

Midterm Solutions

Question	Parts	Points
1: True or False	7	7
2: User authentication	1	8
3: Signatures	4	8
4: RSA	3	15
5: Lab 1	1	10
6: Key exchange	4	17
7: TLS security	1	6
8: Law and technology	1	6
9: Course survey	3	3
Total:		80

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. [7 points] **True or False** (7 parts)

Please answer **T** or **F** for the following. *No justification is needed (nor will be considered).*

- (a) [1 point] Implementation bugs can subvert security even if the designers had the correct goal.

Solution: True

- (b) [1 point] A one-way function must also be collision-resistant.

Solution: False

- (c) [1 point] In practice, we use AES as a collision-resistant hash function.

Solution: False.

- (d) [1 point] There are message authentication codes (MACs) that provide λ -bit security and that have a MAC tag length of λ bits, under standard cryptographic assumptions.

Solution: True.

- (e) [1 point] We believe that Lamport signatures (instantiated with a suitable one-way function) will resist attacks by large-scale quantum computers.

Solution: True.

- (f) [1 point] Let N be an RSA modulus with public exponent $e = 3$. Define the function $F: \mathbb{Z}_N^* \rightarrow \mathbb{Z}_N^*$ as $F(x) := (x + 7)^3 \bmod N$. Under the RSA assumption, the function F is a one-way permutation.

Solution: True.

- (g) [1 point] The hash-then-sign paradigm is useful because it avoids the need to use a collision-resistant signature scheme.

Solution: False.

Problem 2. [8 points] **User authentication** (1 part)

Ben Bitdiddle runs a popular web site, in which users create accounts using their email address as their username. Ben Bitdiddle is worried about the overhead of storing a separate salt for every user's account in his web site's password database. Ben devises the following alternative plan: his web site will store a single global salt s , and for every user, the database will store the user's username (email address) and $H(s||username||password)$.

Is this scheme as good as using individual salts for every password? Explain why, or describe an attack for which this scheme would give an adversary some advantage.

Solution: Yes, this scheme is as secure, because the adversary cannot pre-compute hashes for any passwords before the database is compromised, and every user's password is hashed differently, just as with per-password salts.

Problem 3. [8 points] **Signatures** (4 parts)

- (a) [2 points] What is one benefit of RSA signatures over EC-DSA signatures?

Solution: Fast verification.

- (b) [2 points] What is one benefit of EC-DSA signatures over RSA signatures?

Solution: Short signature size, fast key generation, short public keys.

- (c) [2 points] What is one benefit of Lamport signatures over EC-DSA signatures?

Solution: Post-quantum security, security from OWF only.

- (d) [2 points] What is one benefit of RSA signatures over Lamport signatures?

Solution: Shorter signature size, many-time signatures

Problem 4. [15 points] **RSA** (3 parts)

Let N be an RSA modulus with public exponent $e = 3$ and private exponent d (i.e., $ed \equiv 1 \pmod{\phi(N)}$). The full-domain-hash signature scheme uses a hash function $H: \{0, 1\}^* \rightarrow \mathbb{Z}_N^*$. The scheme computes the signature on a message $m \in \{0, 1\}^*$ as:

$$\sigma \leftarrow H(m)^d \pmod{N}.$$

- (a) [3 points] Your friend (who has taken 6.042 but *not* 6.1600) proposes removing the hash function H from the full-domain-hash signature scheme. With this modified scheme, a signature on a message $m \in \mathbb{Z}_N^*$ is $\sigma \leftarrow m^d \pmod{N}$. Show that an attacker, given only the public key (N, e) , can produce a valid forged signature σ^* on some message $m^* \in \mathbb{Z}_N^*$. Your answer should include a valid message-signature pair: (m^*, σ^*) .

Solution: The message m^* can be any perfect cube with $m < N$. For example (m^*, σ^*) can be $(1, 1)$, $(8, 2)$, $(27, 3)$, ...

- (b) [5 points] Consider the full-domain-hash signature scheme instantiated with some hash function H . Say that an attacker can find two messages $m_0, m_1 \in \{0, 1\}^*$, such that $H(m_0) = H(m_1)^2 \in \mathbb{Z}_N^*$. Explain how the attacker can use (m_0, m_1) to win the signature security game.

Solution:

- The attacker asks for a signature on m_0 . This is the value $\sigma = H(m_0)^d$.
- The attacker outputs $\sigma^* \leftarrow \sigma^2 \in \mathbb{Z}_N^*$ as its forged signature.

The forged signature σ^* is valid since

$$(\sigma^*)^e = ((H(m_0)^d)^2)^e = (H(m_0)^{2d})^e = H(m_0)^{2ed} = H(m_0)^2 \in \mathbb{Z}_N^*.$$

Problem continues on the next page...

- (c) [7 points] Model the hash function $H: \{0, 1\}^* \rightarrow \mathbb{Z}_N^*$ as a truly random function. How many times would an attacker have to evaluate the hash function H , on average, to find messages $m_0, m_1 \in \mathbb{Z}_N^*$ such that $H(m_0) = H(m_1)^2 \in \mathbb{Z}_N^*$.

Solution: The analysis here is similar to the Birthday Problem. The average number is around \sqrt{N} .

Problem 5. [10 points] **Lab 1** (1 part)

Ben Bitdiddle is hired by the 6.1600 course staff to defend the lab 1 key-value store against the attacks covered in the lab. Ben decides to change how key-value leaf nodes are hashed, by inserting a slash separator between the key and the value, as follows:

```
def H_kv(key, val):  
    return H(key + "/" + val)
```

How can you modify the attack in scenario 2 (many fake key-value pairs) so that it still works against Ben's modified design?

Solution: Construct an adversarial tree of 1000 key-value pairs whose two root children hashes have a slash byte in their UTF-8 encoding (for at least one of the two of them). Concatenate these two root children hashes, split them at the slash character, and return the parts before and after the slash, respectively, as the key and value from `attack_key_value()`.

Problem 6. [17 points] **Key exchange** (4 parts)

In his final project for his undergraduate security class, Ralph Merkle proposed a key-exchange protocol based on hash functions.

The protocol uses a hash function $H: \mathbb{Z}_n \rightarrow \{0, 1\}^{256}$, where n is on the order of 2^{60} , and proceeds as follows:

- Alice picks \sqrt{n} random numbers $a_1, \dots, a_{\sqrt{n}} \in \mathbb{Z}_n$ and sends $H(a_1), \dots, H(a_{\sqrt{n}})$ to Bob.
- Bob picks \sqrt{n} random numbers $b_1, \dots, b_{\sqrt{n}} \in \mathbb{Z}_n$ and sends $H(b_1), \dots, H(b_{\sqrt{n}})$ to Alice.
- If there exist $i, j \in \{1, \dots, \sqrt{n}\}$ such that $H(a_i) = H(b_j)$:
 - Alice uses a_i as her shared secret with Bob.
 - Bob uses b_j as his shared secret with Alice.

(If there are many such (i, j) pairs, Alice and Bob use the lexicographically first one.)

Model the hash function as a truly random function.

- (a) [2 points] Explain why Alice and Bob will agree on a shared secret with constant probability.

Solution: The Birthday bound.

- (b) [2 points] How much time does it take Alice to generate her message to Bob? Assume that evaluating $H(\cdot)$ takes a constant amount of time.

Solution: It takes \sqrt{n} time.

Problem continues on the next page...

- (c) [3 points] If an attacker eavesdrops on the communication between Alice and Bob, much time does it take the attacker to recover the shared secret $a_i = b_j$?

Solution: It takes the attacker $\Omega(n)$ time.

(d) [10 points] Define the *goodness* of a key-exchange protocol to be the ratio:

$$\text{goodness} = \frac{\text{the attacker's running time}}{\text{Alice's running time}},$$

where the attacker's running time is the time required to recover the shared secret.

- What is the goodness of Merkle's protocol?
- What is the goodness of Diffie-Hellman key exchange in \mathbb{Z}_p^* for a large prime p , assuming that it takes $O(1)$ time to multiply two integers in \mathbb{Z}_p^* ? (In reality the time for a big-integer multiplication grows with p .)
- What is the goodness of elliptic-curve Diffie-Hellman key exchange in a group of prime order p , assuming that it takes $O(1)$ time to perform a single elliptic-curve point operation \boxplus ?

Solution: Computing an exponentiation by repeated squaring requires $O(\log p)$ multiplications.

- Merkle's protocol has goodness n/\sqrt{n} .
- Diffie-Hellman key exchange has goodness $2^{\log^{1/3} p} / \log p$.
- ECDH has goodness $2^{p/2} / \log p$.

Problem 7. [6 points] **TLS security** (1 part)

Ben Bitdiddle is designing an Android application where users can send money to each other, by username. The application relies on a central server, which authenticates requests from user devices. When the application wants to transfer some amount of money to another user, it opens a TLS connection to the server, and sends the following message:

```
USER: username
PASS: password
REQUEST: transfer
AMOUNT: amount
RECIPIENT: recipient
```

where `username` and `password` authenticate the sender, and the request asks the server to transfer `amount` to `recipient`'s account.

Explain how a network adversary may be able to redirect a transfer to their account. You can assume the adversary knows the user, the fact that the user is transferring money, how much they are transferring, and to whom, but does not know the user's password. Assume that the TLS certificates are correct (i.e., certificate authorities will not issue an incorrect certificate) and that the adversary cannot guess the user's password.

Solution: Create an account that corresponds to some prefix of the recipient's user name. Terminate the TLS connection after that prefix of the recipient's user name is sent. Server will receive the request containing the truncated recipient name and will perform a transfer to that account instead of the intended recipient.

Problem 8. [6 points] **Law and technology** (1 part)

- (a) [6 points] Chris Conley's guest lecture outlined several U.S. acts that regulate federal computer law. Imagine that you have mounted one of the attacks from the 6.1600 labs without permission against a victim server elsewhere on the Internet. Name a U.S. law and an attack from either lab0, lab1, or lab2 that could violate that law. Give a one-sentence explanation for why your chosen attack would violate that law.

Solution:

Performing the lab0 attack to retrieve user's passwords from a hashed list and then accessing said users account and personal information would be a violation of CFAA because it involves unauthorized access of a protected computer. This attack could also be in violation of DMCA if it leads to unauthorized access of copyrighted works.

Performing lab1 attacks to tamper with user's communication with a protected store could violate CFAA because you if you tamper with "gates-up" data via performed unauthorized data manipulation to the Merkle tree server it could be argued you are "recklessly causing damage" to the store.

Performing lab2 attacks to tamper with encrypted wifi packets could violate CFAA as it performs unauthorized data as it "exceeds authorized access" to the server the client is communicating with.

What this question is looking for:

Knowledge of the one of the two main acts Chris Conley discussed. Either: DMCA CFAA Expresses a solid explanation of what DMCA and/or CFAA does via how an attack from lab0,1,2 could be preformed illegally.

Problem 9. [3 points] **Course survey** (3 parts)

Please answer each of these questions in one or two sentences.

- (a) [1 point] What lab assignment or lecture should we get rid of next year, in your opinion?

Solution: 7x Lab 0 (2x tedious; 1x didn't help learn; 1x brute-forcing is obvious). 7x Lab 1 (1x not relevant; 2x too difficult; 2x was nonetheless useful). 6x Lab 2 (1x problem 3d was confusing; 1x too hard; 1x too much math; 1x nerve-wracking true-false). 6x RSA lecture (2x outdated; 1x teach with extended Euclid; 1x too much material). 5x Encryption in practice (1x confusing; 1x not interesting; 1x replace with guest lecture about this topic). 3x Open problems in encryption (1x depressing; 1x replace with guest lecture about this topic). 2x Intro lecture (1x repetitive examples). 2x PKI lecture (1x CAs in particular). 2x Guest lecture (1x disconnected from class). 1x Lamport signatures. 1x Tor. 1x Authenticated encryption.

1x Too much recapping of previous lectures. 1x Want a review/recap lecture. 1x Combine encryption in practice with open questions. 1x Combine CPA and CCA lectures.

- (b) [1 point] Did you think that the lectures so far have been too fast or too slow?

Solution: 17x Too fast (1x not going deep enough into any one scheme; 2x recitations helped; 3x especially math/number theory; 1x recordings help to pause). 30x Just right. 4x Too slow.

- (c) [1 point] What is one topic (on the theme of the course) that you would like to learn more about?

Solution: 11x Elliptic-curve crypto. 8x Applications of key exchange / encryption / authentication. 6x Real-world outside-of-theory mistakes/attacks and how they were addressed. 5x Quantum attacks. 4x Network / Internet security. 2x RSA math, history. 2x TLS. 2x Malware. 1x Legalities of security research. 1x SSH. 1x PIR. 1x Biometrics. 1x Merkle tree applications. 1x Secure messaging. 1x Single sign-on. 1x Denial-of-service. 1x Blockchain. 1x Social engineering. 1x Problems beyond encryption. 1x System security. 1x Hardware security. 1x Side-channel attacks. 1x Tor.
1x Start each lecture with real-world motivation/story.