Some thoughts on the AES process

Lars R. Knudsen

April 15, 1999

Abstract

In this note, it is argued that the Advanced Encryption Standard (AES) should be chosen with a large safety margin. The history of block ciphers shows that the the security levels decrease as a function of the man-years spent in the analysis. Also, we recommend what we think are the candidates best suited for the AES.

1 Security levels and overhead

The AES proposals are required to support at least a block size of 128 bits, and three key sizes of 128, 192, and 256 bits. The hope of NIST is that the end result is a block cipher "with a strength equal to or better than that of Triple-DES and significantly improved efficiency." With the minimum requirements for the key sizes it is clear that an exhaustive key search will be infeasible for many years. Also, with a block size of 128 bits the matching-ciphertext attack requires a huge number of about 2^{64} ciphertext blocks to come into play.

The submitters of most of the algorithms claim a very high level of security. An exhaustive search for the key is often claimed to be the best attack, or it is claimed that an attacker would need all (or more than!!) 2^{128} possible inputs and outputs to succeed.

However, I think that once a few candidates have been selected by NIST, the increased attention of the worlds cryptanalysts will result in new analysis and in levels of security much lower than claimed by the designers.

What do we mean by "security level"? What kind of attacker are we talking about? In my opinion, one should assume the strongest possible attacker. Assume the attacker has access to a black box, which on input any ciphertext or any plaintext of the attackers choice, returns the plaintext respectively ciphertext encrypted under a secret key K. The attacker may ask for any number of texts to be encrypted, his job is to find the secret key K. Also, the attacker knows the description of the block cipher, known as Kerckhoffs' principle, and is allowed to perform encryptions and decryptions of any text and any key of his own choice. One reason to choose such a powerful attacker is, that if even he cannot break the block cipher, then any less powerful, yet more realistic, attacker cannot do it either.

By the *theoretical security level*, I mean the maximum of the numbers in the following triple.

1. The total number of encryptions of the block cipher done by the black-box (that is, the total number of chosen texts.)

2. The total number of encryptions of the block cipher done outside the black-box (that is, the total number of encryptions the attacker does himself.)

3. The total number of words of memory needed by the attack. (One word equals one block in the cipher.)

The history of cryptography is full of examples of systems which are broken despite the designers' *indications* that this should be hard. A few examples follow. In 1917 in an article in *Scientific American* the Vigenère cipher was claimed to be "impossible of translation" [3]. Today it is a student exercise to show that this claim is false. The first version of FEAL [12] had 4 rounds. The first attacks on this version of FEAL required around 100,000 chosen texts. Today, there is an attack which on input 12 chosen texts breaks FEAL with 8 rounds. PES [8] was designed with 8 rounds. Murphy cryptanalysed the cipher and PES was redesigned and renamed IDEA [9]. IDEA was claimed secure after 4 rounds of encryption. Today, there is an attack on 4 rounds of IDEA which is faster than an exhaustive search for the key. Akelarre [1] is a design, which combines features of RC5 and IDEA, and claimed to be secure after 4 rounds. In [4] it was shown that Akelarre with any number of rounds is susceptible to ciphertext-only attacks.

There are many more examples than the above, and they illustrate that the security levels of block ciphers decrease as a function of the man years spent analysing them.

I do not claim to be clairvoyant, and predicting the drop in the security levels of the proposed AES candidates is impossible. However, I state the two following conjectures to make my point and as a thought-provocation.

Conjecture 1 *The theoretical security level of most of the AES candidates will be in the neighborhood of* 2^{100} *or less (with the number of rounds specified by the designers) if about 5 to 10 man-years are spent in serious cryptanalytic effort.*

If this conjecture is true, this would result in security levels lower than the best known key-recovery attacks on triple-DES today.

Conjecture 2 *The theoretical security level of most of the AES candidates will be in the neighborhood of* 2^{80} *or less (with the number of rounds specified by the designers) if about 20 to 30 man-years are spent in serious cryptanalytic effort.*

I could as well have chosen the numbers 2^{90} and 2^{70} or 2^{105} and 2^{87} in the above conjectures. Also, I am aware that "serious cryptanalytic effort" is a vague term, but I have no intentions of even trying to be more specific.

I hope that I have convinced the reader, and NIST in particular, that the algorithm(s) chosen for the AES, should have a large margin of security. The question is, how large this margin should be. All AES candidates are iterated ciphers, where the ciphertext is processed as a function of the plaintext and the key in a number of rounds. Except in a few degenerate cases [4, 2], the security level of an iterated cipher is expected to increase with the number of rounds. As an example, it has been shown that Markov ciphers with independent round keys and primitive transition matrices, are secure against differential cryptanalysis after sufficently many rounds [8, 9]. For ciphers, which do not have independent round keys, one can argue for a similar result, if the key-schedule outputs "random-looking" round keys. In general, the following remains an open question: Given a round function and a key-schedule, how many rounds are needed in the encryption for sufficiently strong encryption?

I recommend the following as a rule of thumb[1] for the AES candidates.

[1] First time I heard a similar statement, was by Massey in his ATS-seminar in 1993 [10]. I do not recall exactly how Massey formulated his recommendation.

Rule of thumb: Let r be the maximum number of rounds, for which there is an attack faster than exhaustive key search. Choose $2r$ rounds for the cipher.

With this rule, several of the AES-candidates need to have an increased number of rounds.

2 Which candidates?

First of all, I will express my concern with the short period of the second round of the AES process. According to the time schedule from NIST, the 15 candidates will be narrowed down to a handful in the summer of 1999, and the winner will be announced August 2000 [11]. This gives about one year to analyse and compare 4-6 candidates. I expect these candidates are the ones for which no serious weaknesses have been reported. In that case, one additional year of research is unlikely to reveal weaknesses in more than one or two of these algorithms, if in any at all.

But, the show must go on. Here are my votes. My favorite candidates are, in alphabetical order:

- RC6 (with 32 rounds)
- Rijndael (with 16 rounds)
- Serpent

A few comments to this recommendation. I prefer "clean" ciphers, which do not mix the group operations, e.g., do not use both exclusive-or and modular additions. Such a mixed use of operations might add more confusion (in Shannon's sense) to a cipher, but also adds more confusion to the designers and cryptanalysts, and it gets harder to be convinced about the security of such proposals. To my knowledge, noone has yet demonstrated to have a clear understanding of how to produce any proof nor "convincing arguments" of the advantage of such an approach.

I like the block cipher Square and therefore also Rijndael, which is a variant hereof. It has a clean and simple design, and it appears to be very strong after sufficiently many rounds. Rijndael is proposed with 10 rounds. I think this will prove to be too optimistic, and I recommend that the number of rounds are increased, e.g., to 16 rounds.

I'm part of the design team of Serpent, so my recommendation of this cipher probably cannot be described as "surprising". However, as opposed to other ciphers, I think Serpent has the correct number of rounds.

Rijndael and Serpent both fall into the "clean" category of ciphers. The simple design and the use of only one group operations, makes it possible to gain some confidence in the design, relatively easier than for other candidates in my opinion.

RC6 uses both exclusive-ors and modular additions, but I'm attracted to the very simple structure of the cipher. Despite the fact that RC6 is not a "clean" cipher, the simplicity of the cipher and its conjectured strength, should make it a candidate for the final five. Also, what speaks in its favor, is that RC6 is the candidate with the easiest-to-remember description (except for the key-schedule). Also for RC6 I suggest to increase the number of rounds, e.g., to 32.

Of the remaining candidates the favourites to make it to the next round are probably Twofish and MARS. Twofish is undoubtedly the most advertised, most mentioned candidate.

And I do believe that MARS will make it to the next round, alone from the impressive list of the authors.

However, no Wiley-book nor any list of authors can hide the fact, that both ciphers are quite complicated using many different operations. I'm not even indicating that they might be weak, but I think that it will take longer time for people to be convinced about the security of these proposals than for the above 3 recommended ciphers.

I do not have a strong opinion about the remaining candidates, except for DEAL [7] which I invented, DFC [5], and LOKI'97 [6]. Without any doubt, DEAL has the most analysed round function of all AES-candidates. However, DEAL was not designed for the AES, but as an alternative for triple-DES [13]. I was encouraged to submit DEAL for AES, and Richard Outerbridge kindly offered to do all the hard work. I think I speak for both of us, when I say, that we never imagined DEAL would be selected as the AES.

May the best candidate win.

I rest my case.

References

[1] G. Alvarez, D. de la Guiaía, F. Montoya, and A. Peinado. Akelarre: a new block cipher algorithm. In *Proceedings of SAC'96, Third Annual Workshop on Selected Areas in Cryptography*, pages 1–14. Queen's University, Kingston, Ontario, 1996.

[2] A. Biryukov and D. Wagner. Slide attacks. In L. R. Knudsen, editor, *Fast Software Encryption, Sixth International Workshop, Rome, Italy, March 1999, LNCS.* Springer Verlag, 1999. To appear.

[3] D.E. Denning. *Cryptography and Data Security.* Addison-Wesley, 1982.

[4] L. R. Knudsen and V. Rijmen. Two rights sometimes make a wrong. Presented at SAC'97. Submitted, 1997.

[5] L. R. Knudsen and V. Rijmen. On the decorrelated fast cipher (DFC) and its theory. In L. R. Knudsen, editor, *Fast Software Encryption, Sixth International Workshop, Rome, Italy, March 1999, LNCS.* Springer Verlag, 1999. To appear.

[6] L. R. Knudsen and V. Rijmen. Weaknesses in LOKI'97. Presented at the 2nd AES Candidate Conference, March, 1999.

[7] L.R. Knudsen. DEAL - a 128-bit block cipher. Technical Report 151, Department of Informatics,University of Bergen, Norway, February 1998. Submitted as an AES candidate by Richard Outerbridge.

[8] X. Lai. On the design and security of block ciphers. In J.L. Massey, editor, *ETH Series in Information Processing*, volume 1. Hartung-Gorre Verlag, Konstanz, 1992.

[9] X. Lai, J.L. Massey, and S. Murphy. Markov ciphers and differential cryptanalysis. In D.W. Davies, editor, *Advances in Cryptology - EUROCRYPT'91, LNCS 547*, pages 17–38. Springer Verlag, 1992.

[10] J.L. Massey. Cryptography: Fundamentals and applications. Copies of transparencies, Advanced Technology Seminars, 1993.

[11] National Institute of Standards and Technology. Advanced encryption algorithm (AES) development effort. `http://www.nist.gov/aes`.

[12] A. Shimizu and S. Miyaguchi. Fast data encipherment algorithm FEAL. In D. Chaum and W.L. Price, editors, *Advances in Cryptology - EUROCRYPT'87, LNCS 304*, pages 267–280. Springer Verlag, 1988.

[13] ANSI X9.F.1. TDEA modes of operation. Draft 5.5, X9.52, March 29, 1996.