Tightness of the Security Bound of CENC

Tetsu Iwata

Department of Computational Science and Engineering,
Nagoya University
Furo-cho, Chikusa-ku, Nagoya, 464-8603, Japan
iwata@cse.nagoya-u.ac.jp
http://www.nuee.nagoya-u.ac.jp/labs/tiwata/

Abstract. CENC (Cipher-based ENCryption) is the recently designed encryption mode for blockciphers. It is provably secure with beyond the birthday bound. In this note, we present a simple distinguishing attack on CENC, and show that the attack is the best attack for some parameter. This proves the tightness of the security bound, and gives a partial answer to the open question posed at FSE 2006.

Key words: Encryption mode, blockcipher, CENC, provable security

1 Introduction

There are many widely used blockcipher encryption modes, including CTR mode, CBC mode, OFB mode and CFB mode. These modes are proved to be secure with the standard birthday bound. That is, if n is the block length of the underlying blockcipher, and σ is the total number of ciphertext blocks that the adversary obtains, then the success probability of a distinguishing attack is $O(\sigma^2/2^n)$ [1, 3]. The analysis is tight. There *is* an adversary that meets the security bound within a constant factor, $\Omega(\sigma^2/2^n)$. This implies that there is no possibility that these modes achieve beyond the birthday bound security, and therefore, the secret key must be updated well before encrypting $2^{n/2}$ blocks of plaintexts.

CENC (Cipher-based ENCryption) was designed to overcome this security limitation [2]. The design goals are: (1) beyond the birthday bound security, (2) security proofs with the standard PRP assumption, (3) highly efficient, (4) single blockcipher key, (5) fully parallelizable, (6) allows precomputation of keystream, and (7) allows random access. CTR mode achieves all the above goals except for the first one, while CENC improves the security of CTR mode without breaking its important advantages.

It was proved in [2] that the success probability of a distinguishing attack against CENC is $O(w\hat{\sigma}^3/2^{2n}+w\hat{\sigma}/2^n)$, where w is a constant (default is $w = 2^8$), and $\hat{\sigma}$ is roughly the same as σ. In [2] the tightness of the security bound was posed as an open question. That is, the question is the existence of a distinguishing attack with success probability $\Omega(w\hat{\sigma}^3/2^{2n} + w\hat{\sigma}/2^n)$, or the proof that the security is better than the above.

Dagstuhl Seminar Proceedings 07021
Symmetric Cryptography
http://drops.dagstuhl.de/opus/volltexte/2007/1016

In this note, we present a simple distinguishing attack on CENC. The attack is based on the observation on the keystream of CENC given in [2], and the success probability is $\Omega(w\sigma/2^n)$, where $\sigma \leq 2^{n/2}$. The strategy is straightforward and obvious. However, it turns out that the attack is the best attack since $O(w\hat{\sigma}^3/2^{2n} + w\hat{\sigma}/2^n)$ is $O(w\hat{\sigma}/2^n)$ when $\hat{\sigma} \leq 2^{n/2}$, and hence it proves the tightness of the security bound. This gives a partial answer to the open question posed at FSE 2006 — the security bound, $O(w\hat{\sigma}^3/2^{2n} + w\hat{\sigma}/2^n)$, is tight for $\sigma \leq 2^{n/2}$.

2 Preliminaries

Notation. If x is a string then $|x|$ denotes its length in bits. If x and y are two equal-length strings, then $x \oplus y$ denotes the xor of x and y. If x and y are strings, then $x \| y$ denotes their concatenation. Let $x \leftarrow y$ denote the assignment of y to x. If X is a set, let $x \xleftarrow{R} X$ denote the process of uniformly selecting at random an element from X and assigning it to x. For a positive integer n, $\{0,1\}^n$ is the set of all strings of n bits. For positive integers n and w, $(\{0,1\}^n)^w$ is the set of all strings of nw bits. For a bit string x and a positive integer n such that $|x| \geq n$, $\mathsf{first}(n, x)$ denote the first n bits of x. For a positive integer n, 0^n denotes the n-times repetition of 0.

Blockciphers. The blockcipher (permutation family) is a function $E : \mathcal{K} \times \{0,1\}^n \to \{0,1\}^n$, where, for any $K \in \mathcal{K}$, $E(K, \cdot) = E_K(\cdot)$ is a permutation on $\{0,1\}^n$. The positive integer n is the block length and an n-bit string is called a block. If $\mathcal{K} = \{0,1\}^k$, then k is the key length.

Let $\mathrm{Perm}(n)$ denote the set of all permutations on $\{0,1\}^n$. This set can be viewed as a blockcipher by considering that each permutation is specified by a unique string. We say P is a random permutation if $P \xleftarrow{R} \mathrm{Perm}(n)$.

The frame, nonce, and counter. CENC takes a positive integer w as a parameter, and it is called a frame width. For fixed positive integer $w \geq 1$ (say, $w = 2^8$), a w-block string is called a frame. A nonce N is a bit string, where for each pair of key and plaintext, it is used only once. The length of the nonce is denoted by ℓ_{nonce}, and it is at most the block length. CENC also uses an n-bit string called a counter, `ctr`. This value is initialized based on the value of the nonce, then it is incremented after each blockcipher invocation. The function for increment is denoted by $\mathsf{inc}(\cdot)$. It takes an n-bit string x (possibly a counter) and returns the incremented x. We assume $\mathsf{inc}(x) = x + 1 \bmod 2^n$, but other implementations also work, e.g., with LFSRs if $x \neq 0^n$.

3 CENC: Cipher-based ENCryption

CENC takes three parameters, a blockcipher, a nonce length, and a frame width.

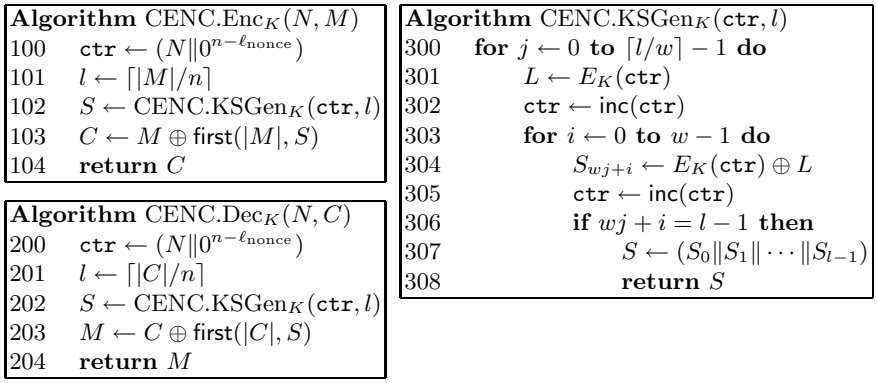

Fig. 1. Definition of the encryption algorithm CENC.Enc (top left), the decryption algorithm CENC.Dec (bottom left), and the keystream generation algorithm CENC.KSGen (right).

Fix the blockcipher $E : \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$, the nonce length ℓ_{nonce} and the frame width w, where $1 \le \ell_{\text{nonce}} < n$. CENC consists of two algorithms, the encryption algorithm (CENC.Enc) and the decryption algorithm (CENC.Dec). Both algorithms internally use the keystream generation algorithm (CENC.KSGen). These algorithms are defined in Figure 1. A picture illustrating CENC.KSGen is given in Figure 2. See [2] for more details of CENC.

Now we are interested in the security of CENC itself, and in the rest of this note we assume that the blockcipher is ideally secure. That is, we assume that the blockcipher is a random permutation $P \xleftarrow{R} \text{Perm}(n)$, and we write CENC.Enc$_P(N, M)$, CENC.Dec$_P(N, C)$, and CENC.KSGen$_P$(ctr, l) where all $E_K(\cdot)$ in Figure 1 and Figure 2 are replaced with $P(\cdot)$.

4 Security of CENC

Security definition for CENC. We consider the strong security notion, which we call indistinguishability from random strings.

An adversary is a probabilistic algorithm (a program) with access to an oracle. Let A be an adversary with access to an oracle $\mathcal{O}(\cdot,\cdot)$, which is either the encryption oracle CENC.Enc$_P(\cdot,\cdot)$ or $\mathcal{R}(\cdot,\cdot)$, and returns a bit. The $\mathcal{R}(\cdot,\cdot)$ oracle, on input (N, M), returns a random string of length $|\text{CENC.Enc}_P(N, M)|$. We say that A is a PRIV-adversary for CENC. Let $(N_0, M_0), \ldots, (N_{q-1}, M_{q-1})$ denote its oracle queries. The adversary is said to be nonce-respecting if $N_0, \ldots, N_{q-1}$ are always distinct, regardless of oracle responses and regardless of A's internal coins. We assume that any PRIV-adversary is nonce-respecting.

The advantage of PRIV-adversary A for CENC is

$$\mathbf{Adv}^{\text{priv}}_{\text{CENC}}(A) \stackrel{\text{def}}{=} \left| \Pr(P \xleftarrow{R} \text{Perm}(n) : A^{\text{CENC.Enc}_P(\cdot,\cdot)} = 1) - \Pr(A^{\mathcal{R}(\cdot,\cdot)} = 1) \right|.$$

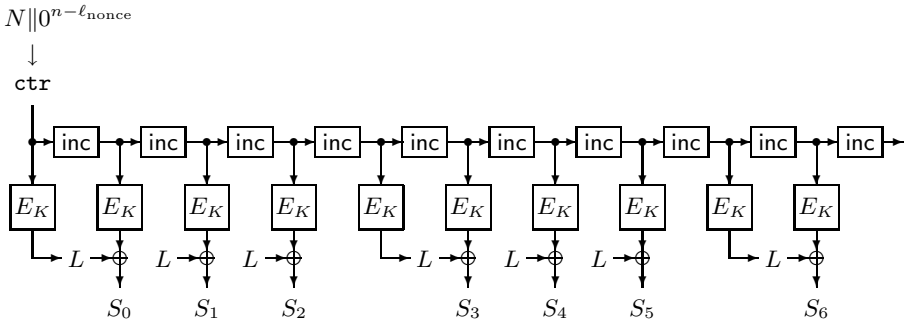

Fig. 2. Illustration of the keystream generation algorithm. This example uses $w = 3$ and outputs $l = 7$ blocks of keystream $S = (S_0, \ldots, S_6)$.

Security result on CENC. Let A be a nonce-respecting PRIV-adversary for CENC, and assume that A makes at most q oracle queries, and the total length of these queries is at most σ blocks, where "the total length of queries" is defined as follows: if A makes q queries $(N_0, M_0), \ldots, (N_{q-1}, M_{q-1})$, then the total length of queries is $\sigma = \lceil |M_0|/n \rceil + \cdots + \lceil |M_{q-1}|/n \rceil$, i.e, the total number of blocks of plaintexts. The following information theoretic result was proved in [2].

Proposition 1. *Let Perm(n), ℓ_{nonce}, and w be the parameters for CENC. Let A be a nonce-respecting PRIV-adversary for CENC making at most q oracle queries, and the total length of these queries is at most σ blocks. Then*

$$\mathbf{Adv}^{\text{priv}}_{\text{CENC}}(A) \leq \frac{(w+1)^4 \hat{\sigma}^3}{w^3 2^{2n+1}} + \frac{(w+1)\hat{\sigma}}{2^{n+1}}, \tag{1}$$

where $\hat{\sigma} = \sigma + qw$.

If we use the rough inequality, $w + 1 \leq 2w$, then we have the simpler form, $\mathbf{Adv}^{\text{priv}}_{\text{CENC}}(A) \leq w\hat{\sigma}^3/2^{2n-3} + w\hat{\sigma}/2^n$. Now if $\hat{\sigma} \leq 2^{(n-3)/2}$, then we have $w\hat{\sigma}^3/2^{2n-3} \leq w\hat{\sigma}/2^n$, while if $\hat{\sigma} > 2^{(n-3)/2}$, then $w\hat{\sigma}^3/2^{2n-3} > w\hat{\sigma}/2^n$. Therefore, the bound in (1) can be written as

$$\begin{cases} \mathbf{Adv}^{\text{priv}}_{\text{CENC}}(A) \leq w\hat{\sigma}/2^{n-1} & \text{if } \hat{\sigma} \leq 2^{(n-3)/2}, \\ \mathbf{Adv}^{\text{priv}}_{\text{CENC}}(A) \leq w\hat{\sigma}^3/2^{2n-4} & \text{otherwise.} \end{cases} \tag{2}$$

5 Distinguishing Attack on CENC

Let Perm(n), ℓ_{nonce}, and w be the parameters for CENC, and fix $q \geq 1$ and $\sigma \geq 1$. For simplicity, we assume σ/w is an integer, but it is easy to handle the general case.

Now we have an obvious observation on the keystream of CENC. That is, if $S_i = (S_i[0], \ldots, S_i[w-1]) \in (\{0,1\}^n)^w$ is one frame of the keystream of CENC, then we always have $S_i[j] \neq S_i[j']$ for any j and j' such that $j \neq j'$. On the other

PRIV-adversary A
100 $M \leftarrow (0^n)^\sigma$
101 $N \leftarrow 0^{\ell_{\text{nonce}}}$
102 $C \leftarrow \mathcal{O}(N, M)$
103 $S \leftarrow M \oplus C$
104 parse S into frames $(S_0, \ldots, S_{\sigma/w-1})$
105 **for** $i = 0$ **to** $\sigma/w - 1$
106 **if** $\text{coll}(S_i) = 1$ **then return** 1
107 **return** 0

Fig. 3. The nonce-respecting PRIV-adversary A for CENC.

hand, if S_i is the truly random string, then $S_i[j] = S_i[j']$ holds with non-zero probability. The above observation was given in [2, p. 315], and this leads to the following simple distinguishing attack.

First, the adversary chooses any message of σ blocks $M = (M_0, \ldots, M_{\sigma-1})$ and asks M to receive $C = (C_0, \ldots, C_{\sigma-1})$, which is either the ciphertext of CENC or the truly random string. Then the adversary retrieves the corresponding keystream $S = (M_0 \oplus C_0, \ldots, M_{\sigma-1} \oplus C_{\sigma-1})$. Now the adversary parses S into σ/w frames, and searches for a collision in the frame. If there is a collision in some frame, then C cannot be the ciphertext of CENC.

More precisely, let A be the nonce-respecting PRIV-adversary for CENC defined in Figure 3. In Figure 3, the oracle $\mathcal{O}(\cdot, \cdot)$ is either the encryption oracle $\text{CENC.Enc}_P(\cdot, \cdot)$ or $\mathcal{R}(\cdot, \cdot)$, and "parse S into frames $(S_0, \ldots, S_{\sigma/w-1})$" is defined by $(S_0, \ldots, S_{\sigma/w-1}) \leftarrow S$, where $|S_i| = wn$ for $0 \le i \le \sigma/w - 1$. The collision search function $\text{coll}(\cdot)$ takes the w block of string $S_i = (S_i[0], \ldots, S_i[w-1]) \in (\{0,1\}^n)^w$ as input, and the output is 1 if and only if $S_i[j] = S_i[j']$ holds for some j and j' such that $j \neq j'$.

Now if $\mathcal{O}(\cdot, \cdot)$ is $\text{CENC.Enc}_P(\cdot, \cdot)$, then the output of the function $\text{coll}(\cdot)$ in line 106 cannot be 1, and thus $\Pr(P \xleftarrow{R} \text{Perm}(n) : A^{\text{CENC.Enc}_P(\cdot,\cdot)} = 1) = 0$. On the other hand, if $\mathcal{O}(\cdot, \cdot)$ is $\mathcal{R}(\cdot, \cdot)$, then S in line 103 is the truly random string of σ blocks since C in line 102 is the truly random string. Since each S_i is the truly random string of w blocks, we have

$$\Pr(\text{coll}(S_i) = 1) = 1 - \prod_{0 \le j \le w-1} \left(1 - \frac{j}{2^n}\right) \ge \left(1 - \frac{1}{e}\right) \frac{w(w-1)}{2^{n+1}} \ge \frac{0.3w(w-1)}{2^n}$$

and thus,

$$\Pr(A^{\mathcal{R}(\cdot,\cdot)} = 1) \ge \frac{0.3w(w-1)}{2^n} \times \frac{\sigma}{w} = \frac{0.3\sigma(w-1)}{2^n}$$

since we have σ/w frames, and these frames are independent. Therefore, we have $\mathbf{Adv}^{\text{priv}}_{\text{CENC}}(A) \ge 0.3\sigma(w-1)/2^n$. This implies that the bound (2) is tight for $\hat{\sigma} \le 2^{(n-3)/2}$, since the bound in (2) is $\mathbf{Adv}^{\text{priv}}_{\text{CENC}}(A) \le w\hat{\sigma}/2^{n-1}$, and hence the above simple distinguisher is the best attack for $\hat{\sigma} \le 2^{(n-3)/2}$.

6 Discussions and Conclusion

Another distinguisher. Another obvious observation on the keystream of CENC is that, if $S = (S_0, \ldots, S_{\sigma-1}) \in (\{0,1\}^n)^\sigma$ is the keystream of CENC, then we always have $S_i \neq 0^n$ for any $0 \leq i \leq \sigma - 1$. On the other hand, if $S = (S_0, \ldots, S_{\sigma-1}) \in (\{0,1\}^n)^\sigma$ is the truly random string, then $S_i = 0^n$ holds with probability $1/2^n$. This observation was also given in [2, p. 315], but it does not give us a better distinguisher than the one presented in Section 5.

Tightness of the security bounds. For any adversary against CENC, the advantage is $O(w\hat{\sigma}^3/2^{2n} + w\hat{\sigma}/2^n)$, which is $O(w\hat{\sigma}/2^n)$ when $\hat{\sigma} \leq 2^{n/2}$, and $O(w\hat{\sigma}^3/2^{2n})$ when $\hat{\sigma} > 2^{n/2}$. We presented a simple distinguisher with advantage $\Omega(w\sigma/2^n)$ for $\sigma \leq 2^{n/2}$, and this implies that (1) the bound $O(w\hat{\sigma}^3/2^{2n} + w\hat{\sigma}/2^n)$ is tight when $\hat{\sigma} \leq 2^{n/2}$, and (2) the simple distinguisher is the best attack for $\hat{\sigma} \leq 2^{n/2}$. The tightness of the bound for $\hat{\sigma} > 2^{n/2}$ is still open. Either we have a distinguisher with advantage $\Omega(w\sigma^3/2^{2n})$, or the bound $O(w\hat{\sigma}^3/2^{2n} + w\hat{\sigma}/2^n)$ can be improved. We conjecture that the bound can be improved to $O(w\hat{\sigma}/2^n)$.

References

1. M. Bellare, A. Desai, E. Jokipii, and P. Rogaway. A Concrete Security Treatment of Symmetric Encryption. Proceedings of *The 38th Annual Symposium on Foundations of Computer Science, FOCS '97,* pp. 394–405, IEEE, 1997.
2. T. Iwata. New Blockcipher Modes of Operation with Beyond the Birthday Bound Security. *Fast Software Encryption, FSE 2006,* LNCS 4047, pp. 310–327, Springer-Verlag, 2006. Full version is available at `http://www.nuee.nagoya-u.ac.jp/labs/tiwata/`.
3. P. Rogaway. Nonce-based Symmetric Encryption. *Fast Software Encryption, FSE 2004,* LNCS 3017, pp. 348–358, Springer-Verlag, 2004.