Homework 3

CS161 Computer Security, Spring 2008

Assigned 2/27/08

Due 3/10/08

1. **Authentication and Key Exchange**

   The Needham-Schroeder protocol has a number of variants; in particular, the version in Figure 12.3 of the Gollman textbook differs from the one presented in lecture. As explained in Section 12.3.3, the textbook variant is vulnerable to a replay attack which tricks B into reusing an old session key (which may have since been compromised).

   (a) (3 points)

   Modify this version of the Needham-Schroeder protocol so that it prevents this attack. You may assume that all three parties have synchronized clocks, but you must not add any messages to the protocol. That is, the protocol should still consist of a message from A to S, S to A, A to B, B to A, and A to B, but the fields in each message may differ from the original protocol. Express your answer in the protocol notation of Figure 12.3 and note any checks conducted by each party after receiving each message.

   (b) (5 points)

   If synchronized clocks are not available, it is also possible to prevent this replay attack by adding additional messages to the protocol and altering the content of other messages. An incomplete version of such a protocol is given on the following page.

$$
\begin{aligned}
&A \rightarrow B : A \\
&B \rightarrow A : \_\_\_\_\_ \\
&A \rightarrow S : A, B, n_a, \_\_\_\_\_ \\
&S \rightarrow A : E_{K_{as}}(n_a, B, K_{ab}, \_\_\_\_\_) \\
&A \rightarrow B : \_\_\_\_\_ \\
&B \rightarrow A : E_{K_{ab}}(n_b) \\
&A \rightarrow B : E_{K_{ab}}(n_b - 1)
\end{aligned}
$$

In this version, A and B exchange an initial round of messages before A contacts S. After that point, the rest of the protocol is very similar to the version of Needham-Schroeder given in Figure 12.3; in particular, the last two messages are the same.

Each of the blanks represents one or more fields that have been omitted. The size of the blanks is not meant to indicate anything about the content of those fields.

Give values for these omitted fields so that the resulting protocol is not vulnerable to the replay attack. Try to avoid including anything beyond what is necessary.

2. **More Authentication and Key Exchange**

Consider the following protocol, which is an alternative to Needham-Schroeder.

$$
\begin{aligned}
&A \rightarrow B : n_a, A, B, E_{K_{as}}(n'_a, n_a, A, B) \\
&B \rightarrow S : n_a, A, B, E_{K_{as}}(n'_a, n_a, A, B), E_{K_{bs}}(n_a, A, B, n_b) \\
&S \rightarrow B : n_a, E_{K_{as}}(n'_a, K_{ab}), E_{K_{bs}}(n_b, K_{ab}) \\
&B \rightarrow A : n_a, E_{K_{as}}(n'_a, K_{ab})
\end{aligned}
$$

In the above, n_a and n'_a are nonces selected by A and n_b is a nonce selected by B.

(a) (2 points)

Do you think this protocol is vulnerable to the same attack as the Needham-Schroeder protocol? That is, is it possible to replay messages in order to cause A or B to reuse an old session key? Briefly explain why you think it is or is not vulnerable to such an attack.

(b) (7 points)

In any case, that protocol is vulnerable to a less severe attack. Specifically, it is possible for an adversary to trick A and B into establishing a session with keys that do not match. That is, A will subsequently attempt to communicate with B using some session key K_{ab} while B will being using some key $K'_{ab} \neq K_{ab}$.

Show how this can be done. Assume the attacker can intercept messages and alter, drop, or replay them, but assume that the attacker has not compromised any session keys or any of the long term keys A and B share with S.

3. **Generating Random Values**

In order to generate a cryptographic key or a seed for a pseudorandom number generator, it is necessary to collect some (truly) random data. Common sources of random data on a system include a high resolution (e.g., nanosecond) clock, latencies between disk seeks, delays between keystrokes or mouse movements, etc.

All of these sources contain some amount of truly unpredictable information, but they are in general not sources of uniform, uncorrelated bits. In order to obtain uniformly distributed random values from such sources, the common practice is to pass them through a cryptographic hash function. However, it is possible to condition random data through simpler methods.

Suppose you wish to generate some number of bits, each of which should be 0 or 1 with equal probability, independent of the other bits. You have a biased coin which turns up tails (0) with some unknown probability p, where $0 < p < 1$, and turns up heads (1) with probability $1-p$. However, you do not have a computer with which to compute a hash function. In this scenario, the following simple algorithm can be used.

1. Flip the coin twice.
2. If the outcome is tails, then heads, write down a 0. If the outcome is heads, then tails, write down a 1. Otherwise, do not write anything.
3. Repeat from Step 1 until the required number of bits has been generated.

If we view the coin flips as generating a string of input bits, this algorithm can be viewed as passing over the string two bits at a time, applying the

following transformation.

"00" $\rightarrow$ " "
"01" $\rightarrow$ "0"
"10" $\rightarrow$ "1"
"11" $\rightarrow$ " "

For example, the input string

"10 00 00 01 01 10 01 10 00 00 01 00 10"

would be translated to

"1 0 0 1 0 1 0 1".

Regardless of the value of p, each output bit is equally likely to be 0 or 1 because the pair of input bits "01" has the same probability of occurring as the pair "10". Furthermore, the probability of each output bit being 0 or 1 will be independent of the other output bits.

(a) (3 points)
Using this method, if the number of input bits is ℓ (which may be assumed to be even) and the probability of each input bit being 0 is p, what is the expected number of output bits generated?

(b) (8 points)
Design a more efficient scheme, that is, one which provides a higher expected ratio of output bits to input bits.
Like the one above, your scheme should take the form of a mapping from groups of input bits to groups of output bits, and you should specify it in this form.
Each of the output bits produced by your scheme should be 0 or 1 with equal probability, independent of the other output bits. For any value of p such that $0 < p < 1$, your scheme should produce more output bits on average than the given scheme.

(c) (4 points)
If the number of inputs bits is ℓ and the probability of each bit being 0 is p, what is the expected number of bits generated by your new scheme?