Midterm Solutions

Question	Parts	Points
1: Security Basics	1	1
2: Lab 0: Passwords	9	9
3: Lab 0: Collision Finding	1	0
4: Lab 1: Merkle Trees	5	5
5: One-Way functions	1	1
6: MACs	3	3
7: Signatures	3	3
8: Bad randomness in ECDSA	3	3
9: Public Key Infrastructure	2	2
10: Encryption	9	9
11: Randomness in Encryption	3	3
Total:		39

Name: _____

- This is an open book exam: you can use your notes or any material released by us this term. You cannot use the internet.
- Any form of collaboration is *strictly* forbidden.
- If you find a question ambiguous, be sure to write down any assumptions you make.

This midterm exam is printed double-sided!

Problem 1. [1 point] **Security Basics** (1 part)

(a) [1 point] Setting a task of recognizing the existence of objects (e.g., bridges) in a set of images before creating a user account is used to try to prevent bots from creating thousands of bogus accounts. An attacker using crowd sourcing to solve object recognition is an example of:

- An insufficient security goal.
- An insufficient attack model.
- A buggy implementation.

Solution: An insufficient attack model.

Problem 2. [9 points] **Lab 0: Passwords** (9 parts)
After lab 0, Rico is not happy with the security that passwords provide. He is building out a web application and wants to come up with a better solution to secure his users.

(a) [1 point] Before Rico tries to develop new technologies, he wants to ensure he understands the current limitations of passwords. He first remembers that he is supposed to hash passwords instead of storing them in plaintext. Rico is trying to pick a good hash function to use. Please select all of the properties below that Rico would like his hash function to have:

- One-way
- Collision resistant
- Fast
- Slow
- Keyed (i.e. $H := \{0,1\}^\lambda \times \{0,1\}^* \to \{0,1\}^\lambda$)

Solution: One-way, collision resistant, slow

(b) [1 point] Rico recalls something about entropy. He knows that higher entropy will make it harder for an adversary to guess his password. He thus suggests enforcing passwords to be longer, with the hope that longer passwords will have more entropy. What is a limitation of this approach in practice? (Select the best answer)

- Our server will now have to store much more data making us a larger target.
- People extend passwords in predictable ways, limiting the expected gain in entropy.
- Since it is hard for people to remember the long passwords, they will write them down and create another attack surface.
- There is no limitation, we get so much more entropy!

Solution: People extend passwords in predictable ways, limiting the expected gain in entropy.

(c) [1 point] Rico now recalls salting. He knows that if he adds salts to each password before hashing them, he can increase the adversaries time to attack all of our users. He wants to ensure that each individual must be attacked separately by ensuring (with high probability) each user has a different salt. If Rico expects $O(2^{32})$ users for his service, what is the minimum number of bits (out of the possible options) needed to ensure this property?

- 10
- 30
- 50
- 100

Solution: 100

(d) [1 point] Rico is not pleased with any of the ideas he has come up with so far. Luckily, he attended all of the 6.160 lectures and learned about public key cryptography and digital signatures. He thinks that he can come up with a protocol that uses digital signatures to authenticate users. He wants to develop a basic scheme to validate his idea. For the rest of question 2, assume Rico and his users have access to a secure digital signature scheme σ that has Gen, Sign, and Ver algorithms as defined in lecture.

Rico first worries about registering a new user. Registration will create an account for a new user and store some data that can be used to authenticate users in the future. Which algorithm should Rico use in registration?

- Gen
- Sign
- Ver

Solution: Gen

(e) [1 point] What information should the `user` remember for future logins:

Solution: The secret key.

(f) [1 point] What information should the `server` remember for future logins:

Solution: The public key.

(g) [1 point] After registration is complete, Rico now must develop his authentication protocol. Authentication will be the process in which the server verifies who the user says that they are. This process can have multiple interactions between the user and the server. What is one idea that could be used as the central idea in Rico's new protocol?

- Hash and sign
- Challenge-Response
- Encrypt then MAC

Solution: Challenge-Response.

(h) [1 point] Given the idea above and σ, please explain the protocol used to authenticate a user:

Solution: 1) Server sends a random nonce r.
2) User signs r with their corresponding secret key.
3) Server runs Ver with the user's public key to check the signature on r.

(i) [1 point] What is one benefit to this protocol over typical passwords?

Solution: Answers may vary. Some potential solutions are below:
1) The server only holds public keys, so it is no longer storing private information (or a function on private information). Thus, it is less valuable for someone to break in to.
2) No one can guess someone's "password" as there is not a password to guess anymore.

Problem 3. [0 points] **Lab 0: Collision Finding** (1 part)
Recall the attack from Part 4 of Lab 0. That is given a hash function

$$H(x) := \{0,1\}^* \to \{0,1\}^\lambda,$$

we run the tortoise and the hare algorithm defined on a graph G such that $V = \{v \in \{0,1\}^\lambda\}$ and $E = \{(u,v) : H(u) = v\}$. Can we run this attack against $H(x) =$ SHA-256(x)?

Solution: No

Problem 4. [5 points] **Lab 1: Merkle Trees** (5 parts)
Recall the implementation of a client and Merkle hash store from Lab 1. The client holds only the Merkle tree's root hash, and can make `get` and `put` requests to the store. To find/insert a key-value pair, the traversal path is determined by the bits of $H_{kv}(\texttt{key})$.

(a) [1 point] Suppose an adversary takes control of an existing Merkle tree store. Without changing the contents of the store, briefly describe how the adversary could convince the client that the lookup of an intermediate node value has succeeded (i.e. convince the client that an intermediate node is a leaf in the Merkle tree).

Solution: The adversarial store will return a truncated proof containing the siblings up to the intermediate node, and return the intermediate node at the end of the proof as a leaf value.

(b) [1 point] Which of the following is a possible defense against this attack?

- Use a collision-resistant hash function
- Include the leaf value at the end of the proof instead of H(leaf)
- Add a prefix of 0 to the leaf node values, 1 to all the intermediate node values

Solution: Add a prefix of 0 to the leaf node values, 1 to all the intermediate node values

(c) [1 point] **T/F** When a client makes the request 'put foo bar' to the store, all the hash values in the Merkle tree must be recomputed.

Solution: False

(d) [1 point] Which properties are needed from the hash function used in a Merkle tree? Select all that apply.

- Collision-resistant
- Random oracle
- Pseudorandom
- One-way

Solution: Collision-resistant, one-way

(e) [1 point] The advantage of Merkle hash over Merkle-Damgard hash is:

- Merkle hash is parallelizable whereas Merkle-Damgard is not.
- Merkle hash takes inputs of arbitrary length whereas Merkle-Damgard does not.
- Merkle hash provides more privacy.

Solution: Merkle hash is parallelizable whereas Merkle-Damgard is not.

Problem 5. [1 point] **One-Way functions** (1 part)
T/F Every one-way function is collision resistant.

Solution: False

Problem 6. [3 points] **MACs** (3 parts)

(a) [1 point] **T/F** If MAC is secure (existentially unforgeable against chosen message attacks) for messages of length 128 then MAC' is secure for messages of length 256, where MAC' is defined by $\mathsf{MAC}'(K, M_1||M_2) = (\mathsf{MAC}(K, M_1), \mathsf{MAC}(K, M_2))$ where M_1 and M_2 are each of length 128 bits.

Solution: False

(b) [1 point] **T/F** The main building block of MAC schemes used in practice is AES.

Solution: True

(c) [1 point] Consider the MAC scheme defined by $\mathsf{MAC}(K, M) = \mathsf{AES}(K, M)$.
Which of the following assumptions on AES is sufficient to make the MAC secure (existentially unforgeable against chosen message attacks)

- Collision resistant hash function
- One-way function
- Pseudorandom function

Solution: Pseudorandom function

Problem 7. [3 points] **Signatures** (3 parts)

(a) [1 point] **T/F** A secure signature scheme (existentially unforgeable against adaptive chosen message attacks) must be randomized.

Solution: False

(b) [1 point] **T/F** A secure digital signature scheme provides privacy and authenticity of the message being signed.

Solution: False

(c) [1 point] **T/F** The RSA signature scheme is known to be secure (existentially unforgeable against chosen plaintext attacks) assuming Factoring is hard.

Solution: False

Problem 8. [3 points] **Bad randomness in ECDSA** (3 parts)
ECDSA is one of the most widely used digital-signature schemes. In lab 2, you had the chance to dive deeper into the math of the algorithm, as well as analyze ways it can be insecure.

(a) [1 point] **T/F** If two different messages are signed using the same nonce in ECDSA and the same secret key, an attacker can derive the secret key efficiently.

Solution: True

(b) [1 point] **T/F** ECDSA is only used for message encryption, not for digital signatures.

Solution: False

(c) [1 point] **T/F** ECDSA is not quantum-resistant, meaning that a sufficiently large quantum computer could break it using Shor's algorithm.

Solution: True

Problem 9. [2 points] **Public Key Infrastructure** (2 parts)

(a) [1 point] Give one advantage of the above scheme.

Solution: Revocation is easier than in conventional PKI, since the Certificate Authority can inform the user/web browser that the public key of the website is no longer valid, as opposed to maintaining revocation lists on the user side. The expiration of the certificate is also easier to implement.

(b) [1 point] Give one disadvantage of the above scheme.

Solution: The load on the Certificate Authority (or Authorities) will be very high, since every user's every visit to a website requires a request to some Certificate Authority (CA) for a public key. This scheme also exposes the CA signing key to direct Internet attacks, and the CA can track/record user activities.

Problem 10. [9 points] **Encryption** (9 parts)

(a) [1 point] Given a (symmetric) CPA secure encryption scheme $(\mathsf{Enc}, \mathsf{Dec})$ and a secure MAC (existentially unforgeable against adaptive chosen message attacks), one can obtain a CCA secure (symmetric) encryption scheme, as follows:

- $\mathsf{Enc}'((K, K'), M) = (\mathsf{Enc}(K, M), \mathsf{MAC}(K', M))$
- $\mathsf{Enc}'((K, K'), M) = (C = \mathsf{Enc}(K, M), \mathsf{MAC}(K', C))$
- $\mathsf{Enc}'((K, K'), M) = (t = \mathsf{MAC}(K', M), \mathsf{Enc}(K, t))$

- None of the above

Solution: $\mathsf{Enc}'((K, K'), M) = (C = \mathsf{Enc}(K, M), \mathsf{MAC}(K', C))$

(b) [1 point] **T/F** If $(\mathsf{Enc}, \mathsf{Dec})$ is a (symmetric) CPA secure encryption for messages of length 128 then $(\mathsf{Enc}', \mathsf{Dec}')$ is a (symmetric) CPA secure for messages of length 256, where Enc' is defined by $\mathsf{Enc}'(K, M_1||M_2) = (\mathsf{Enc}(K, M_1), \mathsf{Enc}(K, M_2))$, where M_1 and M_2 are each of length 128 bits.

Solution: True

(c) [1 point] **T/F** (Symmetric) CCA secure encryption hides the length of the message being encrypted.

Solution: False

(d) [1 point] **T/F** (Symmetric) CPA encryption provides authenticity.

Solution: False

(e) [1 point] **T/F** In practice, symmetric encryption is faster than public key encryption.

Solution: True

(f) [1 point] **T/F** The main building block of symmetric encryption schemes used in practice is SHA-256.

Solution: False

(g) [1 point] Consider the encryption scheme defined by $\mathsf{Enc}(K, M) = (r, \mathsf{AES}(K, r) \oplus M)$. Which of the assumptions on AES is sufficient to make Enc CPA-secure

- Collision resistant hash function
- One-way function
- Pseudorandom function

Solution: Pseudorandom function

(h) [1 point] **T/F** Consider the same encryption scheme defined by $\mathsf{Enc}(K, M) = (r, \mathsf{AES}(K, r) \oplus M)$. It is CCA secure assuming AES is a Random Oracle.

Solution: False

(i) [1 point] **T/F** The Diffie-Hellman key exchange protocol is known to be secure against active attacks.

Solution: False

Problem 11. [3 points] **Randomness in Encryption** (3 parts)
Answer the following questions about randomness in encryption.

(a) [1 point] **T/F** Any CPA secure encryption scheme remains secure if the random string is always fixed to zero.

Solution: False

(b) [1 point] **T/F** Randomness is only required when using symmetric encryption algorithms.

Solution: False

(c) [1 point] **T/F** CPA encryption scheme remains secure even if a random string is reused across multiple encryption operations to save computation time

Solution: False