelization.
- Synchronous Cipher:** For a synchronous cipher to decrypt data correctly, perfect synchronization is required between the transmitter and the receiver. Whenever there is any insertion in the stream, synchronization is lost. Thus synchronous ciphers are immune to Replay Attacks and most Man in the Middle Attacks [9]. Any errors in transmission when using synchronous stream ciphers are restricted to individual bits/words and are not propagated. This is an added advantage in some scenarios where bit errors prevail.

### 5.1.C PERIOD OF KEYSTREAM FOR FASTSTREAM192

FastStream192(K,N) denotes the FastStream192 stream cipher function acting on the 24 Byte key K and the 8 Byte counter sequence N, N + 1, ... N + (2<sup>64</sup> - 1).

FastStream192(192,64) can thus have a sequence length of up to (2<sup>64+8</sup> = 2<sup>72</sup>) bits, as the sequence N, N + 1, ... N + (2<sup>64</sup> - 1) has a length of 2<sup>64</sup> and each round generates 8 words of keystream. Thus it is possible to generate sequences of period 2<sup>72</sup>. However Collisions may occur in FastFlex. No such collisions have been found till now, but it is unreasonable to assume that no collisions will be found in the future. **Following a conservative design approach, the keystream for any given Key-Nonce Pair is thus limited to 2<sup>64</sup> bits,** which is same as the keystream period for AES for a given 128 bit key in counter mode.

### 5.2 FASTHASH256

FastHash256 is an Unkeyed Hash Function. An unkeyed hash function maps a message M of arbitrary but finite length, into a message digest of fixed length N. M >> N in most cases, a one to one mapping is not possible and two or more inputs are bound to produce the same hash N. When this happens a collision has occurred. That is the hash of two or more different messages has collided.

It is theoretically impossible to avoid collisions. Nonetheless a few properties are expected of practical hash functions [10]:

- Given any M, it must be easy to compute N = Hash(M).
- **Preimage Resistance:** Given N, it must be infeasible to find any message M such that Hash(M) = N.
- **2<sup>nd</sup> Preimage Resistance:** Given a Message M, it must be infeasible to find a Message M' such that Hash(M) = Hash(M').
- **Collision Resistance:** It must be infeasible to find two distinct messages M and M' such that Hash(M) = Hash(M').
- **Non-Correlation:** The correlation between the input and output bits of the hash function should be very low.
- **Avalanche:** A hash function should have good avalanche properties such that every bit of input affects every bit of output.
- **Partial Preimage Resistance:** It should be difficult to reconstruct any part of the message from the Hash even if a part of the input message is known.

Now that the desirable properties of Unkeyed Hash Functions have been listed, the specification for FastHash256 can be given.

FastHash256 takes input messages in blocks of 256 bits in the form of 8 words and produces as output a Hash of 256 bits in the form of 8 words. Messages are appended with a right justified binary representation of the length of the original message in bits prior to hashing (Merkle-Damgård

Strengthening [11]). If the length of the message is such that the last block cannot be appended with the length of the message, then the last block is padded with zero and an additional block is inserted, following which Merkle-Damgård Strengthening can be applied.

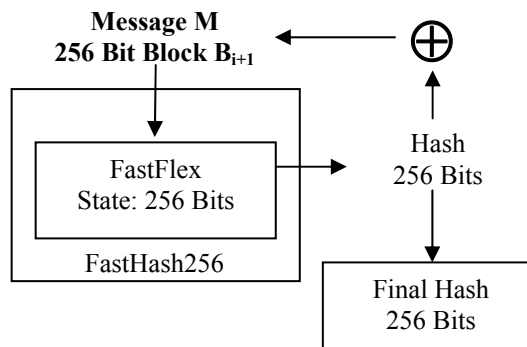


Fig. 6. FastHash256

which can be written as

$$\text{Hash} = \text{FastHash256}(M')$$

where M' is the Message M after MD-strengthening. The procedure for MD-strengthening is as follows:

- The original Message M is divided into (N - 1) 256 bit blocks B<sub>0</sub>, B<sub>1</sub>, ... B<sub>N-1</sub>.
- A length of the Message in bits is appended to the last block B<sub>N-1</sub>. A more thorough explanation of MD-strengthening can be found in [12].
- If the last block cannot hold the length of the message in bits, it is appended with 0's and an additional block is added which is right appended with the Message length in bits. That is, if the bits of the block are {b<sub>0</sub>, b<sub>1</sub>, ... b<sub>255</sub>}, and if K bits are required to represent the Message length, the last message block becomes {b<sub>0</sub>, b<sub>1</sub>, ... b<sub>255-K</sub>}. If the last message block cannot hold K additional bits, an additional block of all 0's is appended to the original message M. The last message block then becomes {0, 0, ... K bits}.
- As a matter of convention, the length K is taken as the word length, that is 32-bits and is appended as 4-bytes in Little Endian.

FastHash256 uses a single iteration of FastFlex. After hashing block B<sub>i</sub>, the hash obtained is XORed with the next block B<sub>i+1</sub> till there are no more blocks remaining.

$$\text{Hash} = \text{FastHash256}(B_i \oplus \text{FastHash256}(B_{i-1} \oplus \text{FastHash256}(B_{i-2} \oplus \text{FastHash256}(B_{i-3} \oplus \dots))))$$



WeakFlex8 is an intentionally weakened version of FastFlex for analysis in academic circles. WeakFlex8 uses a key size of 32 bits, manipulates 8 bit bytes instead of 32 bit words.

The two differences in the structure of WeakFlex8 and FastFlex are:

- Rotation distances in the Row() function of HashCore and in the SAN.
- Use of an 8\*8 S-box derived from Skipjack instead of a 8\*32 S-box as building block of the Virtual S-box.

It should be noted that the structure of FastFlex has not been changed. The numbers of rounds of the HashCore remain the same in WeakFlex8. However, the 48 bit key used in the stream cipher constructed using WeakFlex8 should make cryptanalysis considerably easier and lend more insight into the security properties of FastFlex.

6.1 WEAKFLEX8

Being already familiar with the design of FastFlex, we now describe WeakFlex8.

- The HashCore of WeakFlex8 takes as input eight 8-bit bytes and as output produces eight 8-bit bytes.
- The Row() function of HashCore in WeakFlex8 is:

$$\{Y_0, Y_1, Y_2, Y_3\} = \text{Row}(W_0, W_1, W_2, W_3)$$

where

$$\left\{ \begin{array}{l} Y_1 = W_1 \oplus (W_0 + W_3) \lll 1 \\ Y_2 = W_2 \oplus (Y_1 + W_0) \lll 3 \\ Y_3 = W_3 \oplus (Y_2 + Y_1) \lll 5 \\ Y_0 = W_0 \oplus (Y_3 + Y_2) \lll 8 \end{array} \right\}$$

and + denotes addition Mod  $2^8$ ,  
 $\oplus$  denotes bitwise XOR  
 $x \lll y$  denotes left rotation of byte x by y positions.

- The transformations on elements in the same row are  $\{Y_0, Y_1, Y_2, Y_3\} = \text{Row}(W_0, W_1, W_2, W_3)$   $\{Y_6, Y_7, Y_4, Y_5\} = \text{Row}(W_6, W_7, W_4, W_5)$
- This is followed by transformations on elements in the same column,  $\{Y_i, Y_{i+4}\} = \text{Col}(Y_i, Y_{i+4})$

where

$$\left\{ \begin{array}{l} Y_i = (Y_i \otimes Y_{i+4}) \& 0x00FF \\ Y_{i+4} = S(Y_i \otimes Y_{i+4}) \end{array} \right\}$$

- The HashCore iterates 8 times during a round of WeakFlex8.
- The SAN for WeakFlex8 is defined as a 64:64 bit mapping.
- The SAN executes only once during a round of WeakFlex8.

- The present 8 byte output of HashCore becomes the input for the next round, till no more blocks remain

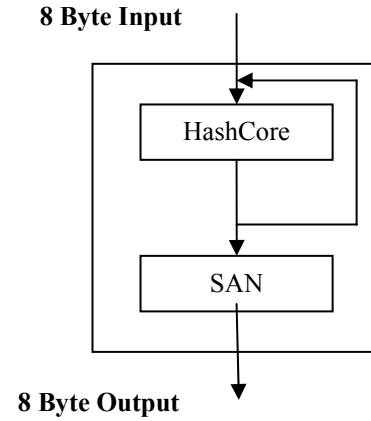


Fig. 8. WeakFlex8

The SAN uses a 8\*8 S-box derived from Skipjack (F-Table) [14]. The key features of the S-box are:

- High Nonlinearity
- 8:8, one to one bit mapping
- No known algebraic structure
- 39.9% entries have non zero elements.


The S-box [ ] is as follows:

0XA3	0XD7	0X09	0X83	0XF8	0X48	0XF6
0XF4	0XB3	0X21	0X15	0X78	0X99	0XB1
0XAF	0XF9	0XE7	0X2D	0X4D	0X8A	0XCE
0X4C	0XCA	0X2E	0X52	0X95	0XD9	0X1E
0X4E	0X38	0X44	0X28	0X0A	0XDF	0X02
0XA0	0X17	0XF1	0X60	0X68	0X12	0XB7
0X7A	0XC3	0XE9	0XFA	0X3D	0X53	0X96
0X84	0X6B	0XBA	0XF2	0X63	0X9A	0X19
0X7C	0XAE	0XE5	0XF5	0XF7	0X16	0X6A
0XA2	0X39	0XB6	0X7B	0X0F	0XC1	0X93
0X81	0X1B	0XEE	0XB4	0X1A	0XEA	0XD0
0X91	0X2F	0XB8	0X55	0XB9	0XDA	0X85
0X3F	0X41	0XBF	0XE0	0X5A	0X58	0X80
0X5F	0X66	0X0B	0XD8	0X90	0X35	0XD5
0XC0	0XA7	0X33	0X06	0X65	0X69	0X45
0X00	0X94	0X56	0X6D	0X98	0X9B	0X76
0X97	0XFC	0XB2	0XC2	0XB0	0XFE	0XDB
0X20	0XE1	0XEB	0XD6	0XE4	0XDD	0X47
0X4A	0X1D	0X42	0XED	0X9E	0X6E	0X49
0X3C	0XCD	0X43	0X27	0XD2	0X07	0XD4
0XDE	0XC7	0X67	0X18	0X89	0XCB	0X30
0X1F	0X8D	0XC6	0X8F	0XAA	0XC8	0X74
0XDC	0XC9	0X5D	0X5C	0X31	0XA4	0X70
0X88	0X61	0X2C	0X9F	0X0D	0X2B	0X87
0X50	0X82	0X54	0X64	0X26	0X7D	0X03
0X40	0X34	0X4B	0X1C	0X73	0XD1	0XC4
0XFD	0X3B	0XCC	0XFB	0X7F	0XAB	0XE6
0X3E	0X5B	0XA5	0XAD	0X04	0X23	0X9C

Fig

0X14	0X51	0X22	0XF0	0X29	0X79	0X71
0X7E	0XFF	0X8C	0X0E	0XE2	0X0C	0XEF
0XBC	0X72	0X75	0X6F	0X37	0XA1	0XEC
0XD3	0X8E	0X62	0X8B	0X86	0X10	0XE8
	0X08	0X77	0X11	0XBE	0X92	0X4F
	0X24	0XC5	0X32	0X36	0X9D	0XCF
	0XF3	0XA6	0XBB	0XAC	0X5E	0X6C
	0XA9	0X13	0X57	0X25	0XB5	0XE3
	0XBD	0XA8	0X3A	0X01	0X05	0X59
	0X2A	0X46				

The SAN is a one way 64:64 bit mapping. It consists of four layers of operations which do not commute. The structure of the SAN is as follows:



**Layer 1:**

$$\begin{aligned}
 A &= [A + S\text{-box}(B) + F] \\
 B &= [B + S\text{-box}(C) + G] \\
 C &= [C + S\text{-box}(D) + H] \\
 D &= [D + S\text{-box}(A) + E]
 \end{aligned}$$

**Layer 2:**

$$\begin{aligned}
 E &= [E + S\text{-box}(A)] \\
 F &= [F + S\text{-box}(B)] \\
 G &= [G + S\text{-box}(C)] \\
 H &= [H + S\text{-box}(D)]
 \end{aligned}$$

**Layer 3:**

$$\begin{aligned}
 A' &= [A \oplus S\text{-box}(G)] & B' &= [B \oplus S\text{-box}(H)] \\
 C' &= [C \oplus S\text{-box}(E)] & D' &= [D \oplus S\text{-box}(F)]
 \end{aligned}$$

**Layer 4:**

$$\begin{aligned}
 E' &= [E + A'] \oplus F \\
 F' &= [F + B'] \oplus G \\
 G' &= [G + C'] \oplus H \\
 H' &= [H + D] \oplus E
 \end{aligned}$$

where

S-box(N) is the to the N<sup>th</sup> entry in S-box[ ],  
 + denotes addition Mod 2<sup>8</sup> and  
 ⊕ denotes bitwise XOR

Let us summarize the design of WeakFlex8:

- WeakFlex8 operates on eight 8-bit bytes.
- WeakFlex8 accepts 8 bytes as input and produces 8 bytes as output for every round.
- A round of WeakFlex8 corresponds to 8 rounds of HashCore and 1 round of SAN.

We now construct a stream cipher providing 48 bit security using WeakFlex8.

### 6.2 WEAKSTREAM48

WeakStream48 is a deliberately weakened synchronous stream cipher structurally similar to FastStream192. Keystream generation in WeakStream48 is done by one round of WeakFlex8. The keystream generator takes two inputs, a 48 bit secret key **K** {**K**<sub>0</sub>, **K**<sub>1</sub>, **K**<sub>2</sub>, **K**<sub>3</sub>, **K**<sub>4</sub>, **K**<sub>5</sub>} and a 16 bit nonce **N** {**N**<sub>0</sub>, **N**<sub>1</sub>}, where **K**<sub>0</sub>, **N**<sub>0</sub> are the least significant bytes of **K** and **N** respectively.

It must be noted that FastStream192 and WeakStream48 are identical in structure. The differences between them are as follows:

- WeakStream48 accepts a 48 bit key and a 16 bit nonce whereas FastStream192 accepts a 192 bit key and a 64 bit nonce.
- WeakStream48 produces keystream 8 bytes at a time, whereas FastStream 192 produces keystream 8 words at a time.
- For a given key-nonce pair, WeakStream48 can be used to generate 2<sup>22</sup> bits of keystream, whereas for FastStream192 2<sup>64</sup> bits of keystream can be generated using a given key-nonce pair. A longer keystream is generated so as to allow thorough study of its properties.

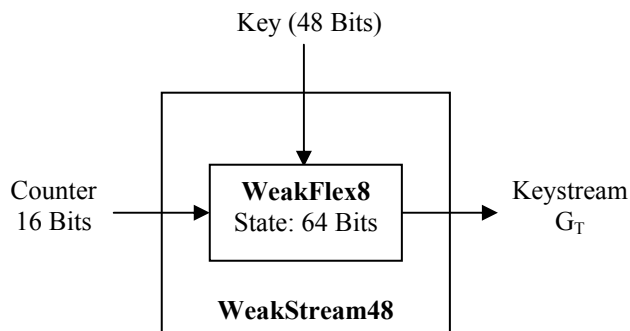


Fig. 10. WeakStream48

**At the Transmitter (Encryption):**

$$C_T = G_T \oplus M_T$$

**At the Receiver (Decryption):**

$$M_T = G_T \oplus C_T$$

where  $M_T$  denotes the message stream,  
 $G_T$  denotes the keystream  
and  $C_T$  denotes the corresponding ciphertext stream.

**6.2.A KEYSTREAM GENERATION IN WEAKSTREAM48**

Keystream is generated in WeakStream48 following the same approach as used in FastStream192. Each round of WeakStream48 uses 1 round of WeakFlex8. The keystream generator takes two inputs, a 48 bit secret key  $K \{K_0, K_1, K_2, K_3, K_4, K_5\}$  and a 16 bit nonce  $N \{N_0, N_1\}$ , where  $K_0, N_0$  are the least significant bytes of  $K$  and  $N$  respectively.

The process of keystream generation is as follows:

- The inner state of the WeakFlex8 is loaded with  $K$  and  $N$  as shown below

$$\left\{ \begin{array}{cccc} N_0 & K_0 & K_1 & K_2 \\ K_3 & K_4 & K_5 & N_1 \end{array} \right\}$$

- WeakFlex8 is allowed to run four times with the output of the present round becoming the inner state of WeakFlex8 for the next round.

$$\left\{ \begin{array}{cccc} N_0 & K_0 & K_1 & K_2 \\ K_3 & K_4 & K_5 & N_1 \end{array} \right\}$$

↓  
4 Rounds of WeakFlex8

$$\left\{ \begin{array}{cccc} A & B & C & D \\ E & F & G & H \end{array} \right\}$$

- The Inner State of WeakFlex8 is loaded with  $\{A,B,C,D,E,F,G,H\}$  as follows:

$$\left\{ \begin{array}{cccc} A & B & C & D \\ E & F & G & H \end{array} \right\}$$

The keystream generator is now ready to produce keystream.

- A single round of WeakFlex8 is then called which produces 8 bytes of keystream.

$$\left\{ \begin{array}{cccc} A & B & C & D \\ E & F & G & H \end{array} \right\}$$

↓  
1 Round of WeakFlex8

$$\left\{ \begin{array}{cccc} G_0 & G_1 & G_2 & G_3 \\ G_4 & G_5 & G_6 & G_7 \end{array} \right\}$$

which can be written as

$$\{G\} = \text{wEAKFlex8}(A,B,C,D,E,F,G,H)$$

- The counter is updated for generating the next 8 bytes of keystream by incrementing GHAB by 1. The value of the counter to be loaded for the next round of encryption thus becomes GHAB + 1.
- On generating every 8 bytes of keystream, the counter is again incremented by 1 before generating the next bytes of the keystream.

**6.2.B PERIOD OF KEYSTREAM FOR WEAKSTREAM48**

**WeakStream48(K,N)** denotes the WeakStream48 stream cipher function acting on the 6 Byte key  $K$  and the 4 Byte sequence  $N, N + 1, \dots, N + (2^{32} - 1)$ .

**WeakStream48 thus has a period of  $2^{32}$ , as the counter rolls over to its initial value after  $2^{32}$  increments, as  $(N+2^{32}) \bmod 2^{32} = N$**

There may be collisions in the keystream and the purpose of WeakStream48 is to study the nature, predictability and frequency of occurrence of such collisions.

**7 COMPARISONS WITH AES**

AES [15] has been the mainstay encryption algorithm for about 3 years now. It has been thoroughly analyzed and even more thoroughly optimized for high performance on a wide variety of platforms, including the AMD Athlon and a Pentium IV. The fastest AES Stream Cipher Implementation in software is reported to take 18 cycles/byte [16]. It however refers to unpublished software which provides no protection against timing attacks. An implementation providing protection against timing leaks should take considerably longer. Our in-house tests suggest encryption at around 24 cycles/byte on an AMD Athlon 2400+ XP Processor Model 8 and at around 23 cycles/byte on an Intel Pentium IV 2.4 GHz. However, it must be noted that the said AES implementation provides only 128-bit key security whereas FastStream192 provides 192-bit key security.

FastStream192 takes 10 cycles/byte for the optimized assembly implementation and 12 cycles/byte for the reference C implementation when encrypting a 32 byte block with a 24 byte key on an AMD Athlon 2400+ XP Processor Model 8 [17]. No key expansion is necessary as in AES. No lookup tables need to be pre-calculated and only one 1KB s-box is used.

On the Pentium IV 2.4 GHz [18], at about 15 cycles/byte for the optimized assembly implementation and 16.8 cycles/byte for the Reference C implementation, FastStream192 takes considerably longer than on an AMD Athlon. This is due to the fact that the Pentium IV is somewhat sluggish when it comes to integer multiplication.

FastHash256 performs considerably faster than AES Hash on all platforms, at around 11.4 cycles/byte for our reference C implementation on an AMD Athlon 2400 + XP Processor Model 8 and at around 16.2 cycles/byte on an Intel Pentium IV. AES Hash takes 30 cycles/byte and 27 cycles/byte respectively on the above platforms [19].

The performance results for constructions using FastFlex for the reference C implementation can be summed up as follows.

Function	Block Size	Key Size	Cycles/Byte A32, A64, PIV
AESHash	128 b		30, -, 27
AES Stream	128 b	128 b	24, -, 23
FastHash256	256 b		11.4, 10.6, 16.2
FastStream 192	256 b	256 b	12, 11.3, 16.8

Table 1. Comparison with AES

QuickStream192, a reduced round variant was also developed for applications requiring very high speed encryption. We can formally define FastFlex as  $(\text{HashCore})^8 \cdot (\text{SAN})$ . That is 8 iterations of the HashCore followed by a single iteration of the SAN. QuickFlex is then  $(\text{HashCore})^4 \cdot (\text{SAN})$ , that is 4 iterations of the HashCore followed by a single iteration of the SAN. It should be noted that as of now both FastFlex and QuickFlex are secure by some margin. Our efforts at cryptanalysis have not been able to break more than 2 iterations of HashCore.

The performance results for QuickStream192 using QuickFlex for the reference C implementation are quite impressive at around 7.3 cycles/byte on an AMD Athlon 2400+ XP Processor Model 8 and at around 10.8 cycles/byte on an Intel Pentium IV 2.4 GHz. In other words, the reference C implementation of QuickStream192 can encrypt a stream at about 3.18 Gbps on an Athlon XP 2400+.

## 8 CONCLUSIONS

A New Fast and Flexible Cryptographic Function, FastFlex has been presented for peer review in this paper.

It has been used to construct a stream cipher and a hash function.

The stream cipher FastStream192 constructed using FastFlex is significantly faster than AES in counter mode by about 12 cycles/byte on an AMD Athlon and 2 cycles/byte on an Intel Pentium IV. The performance degradation on the Pentium IV may be attributed to the fact Integer Multiplication is considerably slower on Pentium IV as compared to AMD Athlon [20].

The keystream generated by FastStream192 has been analysed thoroughly for any known statistical weaknesses using the NIST Statistical Test Suite (NIST PUB 800-022) and no weaknesses have been found in the keystream. Furthermore,

the generated keystream has on average better if not the same statistical randomness as the output of a DRBG using SHA512 (DRBG-SHA512), provided in the NIST test suite itself.

**It is hereby declared that the author has not deliberately inserted any trapdoors in FastFlex, FastStream192, FastHash256 and KeyHash256.**

**The weakened version of FastFlex, WeakFlex8 should only be used for academic purposes. No security claims are made regarding the security of the same.**

**If at all used, WeakFlex8 should only be expected to provide a key security of 48 bits.**

## 9 APPENDICES

### 9.1 DEFINITIONS

**Encryption:** The process of obscuring the information contained in a message in an efficient manner such that only users with the right key to that information can recover the original information. To anyone not having a valid key, the encrypted message should be indistinguishable from a random permutation of bits.

**Key:** A secret password used to encrypt or decrypt messages. A common secret key is shared between the endpoints in symmetric encryption. In an ideal encryption scheme, all security lies only in the key.

**Nonce:** A Number Used Only **Once**, usually paired with a secret key. A nonce need not be random and no security lies in it. It can be declared publicly and is often used as an initialization value for the encryption scheme. The advantage of using a nonce is that in case of loss of synchronization between the endpoints of communication, resynchronization may be achieved by transmitting only a new nonce and not the new key.

**Stream Cipher:** A deterministic algorithm which produces a sequence, usually referred to as keystream, from key and (an optional) nonce, in a manner independent of plaintext. The keystream is later combined with plaintext in a simple, key-independent, and nonce-independent manner to produce ciphertext. A streamcipher should be so that, under the assumption that it is properly used, even if some parts of the plaintext corresponding to (multiple-session) ciphertexts are revealed, no information about unrevealed plaintext corresponding to the rest of the ciphertexts can be deduced.

**Hash Function:** A hash function  $H$  is a transformation that takes a variable-size input  $m$  and returns a fixed-size string, which is called the hash value  $h$  (that is,  $h = H(m)$ ). Hash functions with just this property have a variety of general computational uses, but when employed in cryptography the hash functions are usually chosen to have some additional properties.

**S-box:** A substitution which returns a seemingly unrelated value for any input given to it. S-boxes are used in ciphers to obscure the relationship between plaintext and ciphertext.

**Plaintext:** The original message acted upon by the encryption scheme to produce corresponding ciphertext, from which information contained in the plaintext can only be recovered with the same key as used to encrypt the message.

### 9.2 CRYPTANALYSIS OF FASTFLEX

In this section relevant cryptanalytic techniques are discussed with their possible application to FastFlex (if any).

#### 9.2.A LINEAR CRYPTANALYSIS

Linear cryptanalysis [21] is a known plaintext attack that uses a linear relation or relations between inputs and outputs of the encryption algorithm that hold with a certain probability. This approximation can be used to assign probabilities to the possible keys and locate the most probable one. There however do not seem exist any simple linear relations between the inputs and the outputs of FastFlex.

#### 9.2.B DIFFERENTIAL CRYPTANALYSIS

Differential cryptanalysis [22] involves finding some pairs of plaintext  $\{W, W'\}$  with a small difference such that they produce a small difference in the corresponding output ciphertext  $\{Y, Y'\}$ . FastFlex as used in FastStream192 accepts a 2 word input (Only two words of the counter change) and produces 8 words as output. It seems unreasonable to find any input pair  $\{W, W'\}$  such that the corresponding output  $\{Y, Y'\}$  also has a small difference. This is because while there are only  $(2^{64})$  choices for the input, there are  $(2^{256})$  choices for the output differences. The probability of finding a low difference in both input and corresponding output is equivalent to finding a low difference input pair in  $2^{64}$  which has a corresponding low difference in  $2^{256}$ . Now if we assume that FastHash is a one to one mapping (or a near one to one mapping with a small number of collisions), the  $2^{64}$  inputs will map to only the corresponding  $2^{64}$  outputs spread across  $2^{256}$  possible outputs of FastFlex. The possibility of a low difference being carried over from an input pair to the corresponding output pair is thus very low.

#### 9.2.C SLIDE ATTACKS

Slide Attacks [23] compare the present state of the cipher to the next state. Fortunately, the two states of the  $2*4$  matrix (denoted  $M$ ) in the HashCore do not have the property that  $Col(Row(M)) = Col(Col(M^T)) = Row(Row(M^T))$

Slide attacks are thus not possible as there is no simple relationship between the present and the next state of FastFlex.

### 9.3 STATISTICAL TESTING OF FASTSTREAM92 KEYSTREAM

Ciphertext is generated in FastStream92 by XORing plaintext with corresponding keystream bits. The generated keystream should thus satisfy known tests for randomness for it to withstand statistical analyses by adversaries, to whom the ciphertext should appear indistinguishable from a random sequence.

The keystream generated using FastStream192 was subjected to thorough tests using the NIST Statistical Test Suite [24], which was used by the NIST during the development and testing phase of AES [25], to test the keystreams generated by the block cipher proposals for statistical weaknesses. FastStream192 was compared with DRBG-SHA512 [26, 27, 28], an excellent random sequence generator. A detailed description of the technical background of the tests and the conventions used can be found in [24].

#### 9.3.1 MONOBIT TEST

The purpose of this test is to determine whether the number of ones and zeros in a sequence are approximately the same as would be expected for a truly random sequence.

(Only Reports Number of Streams Failed on average in 1000 independent tests with random keys)  
b: Bits, B: Bytes

StreamLength / No. of Streams	DRBG SHA512	FastStream192
256b/1000	10, 1%	9, 0.9%
512b/1000	6, 0.6%	7, 0.7%
4KB/1000	15, 1.5%	9, 0.9%
8KB/1000	10, 1%	5, 0.5%
1MB/100	2, 2%	0, 0%
1GB/1	-	0, 0%

Table 2. Monobit Test

#### 9.3.2 RUNS TEST

The purpose of the runs test is to determine whether the number of runs of ones and zeros of various lengths is as expected for a random sequence. In particular, this test determines whether the oscillation between such substrings is too fast or too slow and is the observed oscillation what would be expected in a truly random sequence.

(Only Reports Number of Streams Failed)  
b: Bits, B: Bytes

StreamLength / No. of Streams	DRBG SHA512	FastStream192
256b/1000	12, 1.2%	10, 1%
512b/1000	7, 0.7%	6, 0.6%
4KB/1000	9, 0.9%	10, 1%
8KB/1000	17, 1.7%	12, 1.2%
1MB/100	2, 2%	1, 1%

Table 3. Runs Test

#### 9.3.3 LONGEST RUN OF ONES TEST

The purpose of this test is to determine whether the longest run of ones within the tested sequence is consistent with the longest run of ones that would be expected in a random sequence.

(Only Reports Number of Streams Failed)  
b: Bits, B: Bytes

StreamLength / No. of Streams	DRBG SHA512	FastStream192
1MB/100	1, 1%	1, 1%

Table 4. Longest Run of Ones Test

#### 9.3.4 BINARY MATRIX RANK TEST

The purpose of this test is to check for linear dependence among fixed length substrings of the original sequence.

(Only Reports Number of Streams Failed)  
b: Bits, B: Bytes

StreamLength / No. of Streams	DRBG SHA512	FastStream192
8KB/1000	10, 1%	9, 0.9%
1MB/100	1, 1%	1, 1%

Table 5. Binary Matrix Rank Test

#### 9.3.5 DISCRETE FOURIER TRANSFORM TEST

The focus of this test is the peak heights in the Discrete Fourier Transform of the sequence. The purpose of this test is to detect periodic features (i.e., repetitive patterns that are near each other) in the tested sequence that would indicate a deviation from the assumption of randomness. The intention is to detect whether the number of peaks exceeding the 95 % threshold is significantly different than 5 %.

(Only Reports Number of Streams Failed)  
 b: Bits, B: Bytes

StreamLength / No. of Streams	DRBG SHA512	FastStream192
4KB/1000	14, 1.4%	6, 0.6%
8KB/1000	8, 0.8%	7, 0.7%
1MB/100	2, 2%	1, 1%

Table 6. Discrete Fourier Transform Test

9.3.6 CUMULATIVE SUMS TEST

The purpose of the cumulative sums test is to determine whether the sum of the partial sequences occurring in the tested sequence is too large or too small.

(Only Reports Number of Streams Failed)  
 b: Bits, B: Bytes

StreamLength / No. of Streams	DRBG SHA512	FastStream192
256b/1000	9, 0.9%	13, 1.3%
512b/1000	11, 1.1%	16, 1.6%
4KB/1000	29, 2.9%	13, 1.3%
8KB/1000	23, 2.3%	14, 1.4%
1MB/100	3, 3%	0, 0%

Table 7. Cumulative Sums Test

9.3.7 RANDOM EXCURSIONS TEST

The purpose of this test is to determine if the number of visits to a state within a random walk exceeds what one would expect for a random sequence.

(Only Reports Number of Streams Failed)  
 b: Bits, B: Bytes

StreamLength / No. of Streams	DRBG SHA512	FastStream192
1MB/100	5, 5%	2, 2%

Table 8. Random Excursions Test

9.3.8 RANDOM EXCURSIONS VARIANT TEST

The purpose of this test is to detect deviations from the distribution of the number of visits of a random walk to a certain state. Due to the computationally expensive nature of this test, it could only be conducted on a limited number of keystream samples.

Only Reports Number of Streams Failed)  
 b: Bits, B: Bytes

StreamLength / No. of Streams	DRBG SHA512	FastStream192
1MB/100	6, 6%	5, 5%

Table 9. Random Excursions Variant Test

9.3.9 NON OVERLAPPING TEMPLATE MATCHING TEST

The purpose of this test is to reject sequences that exhibit too many occurrences of a given non-periodic (aperiodic) pattern. The results of this test can be downloaded from [fastflex.sourceforge.net](http://fastflex.sourceforge.net)

9.3.10 OVERLAPPING TEMPLATE MATCHING TEST

The purpose of this test is to reject sequences that show deviations from the expected number of runs of ones of a given length. The results of this test can be downloaded from [fastflex.sourceforge.net](http://fastflex.sourceforge.net)

9.3.11 LINEAR COMPLEXITY TEST

The focus of this test is the length of a linear feedback shift register (LFSR). The purpose of this test is to determine whether or not the sequence is complex enough to be considered random. Random sequences are characterized by longer LFSRs. An LFSR that is too short implies non-randomness.

(Only Reports P Value for M = 512)  
 b: Bits, B: Bytes

StreamLength / No. of Streams	DRBG SHA512	FastStream192
128KB/1000	0.41205	0.48388
1MB/100	0.50671	0.56059
8MB/50	0.82407	0.82721

Table 10. Linear Complexity Test

9.3.12 SERIAL TEST

The focus of this test is the frequency of all possible overlapping  $m$ -bit patterns across the entire sequence. The purpose of this test is to determine whether the number of occurrences of the  $2m$ ,  $m$ -bit overlapping patterns is approximately the same as would be expected for a random sequence. Random sequences have uniformity; that is, every  $m$ -bit pattern has the same chance of appearing as every other  $m$ -bit pattern.

(Only Reports Number of Streams Failed)

b: Bits, B: Bytes

StreamLength / No. of Streams	DRBG SHA512	FastStream192
256b/1000	13, 1.3%	14, 1.4%
512b/1000	14, 1.4%	12, 1.2%
4KB/1000	26, 2.6%	20, 2%
8KB/1000	23, 2.3%	16, 1.6%
1MB/100	1, 1%	1, 1%

Table 11. Serial Test

### 9.3.13 APPROXIMATE ENTROPY TEST

The purpose of this test is to determine whether the number of occurrences of  $m$ -bit overlapping patterns is approximately the same as would be expected for a random sequence.

(Only Reports Number of Streams Failed)

b: Bits, B: Bytes, M = 8, 8, 16

StreamLength / No. of Streams	DRBG SHA512	FastStream192
4KB/1000	16, 1.6%	12, 1.5%
8KB/1000	13, 1.3%	18, 1.1%
1MB/100	6, 6%	3, 3%

Table 12. Approximate Entropy Test

### 9.3.14 MAURER'S "UNIVERSAL STATISTICAL" TEST

The focus of this test is the number of bits between matching patterns (a measure that is related to the length of a compressed sequence). The purpose of the test is to detect whether or not the sequence can be significantly compressed without loss of information. A significantly compressible sequence is considered to be non-random.

(Only Reports Number of Streams Failed)

b: Bits, B: Bytes

StreamLength / No. of Streams, L	DRBG SHA512	FastStream192
64KB/1000, 6	15, 1.5%	19, 1.9%
128KB/1000, 7	23, 2.3%	9, 0.9%
256KB/1000, 8	16, 1.6%	16, 1.6%
1MB/100	2, 2%	0, 0%

Table 13. Universal Statistical Test

### 9.3.15 .7Z Compression Test

7z is a new archive format, providing a high compression ratio. A truly random sequence should not have any

redundancy and thus should not be compressible. The purpose of this test was to compress large volumes of keystream generated using FastStream192 and see if any redundancy was found.

Dictionary Size: 64 MB

Method: LZMA

b: Bits, B: Bytes

StreamLength / No. of Streams	Compressed Size
64KB/1000	64.1 MB
128KB/1000	128.1 MB
1MB/1000	1000.1 MB
1GB/10	10.0001 GB

Table 14. .7Z Test

### 9.3 FASTFLEX EXAMPLE

HashCore

**Round 1:**

Initial State:

$$\left\{ \begin{array}{cccc} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{array} \right\}$$

$$\text{Row}(1,0,0,0) = \{10208811, 80, 102000, 41000020\}$$

$$\text{Row}(0,0,1,0) = \{800100, 0, 1, 20000\}$$

After Row transformations:

$$\left\{ \begin{array}{l} 10208811, 80, 102000, 41000020 \\ 1, 20000, 800100, 0 \end{array} \right\}$$

$$\text{Col}(10208811, 1) = \{10208811, 8b9a7465\}$$

$$\text{Col}(80, 20000) = \{1000000, 8b9a7465\}$$

$$\text{Col}(102000, 800100) = \{10200000, 8b9a7c75\}$$

$$\text{Col}(41000020, 0) = (0, 8b9a7465)$$

**After Round 1:**

$$\left\{ \begin{array}{l} 10208811, 100000, 10200000, 0 \\ 8b9a7465, 8b9a7465, 8b9a7c75, 8b9a7465 \end{array} \right\}$$



## REFERENCES

- [1] Dr. Daniel J Bernstein  
The Salsa 20 Specification,  
[cr.yp.to](http://cr.yp.to), [ecrypt.eu.org/stream](http://ecrypt.eu.org/stream)
- [2] Hawkes, Paddon, Rose, de Vries  
Primitive Specification for NLS  
[ecrypt.eu.org/stream](http://ecrypt.eu.org/stream)
- [3] E-Stream Project,  
Call for Stream Cipher Proposals,  
[ecrypt.eu.org/stream](http://ecrypt.eu.org/stream)
- [4] D.J. Wheeler, R.J. Needham  
The Tiny Encryption Algorithm,  
FSE '94
- [5] The Xilinx Spartan III E  
Tech Report,  
[xilinx.com](http://xilinx.com)
- [6] UCL Crypto Group  
Cryptanalysis of Block Ciphers : A Survey  
[dice.ucl.ac.be/crypto/](http://dice.ucl.ac.be/crypto/)
- [7] Dr. Daniel J Bernstein  
Cache Timing Attacks on AES,  
[cr.yp.to](http://cr.yp.to)
- [8] Kahn on Codes  
Macmillan Publishing Co, 1967
- [9] Bruce Schneier  
Applied Cryptography,  
John Wiley & Sons, 1996
- [10] R.C. Merkle  
A Fast Software One-Way Hash Function,  
Crypto '90
- [11] R.C. Merkle  
A Digital Signature Based on a Conventional Encryption  
Function,  
Crypto '87
- [12] MD5 Message Digest  
RFC 1321
- [13] B.S. Kaliski  
A Survey of Encryption Standards,  
IEEE Micro vol. 13, Dec. 1993
- [14] NIST  
Skipjack and KEA Algorithm Specifications,  
[csrc.nist.gov/encryption/skipjack-kea.htm](http://csrc.nist.gov/encryption/skipjack-kea.htm)
- [15] NIST  
Announcing Development of a Federal Information Standard for  
Advanced Encryption Standard,  
Federal Register, v. 62, n. 1, 2 Jan 1997
- [16] AES Speed  
E-Stream Performance Benchmarks  
[www.ecrypt.eu.org/stream/perf/#results](http://www.ecrypt.eu.org/stream/perf/#results)
- [17] AMD Processors x86 Code Optimization Guide  
Vol. 1, Vol. 2 & Vol. 3,  
[amd.com](http://amd.com)
- [18] Intel Architecture, Software Development Manual  
Vol. 2,  
[intel.com](http://intel.com)
- [19] E-Stream Performance Benchmarks  
[www.ecrypt.eu.org/stream/perf/#results](http://www.ecrypt.eu.org/stream/perf/#results)
- [20] The New AMD A64 Processor  
[cpuid.com](http://cpuid.com)
- [21] Linear Cryptanalysis Method for DES Cipher  
M. Matsui,  
Eurocrypt '93
- [22] Differential Cryptanalysis of DES-Like Cryptosystems  
E. Biham, A. Shamir,  
Crypto '90
- [23] Alex Biryukov, David Wagner  
Slide Attacks,  
FSE '99
- [24] NIST  
PUB800-22, NIST Statistical Test Suite for Testing of Random  
& Pseudo Random Number Generators
- [25] NIST  
Federal Information Standard, Advanced Encryption Standard,  
[nist.gov](http://nist.gov)
- [26, 27, 28] [FIPS140-1, FIPS 180-1, FIPS 186] Federal  
Information Processing Standards

