CSE 5852: Lecture 3

Sarah Peck

September 7, 2016

1 Review of Last Class

Last time, we mathematically defined perfect secrecy in terms of probability. Today, we give another definition and show a perfectly secure encryption scheme. After today, we've covered all material needed for the first problem set. Remember that if you work with others, think through problems on your own first, add collaborators' names to homework, and write out answers on your own.

2 Probability

Last class we covered σ-algebras and probability measures. These allow us to define the probability of events in the outcome space. Often, we will care about the probability of functions of events.

Definition 1 (Random Variable)**.** *A random variable, X is a function $X : \Omega \to R$. We write $X = x$ to say the set of all outcomes a such that $X(a) = x$.*

The probability of all things that produce x can be written as $Pr[X = x] = Pr[A]$ where $A = \{a|X(a) = x\}$.

We will use capital letters to denote a random variable (for example, X) and lower case letters to denote a possible value of that random variable (for example, x).

Dice example A die is rolled. If the value a is even, you win \$1. If the value is odd, you lose \$1.

$$X(a) = \begin{cases} 1 & \text{when } a \text{ is even} \\ -1 & \text{when } a \text{ is odd} \end{cases}$$

Definition 2 (Induced Probability Distribution)**.** *The induced probability distribution of X is defined as follows*

$$\Pr[X = x] = \Pr[A] \text{ where } A = \{a|X(a) = x\}.$$

For the above example, the induced probability distribution is $Pr[X = 1] = 0.5$ and $\Pr[X = -1] = -.5$.

Card example A card is drawn. The value x of the card is as follows:

$$x = \begin{cases} 10 & \text{if face card} \\ \text{value of card} & \text{if numeric} \\ 1 & \text{if Ace} \end{cases}$$

Note that this is distinct from the original event space which was a single card. We have projected a card to its "numeric value." The probability that X has a value 10 is $Pr[X = 10] = \frac{4}{13}$.

Definition 3 (Expectation (mean, average))**.** *Consider a random variable X with range in the real numbers, $\mathbb{R}$. The expectation of X, denoted $\mathbb{E}[X]$ is the sum of the values of X multiplied by their probability. That is, $\mathbb{E}[X] = \sum_{x \in X} \Pr[X = x]x$.*

For example, $\mathbb{E}[X]$ of the card example is 6.53.

Independence of random variables Variables are independent when the probability of one variable x doesn't change when the other variable y is known. The two following definitions are equivalent (similarly to the case of events):

$$\forall x, y, \Pr[X = x | Y = y] = \Pr[X = x]$$
$$\forall x, y, \Pr[X = x \cap Y = y] = \Pr[X = x] \Pr[Y = y].$$

3 Secrecy of a Channel

The reason we went through the effort of defining random variables is that there are two choices in the case of an encryption scheme, the choice of the message m and the choice of the key k. The ciphertext is a function of these and can be viewed as a random variable.

Our goal is to now formalize Perfect Secrecy. We said that the message should be independent of the ciphertext. The components of an encryption scheme are:

- Message distribution M
- Key distribution K
- A pair of algorithms Enc and Dec

Correctness A correct channel must:

- Be correct for all messages m

- Work for almost all keys k. In the discussion that follows, we will provide correctness for all keys k this is only because our schemes will provide this.

Definition 4 (Correctness)**.** *Let M be a message space and K be a key space. Let* $\mathsf{Enc}, \mathsf{Dec}$ *be two (possibly randomized) algorithms. We say that* $(\mathsf{Enc}, \mathsf{Dec})$ *is an encryption scheme if*

$$\forall m \in M, k \in K, m = \mathsf{Dec}(k, \mathsf{Enc}(k, m))$$

Notice this definition did not consider the distribution of keys or messages. Could we do the same thing for secrecy?

Can we have secrecy for all distributions K?

No. We may want it for all possible keys, but what about the key distribution K where $Pr[K = 0 \dots 0] = 1$? If we allow such a distribution of keys, we can learn $m = \mathsf{Dec}(0 \dots 0, c)$, so m cannot be secret. A secure scheme needs to include *Enc*, *Dec*, and a key distribution.

Do we want secrecy for all distributions M?

Yes. Since we don't know what our secure channel will be used for, we need to make sure we can securely send for all distributions M. (This is goal, we need to show it is possible.)

Definition 5 (Perfect secrecy)**.** *Let K be a distribution and* Enc *be an algorithm.* Enc *is perfectly secret if for any message m and any message distribution M,*

$$\Pr[M = m | \mathsf{Enc}(K, M) = c] = \Pr[M = m]$$

Remember, in the above definition upper case values are random variables and lower case values are outcomes of these variables. So the expression $\mathsf{Enc}(K, M) = c$ is the fixing the random variable that represents the encryption of an unknown message and unknown key to be c.

For example, if our adversary determined that there was a 1 in 5 chance that our message was $m =$ "Attack," and then the adversary saw the ciphertext c, the probability of the message being "Attack" would remain the same. This should be true no matter what the adversaries belief about the message being "Attack" was before seeing the message. Put another way, the ciphertext should not change the adversaries belief about the message.

Why is *Enc* calculated using the distributions in the equation?

This is the equation of what the adversary sees and knows. The adversary may know the key distribution K and message distribution M, but the adversary does not know the specific key or message. If m is specified in the equation, then $\Pr[M = m] = 1$! It would be known!

Perfect secrecy is a little complicated to work with as you need to prove something for all message distributions. We'll present an alternative definition that seems a little simpler.

Die Roll	Probability M_1	Probability M_2
1	$\frac{1}{6}$	$\frac{1}{3}$
2	$\frac{1}{6}$	0
$\vdots$	$\vdots$	$\vdots$
6	$\frac{1}{6}$	$\frac{1}{6}$

Table 1: Example message distributions. To have perfect secrecy, an encryption scheme should preserve the probabilities in a single column after seeing the ciphertext.

Definition 6 (Shannon secrecy)**.** *Enc is Shannon secure if* $\forall m_1, m_2$*:*

$$Pr[Enc(K, m_1) = c] = Pr[Enc(K, m_2) = c]$$

Note that this definition only talks about a key distribution (not a message distribution). Somewhat surprisingly, these two definitions are equivalent.

Theorem 1. *Shannon secrecy is equivalent to perfect secrecy.*

Proof. Proof 1: Perfect secrecy $\Rightarrow$ Shannon secrecy

We start by assuming that for any $m_1 \in \mathcal{M}$ and any M, $\Pr[M = m_1] = \Pr_{K,M}[M = m_1 | \mathsf{Enc}(K, M)]$.

$$\begin{aligned}
Pr_M[M = m_1] &= Pr_{K,M}[M = m_1 | Enc(K, M) = c] \\
&= \frac{Pr_{K,M}[M = m_1 \cap Enc(K, M) = c]}{Pr_{K,M}[Enc(K, M) = c]} \\
&\quad \text{(by the definition of conditional probability)} \\
&= \frac{Pr_{K,M}[M = m_1 \cap Enc(K, m_1) = c]}{Pr_{K,M}[Enc(K, M) = c]} \\
&\quad \text{(since } m_1 \text{ is fixed is the first part of the numerator,} \\
&\quad \text{we can substitute it in the second part)} \\
&= \frac{Pr_M[M = m_1] Pr_K[Enc(K, m_1) = c]}{Pr_{K,M}[Enc(K, M) = c]} \\
&\quad \text{(the first part of the numerator depends only on } M\text{, the second} \\
&\quad \text{only on } K \text{ and these are independent, thus their probabilities multiply)}
\end{aligned}$$

Simplifying we have that

$$\frac{\Pr_K[\mathsf{Enc}(K, m_1) = c]}{\Pr_{K,M}[\mathsf{Enc}(K, M) = c]} = 1.$$

Note this is also true for any m_2 (the steps above did not depend on m_1) which implies that

$$\frac{\Pr_K[\mathsf{Enc}(K, m_1) = c]}{\Pr_{K,M}[\mathsf{Enc}(K, M) = c]} = \frac{\Pr_K[\mathsf{Enc}(K, m_2) = c]}{\Pr_{K,M}[\mathsf{Enc}(K, M) = c]}.$$

This completes the proof that perfect secrecy implies Shannon secrecy.
Proof 2: Shannon secrecy ⇒ perfect secrecy

Consider some arbitrary distribution M, in the same way as above we can write:

$$\begin{aligned}\Pr_{K,M}[M = m_1 | Enc(K, M) = c] &= \frac{Pr_M[M = m_1]Pr_K[Enc(K, m_1) = c]}{Pr_{K,M}[Enc(K, M) = c]} \\ &= \frac{Pr_M[M = m_1]Pr_K[Enc(K, m_1) = c]}{\sum_{m \in M} Pr_K[Enc(K, m) = c] \Pr[M = m]}\end{aligned}$$

(we can always rewrite a random variable as the sum of each outcome)

$$= \frac{Pr_M[M = m_1]Pr_K[Enc(K, m_1) = c]}{Pr_K[Enc(K, m_1) = c] \sum Pr[M = m]}$$

(we know that this value is the same no matter the message)

$$= Pr_M[M = m]$$

□

One-time pad Based on the above definitions we'll try and construct a secure construction using a single message bit and single key bit. Given that we have a two-bit system, $m = \{0, 1\} \quad k = \{0, 1\}$, we need a function that results in an equal number of ones and zeros for each value of m: XOR. Let k be a uniform n-bit string.

$$\begin{aligned}Enc(k, m) &= k \oplus m = c \\ Dec(k, c) &= k \oplus c = m\end{aligned}$$

Theorem 2. *The one-time pad is Shannon secure (and thus perfectly secure).*

Sketch. Take any c, m pair. There is exactly one key that is consistent. That is, for any ciphertext, the chance that it resulted from an encryption of m is $1/|K|$. □

Problems with the one-time pad

1. Key cannot be reused. What does the distribution $c_1 = k \oplus m_1, c_2 = k \oplus m_2$ look like?
2. Key must be uniform. If it is not, parts of m can be known.
3. Key must be as long as the message. This requires an amount of infrastructure that makes the one-time pad not logistically feasible for most cases.