Beating Shannon requires BOTH efficient adversaries AND non-zero advantage

Yevgeniy Dodis, NYU

In this note we formally show a well known (but not well documented) fact that in order to beat the famous Shannon lower bound on key length for one-time-secure encryption, one must *simultaneously* restrict the attacker to be efficient, and also allow the attacker to break the system with some non-zero (i.e., negligible) probability. Our proof handles probabilistic encryption, as well as a small decryption error.

1 Definitions

Let $(\mathsf{Enc}, \mathsf{Dec})$ be any encryption scheme with key space $\mathcal{K}$ and message space $\mathcal{M}$. In general, we use capital letters for random variables, and lower case letters for specific values; e.g., M, C, S denote appropriately defined random messages, ciphertexts and keys, while m, c, s denote some specific value of those. In the description below, every random variable (e.g., M_1, S, etc.) not explicitly defined in terms of other random variables (e.g., $C = \mathsf{Enc}_S(M_1)$) will always be uniform over its corresponding domain.

Remark 1 *We allow the encryption algorithm* Enc *to be* probabilistic*. However, since all our proofs easily handle this case, we will not explicitly put the randomness R in our notation. I.e., when we write $\mathsf{Enc}_s(m)$, we always really mean a random variable $\mathsf{Enc}_s(m; R)$, even for fixed m and s (let alone when either M or S are random). Similarly, when some encryption is computed inside some probability, we do not explicitly put the choice or R under* $\Pr$*. E.g.,* $\Pr_S[\mathsf{Enc}_S(m) = c]$ *really means* $\Pr_{S,R}[\mathsf{Enc}_S(m; R) = c]$.

Definition 1 A (possibly probabilistic) encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ is called (t, ε)-*secure* if for any message $m_0 \in \mathcal{M}$, and any adversary Eve running in time at most t, it holds

$$\left| \Pr_{S,M_1}[Eve(M_1, \mathsf{Enc}_S(m_0)) = 1] - \Pr_{S,M_1}[Eve(M_1, \mathsf{Enc}_S(M_1)) = 1] \right| \le \varepsilon \tag{1}$$

Namely, Eve cannot tell encryption of m_0 from encryption of uniformly random M_1. $\diamondsuit$

Remark 2 *The above definition is slightly weaker than the more traditional definition stating that for any messages $m_0, m_1 \in \mathcal{M}$, and any adversary Eve running in time at most t, it holds*

$$\left| \Pr_S[Eve(\mathsf{Enc}_S(m_0)) = 1] - \Pr_S[Eve(\mathsf{Enc}_S(m_1)) = 1] \right| \le \varepsilon$$

This is OK, since we are proving a lower bound. In any event, by hybrid argument the gap between "epsilons" is at most a factor of 2.

Definition 2 A (possibly probabilistic) encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ is called γ-*wrong*

$$\Pr_{S,M}[\mathsf{Dec}_S(\mathsf{Enc}_S(M)) = M] \geq 1 - \gamma \tag{2}$$

Namely, decrypting encryption of a random message almost never results in an error. $\diamondsuit$

2 Main Result

According to the values of $t \in [0, \infty]$ and $\varepsilon \in [0, 1]$ one can obtain different notions of security. Here we show that to beat Shannon bound $|\mathcal{K}| \geq |\mathcal{M}|$ (corresponding to $t = \infty$ and $\varepsilon = 0$), we really need both t to be small and ε to be non-zero. Our proof also handles decryption error γ.

Theorem 1 *Assume* $(\mathsf{Enc}, \mathsf{Dec})$ *is at most* γ*-wrong. Then:*

- **Small error needed.** *Let* v *denote maximum bit length of a plaintext plus ciphertext. If* $(\mathsf{Enc}, \mathsf{Dec})$ *is* $(v, 0)$*-secure, then* $|\mathcal{K}| \geq |\mathcal{M}|(1 - \gamma)$*.*
- **Small time needed.** *Let* d *denote maximum decryption time. If* $(\mathsf{Enc}, \mathsf{Dec})$ *is* $(|\mathcal{K}|d, \varepsilon)$*-secure, then* $|\mathcal{K}| \geq |\mathcal{M}|(1 - \varepsilon - \gamma)$*.*

In other word, to beat the Shannon bound in a non-trivial way for any "functional" (e.g., $\gamma < 1 - 1/poly$*) encryption, one must* simultaneously *restrict Eve to be efficient, as well as allow for some non-zero (but possibly negligible) probability* ε *of security failure.*

Proof of First Part. We show that $(v, 0)$-security implies that for any messages $m_0, m_1 \in \mathcal{M}$ and ciphertext c_1, it holds:

$$\Pr_{S,M_1}[M_1 = m_1 \text{ and } \mathsf{Enc}_S(m_0)) = c_1] = \Pr_{S,M_1}[M_1 = m_1 \text{ and } \mathsf{Enc}_S(m_1) = c_1)] \tag{3}$$

To show Equation (3), consider the following $Eve_{m_1,c_1}(m, c)$ running in time $t = v$: *output* 1 *if and only if* $m = m_1$ *and* $c = c_1$. Since $\varepsilon = 0$, it is immediate that Equation (1) $\Rightarrow$ Equation (3) for the $Eve = Eve_{m_1,c_1}$ above. In other words, $(v, 0)$-security implies that distributions $(M_1, \mathsf{Enc}_S(m_0))$ and $(M_1, \mathsf{Enc}_S(M_1))$ are *identical*: $(M_1, \mathsf{Enc}_S(m_0)) \equiv (M_1, \mathsf{Enc}_S(M_1))$.

Now, pick a fresh random key S_1 and look at

$$\Delta = \Pr_{S,M_1,S_1}[\mathsf{Dec}_{S_1}(\mathsf{Enc}_S(m_0)) = M_1] \tag{4}$$

On the one hand, it is clear that

$$\Delta \leq \frac{1}{|\mathcal{M}|} \tag{5}$$

Indeed, if we let $M = \mathsf{Dec}_{S_1}(\mathsf{Enc}_S(m_0))$, then M_1 is perfectly uniform and independent of M. So $\Pr[M = M_1] \leq \frac{1}{|\mathcal{M}|}$, indeed. On the other hand, by Equation (3), since $(M_1, \mathsf{Enc}_S(m_0)) \equiv (M_1, \mathsf{Enc}_S(M_1))$, we can rewrite Equation (4) as

$$\Delta = \Pr_{S,M_1,S_1}[\mathsf{Dec}_{S_1}(\mathsf{Enc}_S(M_1)) = M_1] \tag{6}$$

But then

$$\begin{aligned}
\Delta &= \Pr_{S,M_1,S_1}[\mathsf{Dec}_{S_1}(\mathsf{Enc}_S(M_1)) = M_1] \\
&\geq \Pr[S = S_1] \cdot \Pr_{M_1,S_1}[\mathsf{Dec}_{S_1}(\mathsf{Enc}_{S_1}(M_1)) = M_1] \\
&\geq \frac{1}{|\mathcal{K}|} \cdot (1-\gamma)
\end{aligned}$$

where the last inequality used Equation (2). Comparing the inequality above with Equation (5), we get $\frac{1}{|\mathcal{K}|} \cdot (1-\gamma) \leq \Delta \leq \frac{1}{|\mathcal{M}|}$, which implies $|\mathcal{K}| \geq (1-\gamma)|\mathcal{M}|$. $\square$

Proof of Second Part. We show that $(|\mathcal{K}|d, \varepsilon)$-security implies $|\mathcal{K}| \geq |\mathcal{M}|(1-\varepsilon-\gamma)$. For that, consider the following attacker Eve of complexity $t = |\mathcal{K}|d$:

$Eve(m_1, c_1)$: *output* 1 *if and only if there exists* $s_1 \in \mathcal{K}$ *s.t.* $\mathsf{Dec}_{s_1}(c_1) = m_1$.

Now, let us compute both probabilities when we apply Equation (1) to this Eve. First,

$$\begin{aligned}
\Pr_{S,M_1}[Eve(M_1, \mathsf{Enc}_S(M_1)) = 1] &= \Pr_{S,M_1}[\exists s_1 \text{ s.t. } \mathsf{Dec}_{s_1}(\mathsf{Enc}_S(M_1)) = M_1] \\
&\geq \Pr_{S,M_1}[\mathsf{Dec}_S(\mathsf{Enc}_S(M_1)) = M_1] \\
&\geq 1-\gamma
\end{aligned}$$

where the last inequality used Equation (2). By Equation (1), we get

$$\Pr_{S,M_1}[Eve(M_1, \mathsf{Enc}_S(m_0)) = 1] \geq \Pr_{S,M_1}[Eve(M_1, \mathsf{Enc}_S(M_1)) = 1] - \varepsilon \geq 1-\varepsilon-\gamma \tag{7}$$

On the other hand,

$$\begin{aligned}
\Pr_{S,M_1}[Eve(M_1, \mathsf{Enc}_S(m_0)) = 1] &= \Pr_{S,M_1}[\exists s_1 \text{ s.t. } \mathsf{Dec}_{s_1}(\mathsf{Enc}_S(m_0)) = M_1] \\
&\leq \sum_{s_1} \Pr_{S,M_1}[\mathsf{Dec}_{s_1}(\mathsf{Enc}_S(m_0)) = M_1]
\end{aligned}$$

However, for any s_1, if we let $M = \mathsf{Dec}_{s_1}(\mathsf{Enc}_S(m_0))$, then M_1 is perfectly uniform and independent of M. So $\Pr[M = M_1] \leq \frac{1}{|\mathcal{M}|}$, which means that

$$\Pr_{S,M_1}[Eve(M_1, \mathsf{Enc}_S(m_0)) = 1] \leq \sum_{s_1} \frac{1}{|\mathcal{M}|} = \frac{|\mathcal{K}|}{|\mathcal{M}|} \tag{8}$$

Combining Equation (7) and Equation (8), we get $1-\varepsilon-\gamma \leq \frac{|\mathcal{K}|}{|\mathcal{M}|}$ or $|\mathcal{K}| \geq |\mathcal{M}|(1-\varepsilon-\gamma)$. $\square$

Tightness: Both bounds are nearly tight, which can be shown by tweaking the generalization of the one-time pad (OTP) encryption for general cardinality N message spaces (not just the power of 2, which can be accomplished by addition modulo N). For simplicity, we only do it for two special cases $\varepsilon = 0$ and $\gamma = 0$, leaving the common generalization as a (tedious) exercise.

First, assume $\varepsilon = 0$. Take any $|\mathcal{M}|$ of cardinality N, and any subset $\mathcal{M}_0 \subseteq \mathcal{M}$ of cardinality $N(1-\gamma)$. Start with the OTP scheme over $\mathcal{M}_0$ (so that $|\mathcal{K}| = N(1-\gamma)$ as well), and enlarge it to all of $\mathcal{M}$ by taking any fixed $m_0 \in \mathcal{M}_0$ and defining $\mathsf{Enc}_s(m_1) = \mathsf{Enc}_s(m_0)$, for $m_1 \in \mathcal{M} \backslash \mathcal{M}_0$. The addition of these γN messages (which decrypt incorrectly) to our OTP does not affect the security of the scheme (since $\mathsf{Enc}(m_0)$ is perfectly secure), but creates a decryption error with probability γ with $|\mathcal{K}| = |\mathcal{M}|(1-\gamma)$.

Second, assume $\gamma = 0$. Now, for any $\mathcal{M}$ of cardinality N, take the OTP for $\mathcal{M}$ (so that $|\mathcal{K}| = N$), and simply remove $\varepsilon N/2$ keys from $\mathcal{K}$, defining the actual set $\mathcal{K}_0$ of $N(1-\varepsilon/2)$ keys. One can imagine sampling a key $s \leftarrow \mathcal{K}_0$ but first sampling the key $s \leftarrow \mathcal{K}$ and claiming that Eve unconditionally won the game if $s \in \mathcal{K} \backslash \mathcal{K}_0$. Equivalently, we can always actually run Eve on a fully uniform key s from $\mathcal{K}$, but then declare Eve victorious anyway if $s \in \mathcal{K} \backslash \mathcal{K}_0$. Clearly, when s is fully uniform, Eve has probability exactly $1/2$ telling apart encryptions of m_0 from M_1, so now her probability is at most $1/2 + \varepsilon/2$, creating distinguishing advantage at most ε with $|\mathcal{K}_0| = |\mathcal{M}|(1-\varepsilon/2)$.