PERFECT SECRECY AND ADVERSARIAL INDISTINGUISHABILITY

BURTON ROSENBERG
UNIVERSITY OF MIAMI

Contents

1. Perfect Secrecy

Claude Shannon introduced an entropy model for information, and applied it to secrecy in communications. It supposes a source of information, Alice, which chooses among a set of possible messages. There is associated with this choice a likelihood that Alice would chose a particular message. Symbols are then sent across a channel to Bob. These symbols should refine Bob's likelihood function, emphasizing the likelihood of Alice's chosen message. Alice and Bob share a secret key, but this key is not shared with the eavesdropper Eve. Eve sees the symbols on the channel, and understands as well the likelihood by which Alice chooses messages. However, because Eve does not share the secret key, Eve should find no use for these symbols. Her likelihood function should not be refined.

Alice's likelihood is represented as a probability distribution over a message space. The messages space M is assumed finite. A probability distribution $P(M)$ is a map from M to $[0, 1]$, satisfying the axioms of a probability distribution; but might be better to think of $P(M)$ as a map from events in M, that is, subsets of M, to $[0, 1]$. Events are things we can learn about the message, such as "the event that the message contains a vowel". Generally, for every message m, the event "the message

Date: January 23, 2016.

is m" is an admissible event, and so there is no difference between $P : M \to [0,1]$ and $P : Pwr(M) \to [0,1]$.

Complete uncertainty on Alice's choice corresponds to the uniform distribution: $P(M = m) = 1/|M|$. In this case, Bob will have no preferred message that he can act on in advance of any symbols. Complete certainty corresponds to,

$$P(M = m) = \begin{cases} 1 & \text{if } m = m^* \\ 0 & \text{else} \end{cases}$$

In this case, each time Alice picks a message, that message is m^*, and it is even unnecessary that she sends symbols. Bob can act in advance on the knowledge that when Alice chooses, she will choose m^*.

If the symbols placed on the channel are from the space C, the ciphertext, we wish that Bob learns from this symbol. Consequently the probability is updated. However, as Eve learns nothing, the probability on M conditioned on C should be unchanged. This is the Shannon definition of Perfect Secrecy:

Definition 1.1. An encryption scheme as *Perfect Secrecy* if for every probability distribution $P(M)$ and for every $c \in C$, the probability distribution $P(M \mid c)$ is the same as the a priori likelihood distribution $P(M)$.

1.1. **A Perfectly Secret Cipher.** The *Vernam Cipher*, or *One Time Pad*, is an example of a perfectly secret cipher. It works on a message space of bits, and the key is a stream of bits matching the length of the message. We discuss the case of a one bit message.

Alice and Bob flip a fair coin (or a coin of bias β). The result $k \in \{0,1\}$ is their secret key. Given the message $m \in \{0,1\}$, Alice forms ciphertext $c = k \oplus m$. Bob receives c and recovers m by,

$$c \oplus k = m \oplus k \oplus k = m \oplus 0 = m.$$

Theorem 1.1. The Vernam Cipher has Prefect Secrecy if and only if $\beta = 1/2$.

Proof. Because $\beta = 1/2$, half of the 0 messages end up transmitting a 0, and half of the 1 messages end up transmitting a 0; so half of the transmissions are 0. Leaving that the other half of the transmissions are a 1.

$$P(C = 0) = P(C = 1) = 1/2$$

Ciphertext c obtained from message m exactly when $k = c \oplus m$,

$$P(C = c \mid M = m) = P(k = c \oplus m) = 1/2.$$

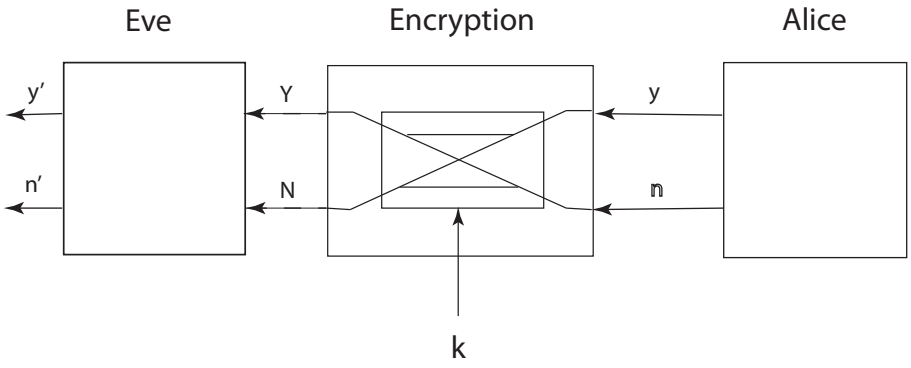

FIGURE 1. Vernam Cipher.

Using Bayes Theorem to put these facts together,

$$\begin{aligned} P(M = m \mid C = c) &= P(C = c \mid M = m)P(M = m)/P(C = c) \\ &= (1/2)P(M = m)/(1/2) \\ &= P(M = m) \end{aligned}$$

satisfying the definition of Perfect Secrecy.

For the situation when $\beta \neq 1/2$, see below. □

Hence, to Eve, without access to k, the likelihood of a message after observation of C is the same as without any observation. The channel reveals nothing because Eve's beliefs can ignore it entirely without loss. However, Bob can update his probability to achieve certainty about which message was chosen by Alice.

1.2. **Odds Ratio and Bias.** Assume, through some fault of key generation that the key coin is not fair, $\beta \neq 1/2$. Then the system does not have perfect secrecy. In the case of a probability distribution that places weight $1/2$ on each of two elements in M, the ciphertext is a clear hint about the message chosen. Into the encryption box, both a 0 or a 1 are equally likely, but out of the encryption box, it is more likely a 0 if the input were a 0, and more likely a 1 if the input were a 1.

We calculate that $P(m = c) = \beta$, in this case. We can do this directly from the equation $k = m \oplus c$.

However, what of perfect secrecy in the case were the message distribution is not uniform. We can assume w.l.o.g. $\beta \geq 1/2$ and $\delta = P(M = 0) \geq P(M = 1)$. What

should Eve do? Let's assume that Eve chooses the most likely input to the encryption box given the observed output. We express this as an odds ratio:

$$P(M = 0 \mid C = c)/P(M = 1 \mid C = c)$$

Intuitively, if $c = 0$, then it is more likely that $m = 0$. $M = 0$ is already the a priori guess, and a biased coin only adds likelihood to this guess if $c = 0$.

So:

$$\begin{aligned}\frac{P(M = 1 \mid C = 1)}{P(M = 0 \mid C = 1)} &= \frac{P(C = 1 \mid M = 1)P(M = 1)/P(C = 1)}{P(C = 1 \mid M = 0)P(M = 0)/P(C = 1)} \\ &= \frac{P(C = 1 \mid M = 1)}{P(C = 1 \mid M = 0)} \frac{P(M = 1)}{P(M = 0)} \\ &= \frac{\beta}{1-\beta} \frac{1-\delta}{\delta}\end{aligned}$$

Eve will guess 1 if,

$$\frac{\beta}{1-\beta}\frac{1-\delta}{\delta} > 1,$$

which reduces to $\beta > \delta$.

Hence Eve continues to guess the more a priori outcome and achieves success probability $\delta = P(M = 0)$ unless the coin bias β rises about δ, in which case Eve guesses that m is c and achieves success probability β.

1.3. **Conditions for Perfect Secrecy.** A few necessary conditions for perfect secrecy are immediate. It must be that the key space is at least the size the space of ciphertext messages. If not, then given a ciphertext, decrypting by each key will give the space of possible messages. Some message cannot be achieved, because there are too many messages to be covered by the keys. An uncovered message will have a posteri probability 0, and for the purposes of the proof, a distribution on messages can be assumed with a priori probability non-zero for this message.

The space of ciphertext messages must be at least as large as the space of messages, else there will be some two messages encrypting to the same ciphertext, and the requirement that decryption be the inverse of encrypting cannot be achieve.

Shannon's theorem for perfect secrecy assumes equal sizes for the key space, message space, and ciphertext space and gives two conditions necessary and sufficient for perfect secrecy.

(1) The choice if k from K is made uniformly at random; and
(2) For each m in M and c in C there is a unique k in K such that $c = E_k(m)$.

One can check that these theorem holds for the Vernam Cipher described above.

The exclusive-or over a 0-1 space can be replaced with a randomly chosen shift in the space of modular integers. If each shift were chosen independently, uniformly

at random for each character in a text, then the requirements of Perfect Secrecy are satisfied.

However, shift ciphers do not have completely random choices for keys for each column. While the Vignere cipher improves this by having several groups of columns which independently chosen shifts, it is not sufficient. In a simple shift cipher seeing a ciphertext with a repeated cipher symbol reduces to zero the likelihood of the message having differing letters in arranged where the cipher symbols are the same.

2. Adversarial Indistinguishability

Shannon introduced encryption in the context of entropy and information. The model gives absolute answers. It does not depend on any assumptions about the attacker. The encryption keeps the message secret because the channel symbols support equally any hypothesis about which message is more likely, in the presence of that symbol.

However, the price is that keys must be a large as message. This is not practical. Computational complexity allows for the possibility of a practical scheme that is effectively secure. A claim can be made that, although possible, it is not with the capacity of a practical computation to extract information from the ciphertext. This is possible because some functions are (hypothesized to be) one way: a quick calculation can lead to a result whose inversion is (possibly) intractable.

To help introduce these notions we rephrase perfect security with an adversary, that first will have unbounded power but then we will ask that the adversary fit within computational bounds.

Adversarial Indistinguishability is comprised of an encryption system,

$$\Pi = (Gen, Enc, Dec),$$

and an Adversary $\mathcal{A}$, and a game where:

(1) The adversary chooses two message, m_0 and m_1 from the message space.
(2) A key k is generated at random by Gen, and a fair coin chooses $b \in \{0, 1\}$.
(3) The adversary $\mathcal{A}$ is presented $Enc(k, m_b)$
(4) The adversary returns b'.

The value of $PrivK^{eav}_{\mathcal{A},\Pi}$ is $(b == b')$.

Definition 2.1. An encryption scheme Π has *Perfect Adversarial Indistinguishability* if, for any adversary $\mathcal{A}$,

$$P(PrivK^{eav}_{\mathcal{A},\Pi} = 1) = 1/2,$$

where the probability space includes the distribution of k by Gen, the coin flip b, the choice of m_0 and m_1 by $\mathcal{A}$, and any randomness in the functions $\mathcal{A}$ and Enc.

2.1. **Prefect Secrecy and Indistinguishability are Equivalent.**

Theorem 2.1. An encryption scheme has Perfect Adversarial Indistinguishability if and only if it has Perfect Secrecy.

Proof. Assume the systems does not have Perfect Adversarial Indistinguishability. There exists an Adversary $\mathcal{A}$ that has an advantage. If the adversary $\mathcal{A}$ has any advantage then it must have an advantage for some two message m_0, m_1. Fix these as the message choice by $\mathcal{A}$.

We can consider the adversary is deterministic, because there is no gain to randomizing its guesses. Hence $\mathcal{A}$ reduces to a guess function $\mathcal{A} : C \to \{0, 1\}$. The advantage is then expressed as,

$$P(\mathcal{A}(E_k(m_b)) = b) > 1/2$$

For any given c, and $i \in [0, 1]$, define I_c^i as,

$$I_c^i = \{k \mid E_k(m_i) = c\}$$

and I_c as,

$$I_c = I_c^0 \cup I_c^1$$

Since the overall advantage of $\mathcal{A}$ is a sum of advantages for each c, weighted by the $P(c)$, there must be some c such that $\mathcal{A}(c) = i$, and i is more like the correct answer to challenge c,

$$P(I_c^i \mid I_c) > P(I_c^{1-i} \mid I_c)$$

Since $P(I_c^i \mid I_c) + P(I_c^{1-i} \mid I_c) = 1$, then neither can be $1/2$. So

$$P(I_c^0 \mid I_c) = P(M = m_0 \mid C = c) \neq 1/2$$

yet

$$P(M = m_0) = 1/2.$$

So $P(M = m_0 \mid C = c) \neq P(M = m_0)$ and the scheme does not have Perfect Secrecy.

Assume that the scheme does not have Perfect Secrecy. An equivalent definition for Perfect Secrecy is that $\forall m_0, m_1 \in M$ and $c \in C$,

$$P(C = c \mid M = m_0) = P(C = c \mid M = m_1)$$

because if, $P(m \mid c) = P(m)$, then

$$P(c \mid m) = P(m \mid c)P(c)/P(m) = P(c) = P(c \mid m').$$

Therefore since the scheme is assumed not perfectly secret, there exists m_0, m_1 and c such that

$$P(C = c \mid M = m_0) \neq P(C = c \mid M = m_1).$$

Construct $\mathcal{A}$ as follows: the adversary offers messages m_0 and m_1. On receiving c', if $c' \neq c$ it answers a random b', else if $c' = c$ it answers 0. Note:

$$\begin{aligned} P(C = c) &= P(C = c \mid b = 0)P(b = 0) + P(C = c \mid b = 1)P(b = 1) \\ &= 1/2\,(P(C = c \mid b = 0) + P(C = c \mid b = 1)) \\ &\neq 1/2\,(P(C = c \mid b = 0) + P(C = c \mid b = 0)) \\ &= P(C = c \mid b = 0). \end{aligned}$$

When $c' \neq c$, the adversary is correct with probability $1/2$. However, when $c' = c$, then

$$\begin{aligned} P(PrivK^{eav}_{\mathcal{A},\Pi} = 1 \mid C = c) &= P(b = 0 \mid C = c) \\ &= P(C = c \mid b = 0)P(b = 0)/P(C = c) \\ &\neq P(b = 0) = 1/2 \end{aligned}$$

The advantage when $c' = c$ is $1/2$, but in the case that $c' \neq c$, the advantage is not $1/2$. Averaging over the cases, the average is not $1/2$ and so the scheme is not Adversary Indistinguishable.

□