The Journal of Systems and Software 83 (2010) 543–547

Contents lists available at ScienceDirect

The Journal of Systems and Software

journal homepage: www.elsevier.com/locate/jss

A practical distinguisher for the Shannon cipher ☆

Zahra Ahmadian [a,*], Javad Mohajeri [b], Mahmoud Salmasizadeh [b], Risto M. Hakala [c], Kaisa Nyberg [c]

[a] *Department of Electrical Engineering, Sharif University of Technology, Tehran, Iran*
[b] *Electronics Research Center, Sharif University of Technology, Tehran, Iran*
[c] *Department of Information and Computer Science, Helsinki University of Technology, Finland*

ARTICLE INFO

Article history:
Received 29 November 2008
Received in revised form 29 August 2009
Accepted 13 October 2009
Available online 28 October 2009

Keywords:
Stream ciphers
Shannon
Distinguishing attack
Crossword Puzzle attack
Linear cryptanalysis

ABSTRACT

In this paper, we present a practical linear distinguisher on the Shannon stream cipher. Shannon is a synchronous stream cipher that uses at most 256-bit secret key. In the specification for Shannon, designers state that the intention of the design is to make sure that there are no distinguishing attacks on Shannon requiring less than 2^{80} keystream words and less than 2^{128} computations. In this work we use the Crossword Puzzle attack technique to construct a distinguisher which requires a keystream of length about 2^{31} words with workload about 2^{31}.

1. Introduction

Stream ciphers form an important class of symmetric encryption primitives. Since the standardization of the block cipher AES (NIST, 2001) a lot of attention in symmetric cryptology has moved to the area of stream ciphers. Stream ciphers can offer advantages of different kinds, for example, low power consumption, low hardware complexity or extreme software efficiency.

An ideal keystream is a truly random sequence. Such a keystream provides perfect secrecy. It means that having an infinite computing power, an adversary is unable to obtain any information of the plaintext (Shannon, 1949). Thus one of the most important criteria for the security of a keystream generator is that the keystream can be made to appear random. To analyze this property, statistical and algebraic distinguishing attacks have been developed. Distinguishing attacks does not result in any information leakage about the key or internal state of the cryptographic primitive. In such attacks, an adversary tries to determine whether a sequence has been produced by a specific cipher or seems to be a random sequence. In fact the data and computational complexity that is required for a successful distinguishing attack on a stream cipher, gives a criteria for the strength of that cipher.

Linear distinguishing attacks have been presented to numbers of stream ciphers such as SNOW 1.0 (Coppersmith et al., 2002), SNOW 2.0 (Watanabe et al., 2004; Nyberg and Wallen, 2006), SOBER-128 (Cho and Pieprzyk, 2006a) and NLS (Cho and Pieprzyk, 2006b).

The Shannon stream cipher was designed by Rose et al. (2007) of Qualcomm Australia as a new member of SOBER family of stream ciphers. It has been designed according to the ECRYPT NoE call for stream cipher primitives, profile 1A (2005) (but well after the call).

Shannon uses a secret key that may be up to 256 bits in length. In addition to keystream generation, Shannon offers message authentication functionality. In this paper, we consider only the keystream generation part, and propose a distinguisher for this mode of operation of Shannon. The primitive is constructed from a nonlinear feedback shift register (NFSR) and a nonlinear output filter (NLF).

The first attack on the Shannon was a multidimensional linear distinguishing attack by Hakala and Nyberg (2008, 2009) which requires a keystream of length 2^{107} words. The second one is a differential distinguishing attack based on the fault analysis (Hassanzadeh et al., 2008). This attack has an extremely low data complexity, 11 keystream words for two differential pairs of the initial state, and its computational complexity is equivalent to the complexity of four times running the Shannon. While being a successful differential analysis of Shannon, the authors had an additional assumption: access to the initial state, which renders the attack unrealistic. A distinguisher can be assumed to have access only to the keystream sequence but not to the internal state of the algorithm.

In our attack, we exploit the Crossword Puzzle attack technique to construct an efficient distinguisher for the Shannon cipher. The

☆ This project is partially supported by Iran Telecommunication Research Center grant number T/500/10651.

* Corresponding author. Tel.: +98 21 66164904.
E-mail addresses: ahmadian@ee.sharif.edu (Z. Ahmadian), mohajer@sharif.edu (J. Mohajeri), salmasi@sharif.edu (M. Salmasizadeh), risto.m.hakala@tkk.fi (R.M. Hakala), kaisa.nyberg@tkk.fi (K. Nyberg).

0164-1212/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.jss.2009.10.019

Crossword Puzzle attack proposed by Cho and Pieprzyk (2006b) in a distinguishing attack on NLS is a method for finding linear relations of the output keystream bits for NLF and NFSR based primitives. This method is an extension of the linear masking method that was introduced by Coppersmith et al. (2002) for a distinguishing attack on traditional LFSR based stream ciphers such as SNOW 1.0 (Coppersmith et al., 2002) and SOBER-128 (Cho and Pieprzyk, 2006a).

In a Crossword Puzzle attack, one first derives linear approximations of both NFSR and NLF, the next step has been to combine proper sets of linear approximation of NFSR and NLF to eliminate all the internal states and achieve a new approximation of only keystream bits with reasonable bias. The first attack on Shannon by Hakala and Nyberg also used the Crossword Puzzle technique, but their attack was far from being optimal. In this paper, we introduce a new Crossword Puzzle distinguisher for the Shannon keystream generator with an estimated bias of $2^{-18.13}$. This means that the distinguishing attack is expected to succeed given about $2^{31.3}$ words of the keystream. We implemented the distinguisher in practice. The observed bias was $2^{-19.12}$ and the distinguisher succeeded with high probability after about $2^{33.9}$ keystream words was available. This should be compared with the specification for Shannon (Rose et al., 2007) which stated that the intention of the design is to make sure that there are no distinguishing attacks on Shannon requiring less than 2^{80} keystream words and less than 2^{128} computations.

The paper is organized as follows: Section 2 briefly introduces to the statistical principles of distinguishers. In Section 3, we give a description of the Shannon stream cipher without its MAC functionality. In Section 4, we present the linear distinguishing attack on the Shannon. Section 5 includes the results of practical experiments on the full version of the cipher. Some properties of the Shannon stream cipher is presented in Section 6. Finally, Section 7 concludes our work, including a comparison between our attack and previous attacks.

2. Linear distinguisher

A distinguishing attack on a stream cipher is a statistical hypothesis test, which for a given binary sequence $\{v_t\}$ of length N decides between the following two hypotheses with high confidence level:

- Cipher, the sequence is actually produced from the given cipher.
- Random, the sequence is uniformly random.

A linear distinguisher at first applies a linear transformation to the input sequence $\{v_t\}$ to get a new sequence $\{\hat{v}_t\}$ in such a way that $\hat{v}_t$ is a biased binary random variable, i.e., $\varepsilon = \Pr(\hat{v}_t = 0) - 1/2 \neq 0$ if the sequence originates from the cipher, and $\hat{v}_t$ is unbiased, i.e., $\varepsilon = \Pr(\hat{v}_t = 0) - 1/2 = 0$ otherwise. The distinguishing between two hypotheses is based on the distance $|\varepsilon|$ between the expected values of the biased distribution and the uniform distribution.

Assume now that $\varepsilon > 0$, which will be the case in this paper. After sampling N bits of the sequence, a linear distinguisher decides between two hypothesis as follows. Let u be the number of bits $\hat{v}_t = 0$. Then the output of the distinguisher is cipher, if $u \geqslant N \cdot \left(\frac{1}{2} + \frac{\varepsilon}{2}\right)$, otherwise the output is Random (Junod, 2003).

A linear distinguisher can perform this test reliably, if it has sufficient number of samples $\{\hat{v}_t\}$. It was shown by Matsui (1994) that the number of required samples for a linear attack to make a reliable decision is $N \approx 1/\varepsilon^2$. Later Coppersmith et al. (2002) showed that the same is true for a linear distinguishing attack. More precisely, the probability that a sequence originating from the cipher is correctly identified is $P_S = 1 - \Phi\left(-\varepsilon\sqrt{N}\right)$, where Φ is the cumulative distribution function of the standard normal distribution. For $N = 1/\varepsilon^2$, we have $P_S = 0.921$.

3. Brief description of Shannon

Shannon is constructed from a nonlinear feedback shift register and a nonlinear output filter. It is based on 32-bit operations, and 32-bit words. We denote the state of the register at time t, by $\sigma_t = (r_t[0], r_t[1], \ldots, r_t[15])$.

The structure of the Shannon keystream generator is depicted in Fig. 1. In this mode of operation the transition from the state σ_t $(t \geqslant 0)$ to the state σ_{t+1}, and the generation of the output keystream word, v_t, is defined as follows:

$$r_{t+1}[i] = r_t[i+1] \quad \text{for } 1 \leqslant i \leqslant 14 \tag{1}$$

$$r_{t+1}[15] = f_1(r_t[12] \oplus r_t[13] \oplus Konst) \oplus (r_t[0] \lll 1) \tag{2}$$

$$r_{t+1}[0] = r_t[1] \oplus f_2(r_{t+1}[2] \oplus r_{t+1}[15]) \tag{3}$$

$$v_t = f_2(r_{t+1}[2] \oplus r_{t+1}[15]) \oplus r_{t+1}[8] \oplus r_{t+1}[12] \tag{4}$$

where f_1 and f_2 are nonlinear functions and $Konst$ is a 32-bit secret constant that is derived in the initialization process. The notation $x \lll n$ is used to denote n-bit left rotation of the 32-bit word x. Note that Eqs. (2) and (3) specify the NFSR and Eq. (4) specifies the NLF part of the primitive.

The functions f_1 and f_2 are defined by

$$\begin{aligned} f_1(x) &= g(g(x,5,7),19,22) \\ f_2(x) &= g(g(x,7,22),5,19) \end{aligned} \tag{5}$$

where the function g is defined by

$$g(x,a,b) = x \oplus ((x \lll a) \vee (x \lll b)) \tag{6}$$

and $x \vee y$ means bitwise OR of 32-bit words x and y.

For further details, such as initialization procedure and MAC functionality, we refer the reader to the specification of Shannon (Rose et al., 2007).

4. Linear distinguishing attack on Shannon

Distinguishing attacks are based on a statistical model of the cipher, so a set of standard assumptions are always accepted:

1. Distinguishing attack is a known plaintext attack and the adversary only has a sufficient length of ciphertext sequence and corresponding plaintext.
2. The contents of the shift register are statistically independent and uniformly distributed 32-bit random variables.
3. Linear approximations of the nonlinear functions of the cipher are statistically independent if the involved variables are statistically independent.

Assumption 2 is in a crucial role when establishing statistical independence of linear approximations according to Assumption

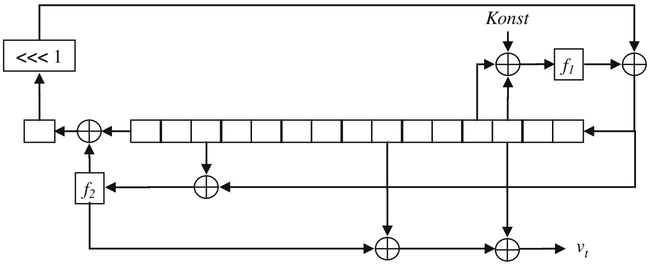

Fig. 1. The Shannon keystream generator.