ECE 3170 Cryptographic Hardware for Embedded Systems

Fall 2025

Assoc. Prof. Vincent John Mooney III

Georgia Institute of Technology

Homework 4, 100 pts.

Due Friday Sept. 12 prior to 11:55pm

(please turn in homework electronically on Canvas)

As always this semester, you are required to solve any and all homework questions alone.

Notation (copied from the lecture notes for ease of reference):

- $\{X\}$ is a set of elements of type X
- m is a message in plaintext
  - m is composed of smaller blocks m_i suitable for individual encryption steps
  - $m = \{m_i\}$
- c_i is ciphertext corresponding to message block m_i
- c is ciphertext corresponding to message m
- Enc_k is encryption with key k
  - $c \leftarrow Enc_k(m)$
- Dec_k is decryption with key k
  - $m \leftarrow Dec_k(c)$
- MAC_k is generation of a message authentication code t with key k
  - $t \leftarrow Mac_k(m)$ or, alternatively, $t \leftarrow Mac_k(c)$
- $<a,b>$ is a concatenation of a followed by b

1) (20 pts.) Hash functions and collisions.
   a. (10 pts.) Use your own words (do **not** copy from **any** source) to describe what is *collision resistance* for a hash function.
   b. (5 pts.) Consider MD5 which takes a 512-bit input and produces a 128-bit hash output. Is it possible for each unique input to be associated with a unique output? If your answer is no, for MD5 with a 512-bit input, do at least the majority of inputs have unique associated hash outputs? If your answer is no the majority of inputs do not have unique associated hash outputs, do at least 10% of inputs have unique associated hash outputs? If your answer is that 10% of the inputs do not have unique associated hash outputs, do at least 1% of inputs have unique associated hash outputs? Please explain your answers; a correct sequence of "yes" or "no" answers with no explanations will receive zero points.
   c. (5 pts.) Does *target-collision resistance* (also known as *second preimage resistance*) imply collision resistance? Why or why not?

2) (30 pts.) Consider the following scenario describing an attack on documentation signed by a hash function.

You are given a legitimate message x_1, a fraudulent message x_2, and a one-way hash function (keyless) denoted by $h()$ which produces an m-bit output where $m = 10$.

The goal of this attack is to substitute the fraudulent message for the legitimate; for example, if $h(x_1) = h(x_2)$ then the attacker is done! However, since this is extremely unlikely, the plan then is for the attacker to modify x_1 and x_2 in minor ways – e.g., introducing meaningless spaces, tabs or extra punctuation – to produce messages x_1' and x_2' with the result that $h(x_1\text{'}) = h(x_2\text{'})$. If this attack is successful, the attacker aims to convince an unsuspecting party to sign x_1' (or x_1) but later implement x_2'.

If $h(x_1) \neq h(x_2)$, the following approach is taken: first a modification x_1' is tried, and then second a modification x_2' is tried. Third, a different modification x_1'' is tried, and then fourth a different modification x_2'' is tried (i.e., the modifications alternate). Assume that all $h(x_1)$, $h(x_2)$, $h(x_1\text{'})$, $h(x_2\text{'})$, $h(x_1\text{''})$, $h(x_2\text{''})$, etc., values are stored in an efficient manner (in other words, ignore the time it takes to search the stored values). Further assume that each output of the hash function is random, i.e., each time $h()$ is calculated for any x_1, x_2, x_1', x_2', x_1'', x_2'', x_1''', x_2''', etc., the m-bit hash output is random and has no correlation with previous hash outputs. Do not assume that the output of the hash function is necessarily unique; after all, a truly random process does have a non-zero possibility of repeating an output bit pattern.

a. (5 pts.) Consider the starting point with documents x_1 and x_2, and recall that the output distribution of $h()$ is random. What is the probability that $h(x_1) = h(x_2)$? (HINT: the obvious answer is correct!)

b. (5 pts.) Let $\{h(x_1\text{'})\}$ denote the set of all hash values of x_1 or of any of the attempted modifications of x_1. Similarly, let $\{h(x_2\text{'})\}$ denote the set of all hash values of x_2 or of any of the attempted modifications of x_2. After the fourth step described above (i.e., after a different modification x_2'' is tried), what is the size of $\{h(x_1\text{'})\}$ and what is the size of $\{h(x_2\text{'})\}$ (i.e., how many elements are contained in each set)?

c. (15 pts.) How many tries are needed to arrive at a 50% chance that one of the hash entries in $\{h(x_1\text{'})\}$ matches a hash entry in $\{h(x_2\text{'})\}$?

d. (5 pts.) As with earlier homeworks, there may be multiple valid assumptions for the answer to part c (just prior to this part, part d, of question 2 on homework 4). Please clearly state at least one important assumption you made in your answer to part c above. NOTE: full credit will be given only for reasonable assumptions (i.e., clearly incorrect assumptions will lose all or nearly all of the points).

[PLEASE TURN TO THE NEXT PAGE!]

3) (50 pts.) The lecture “Crypto VIII: Two Attacks on Encryption” covered the padding-oracle attack on Construction 3.30.
   a. (10 pts.) Describe first stage of the padding-oracle attack on PKCS #5 padding. More specifically, use the notation on the first page of this homework and also use your own words (do not copy from the book chapter section that covers this attack; a recommendation is to use your own notes taken while listening to Lecture 14 Cryptography Part VIII) to describe the first stage of the attack where *b*, the amount of padding, is learned.
   b. (15 pts.) Continue to describe the padding-oracle attack. In particular, use the notation above and your own words to describe the second stage of the attack where *B0*, the final byte of the message, is learned by the attack. You should assume that *b*, the amount of padding, has already been learned.
   c. (25 pts.) Extend the padding-oracle attack to the byte *B1* which is the byte just before the final byte *B0* of the message. You should assume that *B0*, the final byte of the message, and *b*, the amount of padding, have already been learned.

YOU MAY NOT CONSULT HOMEWORK SOLUTIONS OF THESE EXACT PROBLEMS FROM OTHER COURSES. ALL LABORATORY SUBMISSIONS MUST INCLUDE YOUR NAME, COURSE NUMBER, SECTION, AND THE HOMEWORK SET NUMBER. ALL SHEETS MUST BE STAPLED TOGETHER. ALL WRITING MUST BE EASY TO READ (FOR EXAMPLE, YOU MAY HAVE TO WRITE ONLY ON ONE SIDE OF EACH SHEET OF PAPER THAT YOU SUBMIT AND MAY NOT BE ABLE TO USE RECYCLED PAPER). FAILURE TO FOLLOW INSTRUCTIONS MAY RESULT IN ZERO POINTS. ALL WORK MUST BE YOUR OWN. NO PLAGIARISM IS ALLOWED, AND YOU MUST PROPERLY REFERENCE ALL SOURCES OF YOUR INFORMATION – ALTHOUGH YOU SHOULD NOT LOOK FOR AND MAY NOT CONSULT “SOLUTIONS” AVAILABLE FROM OTHER SOURCES (TO REPEAT, YOU MAY NOT CONSULT HOMEWORK SOLUTIONS OF THESE EXACT PROBLEMS FROM OTHER COURSES!).