On the Applicability of Distinguishing Attacks Against Stream Ciphers

Greg Rose, Philip Hawkes

QUALCOMM Australia
{ggr, phawkes}@qualcomm.com

Abstract. We demonstrate that the existence of distinguishing attacks against stream ciphers is unrelated to their security in practical use, and in particular that the amount of data required to perform a distinguishing attack is unrelated to the key length of the cipher. The implication for the NESSIE Project is that no submitted symmetric cipher would be accepted under the unpublished rules for distinguishing attacks, not even the block ciphers in Counter Mode or Output Feedback Mode.

Keywords. Distinguishing attack, stream cipher**.**

1 Introduction

NESSIE is a project within the Information Societies Technology (IST) Programme of the European Commission.

Quoting from https://www.cosic.esat.kuleuven.ac.be/nessie/:

> "The main objective of the project is to put forward a portfolio of strong cryptographic primitives that has been obtained after an open call and been evaluated using a transparent and open process."

Of the stream ciphers submitted, "All have problems to a greater or lesser degree." (Bart Preneel, at the EuroCrypt Rump Session, 2002). Some of these problems are distinguishing attacks with a computational complexity less than is required for an enumeration attack on the key.

We argue below that

1. Distinguishing attacks are properly related to the amount of data available, not to the key length,
2. In most cases, distinguishing attacks on stream ciphers have no security implications in the context of use of the cipher,
3. In practice, distinguishing attacks on stream ciphers are impossible to mount.

In section 2, we define distinguishing attacks and refer to some examples. Section 3 considers weak attacks against block ciphers, while Section 4 looks in more detail at

some attacks against true stream ciphers. Section 5 examines the context in which stream ciphers are used.

2 Distinguishing Attacks

A distinguishing attack on a cipher relates to the formal model of security, where an adversary can distinguish between the output of a particular cipher and the output of a truly random process, with a non-negligible probability. If attackers cannot make this distinction, an algorithmically derived stream cipher will look to them like a Vernam Cipher (often called a One Time Pad), and will be information-theoretically secure. Of course it is true that there is always a distinguishing attack against any algorithmic cipher; since it must have a finite key, and so brute-force key enumeration will yield a distinguishing attack of complexity 2^{k-1} where k is the key length.

In the case of block ciphers, being able to identify some distinguishing characteristic of the output might lead to an attack that reveals information about the key of the cipher. For example, Differential Cryptanalysis [2] of DES proceeds by identifying a distinguishing characteristic of the first rounds of DES, and then uses that characteristic to verify guesses about the key bits.

While it is conceivable that distinguishing attacks applied to stream ciphers might yield information about the key, or about future keystream, many distinguishing attacks on stream ciphers do not, by themselves, compromise the cipher in its normal operation. For example, RC4 is vulnerable to a distinguishing attack of order 2^{31} bytes [1], while DES in OFB or CTR mode has straightforward distinguishing attacks of order 2^{32} blocks (2^{35} bytes). NIST recommends both OFB and CTR modes for use with AES with 256-bit keys, despite the fact that the same straightforward distinguishing attacks have complexity 2^{64} blocks (2^{68} bytes). And yet, nowhere in the published literature are these attacks referred to as weaknesses in those ciphers; instead, they are used as justification for the "rule of thumb" that one should not generate too much keystream before rekeying.

In this paper we would like to draw an informal distinction between those attacks that yield useful cryptanalytic information such as information about the key or unknown keystream, and attacks that do not yield such useful information. We will call these *powerful* and *weak* distinguishing attacks respectively.

In the discussion below, it will be important to remember that the attacks we refer to require large amounts of known plaintext and the corresponding ciphertext. We will examine some cases where there might be some uncertainty in the known plaintext.

3 Distinguishing Attacks Against Block Ciphers

The existence of distinguishing attacks against block ciphers, particularly in an iterated design where the distinguisher applies to all but one round, might be usable to derive part of the last round scheduled key, and expose the cipher to a divide and conquer attack. Again, referring back to DES in CTR mode as an example, distinguishing

it from random after 2^{32} blocks of output does not help to recover its key, nor to predict with any useful degree of accuracy subsequent outputs. That DES in OFB mode would be considered broken at that point is an artifact of its being a permutation, not of the distinguishing attack *per se*.

To further illustrate this point, we would like to examine the ramifications of the weak distinguishing attack against a block cipher used in Counter Mode. Let F be a permutation of b-bit blocks, drawn at random from the set of all such permutations.

$$F: \{0,1\}^b \rightarrow \{0,1\}^b \text{ s.t. } F \text{ is bijective} \quad (1)$$

Now, there are $2^b!$ such permutations, so it is arguable that a key used to specify any particular one would be $\lceil \log_2(2^b!) \rceil$ bits long. By Sterling's Approximation, such a key would be approximately $b.2^b$ bits long. This is very long indeed. In fact, it is approximately the same size as a codebook for the cipher. The work to build this codebook is very much less than the work to enumerate the "keys".

The simple distinguishing attack against F used in Counter Mode is based on the fact that F is a permutation of the possible outputs, and not a truly random function, so the birthday paradox provides a distinguisher. The expected number of collisions (that is, identical b-bit blocks) in a stream of m random output blocks is:

$$\frac{m(m-1)}{2^{b+1}} \quad (2)$$

The attack is to observe $2^{b/2}$ output blocks, and with high probability, there will be a duplicate block if the output stream comes from a random function. If no duplicate block is present, the output stream is much more likely to be from F.

But now consider what this distinguishing attack has determined about the "key", or about upcoming keystream. A very tiny proportion of the possible "keys" have been eliminated. The next block of keystream to be output from F is slightly constrained, in that the probability of it being the same as one of the observed output blocks is zero, but $Pr[b_{out} = x] = 1/(2^b - 2^{b/2})$ for all blocks x which have not yet been observed. The entropy $H(b_{out})$ of the next output block was b bits without taking into account the blocks seen. But the entropy of the next output block given all the blocks already seen is:

$$H(b_{out} \mid \{b_i, 0 \le i \le 2^{b/2}\}) = b.\log_2(1 - 2^{-b/2}) \quad (3)$$

Thus the loss in entropy from knowing enough outputs to reliably distinguish the counter mode output from random is a negligible fraction of a bit.

Shannon in his seminal paper [3] introduced the concept of *equivocation*, which is the amount of uncertainty about a message given knowledge of the corresponding received message after transmission over a noisy channel. While the situation here is not really equivalent, we asked the question "How much would have to be wrong with the keystream, before this distinguishing attack will fail?"

We performed a relatively simple experiment. Using a block size b of 32 bits, we examined 100 streams of output, each 2^{18} blocks long, from a stream cipher that passes all statistical tests used so far [4]. 778 collisions were detected, compared to the expected value of 800 given by Equation 2. Using a 32-bit block cipher [5] there

were of course no collisions detected within any stream. However, when exactly one bit from each block of output from the block cipher was flipped at (pseudo[1]-)random, a total of 822 collisions were counted from 100 similar output streams. Comparing the distributions of the collisions with the Kolmogorov-Smirnov test gave the result that the two data sets were from the same statistical distribution with 99.4% probability. It would be interesting to extend these results (our computers began page-thrashing), but the clear result is that the distinguishing attack is very sensitive to the accuracy of the keystream, and hence relies heavily on the known plaintext assumption.

4 Distinguishing Attacks Against Stream Ciphers

In the context of the stream cipher attack on RC4 mentioned above, and of published distinguishing attacks on other ciphers such as SNOW and SOBER [6,7], these distinguishing attacks do not yield any usable information about the state of the cipher generator, and cannot be used to attack the generator itself. RC4 is still considered perfectly adequate for encryption in SSL and TLS, with 128-bit keys, despite the fairly powerful distinguishing attack mentioned above. For another example (the following discussion assumes the attack on SNOW [6]), given 2^{95} words of known keystream output from SNOW, and a large amount of ciphertext encrypted with the subsequent output from the same generated keystream, there is still no known or hypothesized attack that would reveal any useful information at all about the corresponding unknown plaintext. Neither the content of the ciphertext, state of the generator, nor the input key, is compromised by these attacks.

One way of looking at this is that 2^{95} words of known plaintext was required to recover one bit of information about the ciphertext.

There is no clear relationship between the key length and the amount of generated keystream for which such a weak distinguishing attack would justify calling a stream cipher "broken". This decision depends more on the amount of data to be encrypted than upon the key length. Conversely, the strong distinguishing attack corresponding to key enumeration requires no known plaintext at all, and only enough ciphertext to exceed the unicity distance of the input language and key size. We believe the phrase "distinguishing attack [...] much faster than exhaustive key search" to be like comparing apples to stream ciphers.

Context of Stream Ciphers

Some cryptographers have posed the question “Why do we need stream ciphers, when we can use block ciphers in Counter Mode?” The primary answer to this question is that synchronous stream ciphers can presumably be made more efficient. Rijndael was selected for the Advanced Encryption Standard algorithm primarily for reasons of

[1] The least significant 5 bits of the previous output block were used to choose a bit to invert in the current block. There should be no observable correlation between the encryptions of distinct counter values.

efficiency with respect to the other candidates, but there are a number of stream ciphers (which we argue are comparably secure to a block cipher in Counter Mode) that are faster in software and smaller in hardware than Rijndael.

Stream ciphers are usually used, then, in applications where large amounts of data are employed, or extremely high throughput is needed, or low complexity hardware is a requirement. Most cutting-edge applications with these requirements are in multimedia applications, for example mobile phones, music and video. While the raw input may be highly redundant, in all of these cases the data is highly compressed before transmission, giving data with high entropy rate per bit. The recent release (to movie theatres only) of the Star Wars movie *Attack of the Clones* [9] consisted of nearly one terabyte of compressed video and audio data, or for simplicity call it 2^{40} bytes. The known plaintext required for the distinguishing attack on SNOW is 2^{55} such movies, certainly more than will be produced in the next century. And with the value of that one movie being around US$\$2^{24}$, what's the chance that there will be that much known plaintext? As an aside, we will mention that the three component video signals and the audio signal are encrypted using triple-DES in Interleaved Output Feedback Mode, requiring 10 DES cores in hardware (the audio being low enough bandwidth that a single core can cope with it), and new keys are required for the streams for every few minutes of movie. The hypothetical 2^{55} movies mentioned above will probably be encrypted with at least some different keys, again making the distinguishing attack inapplicable.

Data to be encrypted has some inherent entropy, for without entropy it is without value. The entropy in the plaintext, even if only a small proportion of the bits transmitted, will generally frustrate these kinds of distinguishing attack. This is especially true when considering the kind of data for which high-speed stream ciphers are most desirable, such as highly compressed streaming video. We are concerned that while we cryptographers are rejecting stream ciphers based on distinguishing attacks, other communities are considering pseudo-randomly changing the sign bit of elements of discrete cosine transforms [8] to be sufficient encryption! In that document, some pictures are recognizable despite the "encryption" – that's a real distinguishing attack. A scheme based on this one is likely to be standardized for MPEG "encryption".

Conclusion

In summary, we feel that weak distinguishing attacks with large known plaintext requirements do not represent a security problem in practice.

There seems to be a need for a better understanding of the relationship between data complexity of cryptanalytic algorithms and their true strength. This seems to hinge on the question of how much useful information results from the attack.

References

1. Scott R Fluhrer and David A McGrew, "*Statistical Analysis of the Alleged RC4 Keystream Generator*", Fast Software Encryption Seventh International Workshop, Springer, 2000.
2. E. Biham, A. Shamir, *Differential Cryptanalysis of DES-like Cryptosystems*. Journal of Cryptology, Vol. 4 No. 1 1991.
3. C. E. Shannon, *A Mathematical Theory of Communication*, Bell System Technical Journal, Volume 27, July-October 1948.
4. G. Rose, P. Hawkes, *Turing: a fast software stream cipher*, Rump session of Crypto 2002, http://people.qualcomm.com/ggr/QC/Turing.tgz .
5. G. Rose, *skip32: a 32-bit block cipher based on Skipjack*, http://people.qualcomm.com/ggr/QC/skip32.c
6. Coppersmith *et al*, *Cryptanalysis of Stream Ciphers with Linear Masking*, proc. Crypto 2002, LNCS 2442, Springer 2002.
7 P. Ekdahl, T. Johansson, *distinguishing attacks on SOBER-t16 and t32*, proc. Fast Software Encryption, Springer 2002.
8. See http://www.cs.purdue.edu/homes/bb/security99.ps.
9. QUALCOMM proprietary documents.