

Final

Question	Parts	Points
1: Instructions	1	0
2: True or False	10	20
3: User authentication	3	15
4: Symmetric Encryption	2	10
5: Public-Key Encryption	1	10
6: Compress then encrypt	3	15
7: Certificate revocation	2	10
8: Access control	1	5
9: Secure boot	1	5
10: LPN Error Correction	1	5
11: Isolation	2	10
12: Fuzzing	2	10
13: Symbolic execution	2	15
14: Sandbox Sharing	3	15
15: Differential Privacy	1	5
16: Timing Side Channel	4	20
17: Course Survey	4	10
Total:		180

Name: _____

Problem 1. [0 points] **Instructions** (1 part)

- This is an open book exam: you can use your notes from this class, or any material released by us this term. You cannot use the internet. Use of any material not released by us this term is *strictly* forbidden.
- Any form of collaboration is *strictly* forbidden.
- If you need assistance clarifying a question in the exam, raise your hand and a proctor will come by.
- Point totals correspond roughly to how much time we expect you to spend on each problem (part).

Problem 2. [20 points] **True or False** (10 parts)

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

- (a) [2 points] One way to construct a CPA-secure encryption scheme is with the “hash-and-encrypt” paradigm.

- (b) [2 points] The RSA signature scheme is post-quantum secure.

- (c) [2 points] One way to construct a secure signature scheme for arbitrary length messages is with the “hash-and-sign” paradigm.

- (d) [2 points] Transport-layer security is the protocol that underlies HTTPS.

- (e) [2 points] A CCA-secure encryption scheme must also be CPA-secure.

- (f) [2 points] If a client and server communicate using transport layer security (TLS), an eavesdropper in the middle of the network can learn the IP addresses of the client and server.

- (g) [2 points] A CCA-secure encryption scheme hides the length of the plaintext.

- (h) [2 points] It is impossible to update the lowest level Boot Read-Only Memory (ROM) in an iOS device.

- (i) [2 points] If you live in Belgium and connect to the Internet via a U.S.-based VPN provider, your Belgian ISP learns no information about your browsing behavior.

- (j) [2 points] You connect to a Google server via an encrypted TLS connection. Your ISP can distinguish between the cases in which (a) you watch a 30-minute episode of Stephen Colbert on YouTube and (b) you watch a 30-second cat video.

Problem 4. [10 points] **Symmetric Encryption** (2 parts)

- (a) [5 points] Construct a CPA secure symmetric encryption scheme for messages in $\{0, 1\}^{256}$ given a CPA secure symmetric encryption scheme for messages in $\{0, 1\}^{128}$.
- (b) [5 points] Show how to upgrade the security to CCA security given a MAC scheme for messages of arbitrary length.

Problem 5. [10 points] **Public-Key Encryption** (1 part)

Consider a 2-message key exchange protocol, where party 1 chooses randomness r_1 and sends a message m_1 , which is a function of r_1 (i.e., $m_1 = m_1(r_1)$), party 2 chooses randomness r_2 and replies with a message m_2 , which is a function of m_1 and r_2 (i.e., $m_2 = m_2(r_2, m_1)$), and the shared secret s can be efficiently computed from (m_1, r_2) or (m_2, r_1) .

Construct a CPA secure public-key encryption scheme from the key exchange protocol above, where the *weak* security guarantee of the key-exchange protocol is that given (m_1, m_2) generated as above, it is hard to compute the shared secret s (but s is not necessarily indistinguishable from uniform given (m_1, m_2)).

You can assume (for simplicity) that the secret key s is a binary string of length k (for some k), and that there is a secure hash function $H : \{0, 1\}^k \rightarrow \{0, 1\}^{k'}$ (modelled as a Random Oracle), and construct an encryption scheme for messages in $\{0, 1\}^{k'}$.

Problem 6. [15 points] **Compress then encrypt** (3 parts)

You browse to `evil.com` while using an evil WiFi network, so the attacker can run JavaScript in your browser *and* observe your network traffic.

While you are on this evil website, the attacker's JavaScript causes your browser to load a sequence of URLs from Google. That is, the attacker can cause your browser to send the following type of GET request to Google.

```
GET /ATTACKER_CHOOSSES_THIS_URL HTTP/1.1
Host: google.com
Accept: */*
Cookie: SECRET_VALUE_HERE
```

The Google cookie value is a secret 32-byte hexadecimal string—e.g., `1FAB8AEEC9...9809D9BC02`. (Each hexadecimal character in ASCII takes one byte to represent.)

The GET request travels to Google over an encrypted TLS connection. Recall that TLS encryption increases the length of the plaintext by a fixed amount.

Say that your browser *compresses* the GET request string before encrypting it. The compression algorithm has the property that:

- If there is no repeated 4-byte string in the GET request, the output is L bytes long.
 - If there are r repeated 4-byte strings in the GET request, the output is $L - r$ bytes long.
- (a) [5 points] Show that if the attacker can guess the first 4 bytes of your Google cookie value, it can confirm that its guess was correct.

(b) [5 points] Show that if the attacker can cause your browser to load 2^{16} URLs, it can recover the first 4 bytes of your Google cookie.

(c) [5 points] Show how the attacker can recover your entire Google cookie by loading $\ll 2^{20}$ URLs. You may assume that the cookie itself contains no repeated 4-byte strings.

Problem 7. [10 points] **Certificate revocation** (2 parts)

This problem deals with certificate revocation in the public-key infrastructure.

- (a) [5 points] The *online certificate status protocol* (OCSP) works as follows: when your browser gets a public-key certificate from a server, the certificate contains the URL of an OCSP server, typically run by the certificate authority who issued the certificate. The client then asks the OCSP server “Is this certificate `<hash of cert>` still valid?” The OCSP server responds with a signed message indicating YES or NO.

Unfortunately, the OCSP server learns which websites the client is visiting. A more recent technology, called *OCSP stapling* has the web server (e.g., `nytimes.com`) fetch a signed OCSP response from its certificate authority indicating that its certificate is still valid. This eliminates the privacy concern from the original scheme, but requires the web server to periodically fetch a new signed attestation from the certificate authority. The OCSP response includes a timestamp and clients will reject the response if it is too old.

Yet another concern with OCSP that it adds latency to the client’s web browsing experience, since the client must contact the OCSP server before it completes its connection to the web server.

To reduce client latency, one option would be to first download the webpage from the web server (e.g., `nytimes.com`) and then use OCSP to check the validity of the certificate. If the certificate is valid, the browser would tear down the browser window and display an error.

What security problems could arise from using OCSP like this?

- (b) [5 points] An alternative to OCSP is certificate revocation lists (CRLs). When using CRLs, the certificate authority includes in each certificate the URL of a CRL file. The client can fetch this signed CRL file, which contains a list of all certificates that the certificate authority has revoked. In this way, the client can check whether a particular certificate has been revoked. Explain how a malicious certificate authority could abuse this process to learn which clients are visiting which websites.

Problem 8. [5 points] **Access control** (1 part)

Ben Bitdiddle is designing a data storage system, and thinks capabilities are great because they allow for easy delegation. He uses capabilities as the basic access control primitive in his storage server, in lieu of traditional access control lists: clients must have a capability in order to fetch an object from his data store.

What problems might Ben run into by using only capabilities?

Problem 9. [5 points] **Secure boot** (1 part)

Alyssa P. Hacker is building a game console and wants to make sure that the game console hardware can only be used to run authentic games approved by her company. Alyssa wants to achieve this using secure boot, but wants to avoid checking many signatures during boot, and wants to avoid having to hard-code a public key at each step of the boot process. Alyssa's idea is to have each boot step store a hash of the next step in a designated memory location, and then the final boot loader will be responsible for checking a signature on all of these hashes together before jumping to the OS kernel.

Specifically, the boot ROM loads the boot loader from disk, computes a hash of the boot loader, h_{boot} and stores it in a designated memory location m_{boot} , before jumping to the boot loader. The boot loader then loads the OS kernel from disk, computes a hash h_{os} and stores it in designated memory location m_{os} . Then, before jumping to the OS kernel, the boot loader loads the signature from disk and checks that it's a valid signature, under the public key embedded in the boot loader, of $h_{boot} || h_{os}$. If the signature matches, the boot loader jumps to the OS kernel and proceeds booting. Otherwise, the boot stops.

Is Alyssa's scheme secure? Explain why or why not.

Problem 10. [5 points] **LPN Error Correction** (1 part)

Recall the LPN based error correction scheme (Lecture 17). n refers to the number of bits in the secret key s , m is the number of ring oscillator pairs and the number of bits in the helper data b and noise vector e .

The **Gen** step determines the b vector and exposes it. The **Rep** step chooses the n most stable bits (choosing the n out of m counter values that are the largest absolute values regardless of sign), and uses the corresponding n equations to solve for s .

Ben Bitdiddle (remember him from 6.004?) decides that he can retain his five customers by marketing his new scheme as providing **independent** secrets $s^{(k)}$ for each of his customers, k varying from 1 to 5. That is, he will use the same ring oscillator circuit array to “encode” different $s^{(k)}$ values into m -bit $b^{(k)}$ vectors in the **Gen** step. The **Rep** step is unchanged, given a $b^{(k)}$ vector the circuit regenerates $s^{(k)}$. Ben claims that his scheme is more general than the original scheme and has equivalent security. Is Ben correct? Provide a brief explanation.

Problem 11. [10 points] **Isolation** (2 parts)

- (a) [5 points] Ben Bitdiddle is using virtual machines for isolation. He notices that the cost of saving and restoring registers when switching between virtual machines is a significant source of overhead for his workload. Furthermore, he notices that a large part of that cost is due to saving large floating-point registers, even if the virtual machine didn't use any floating-point code.

Ben implements an optimization: instead of saving and restoring floating-point registers when switching between virtual machines, his virtual machine monitor lets the floating-point registers remain in-place. However, he configures the processor to trap into the virtual machine monitor if the virtual machine runs any instructions that access a floating-point register. At that point, his virtual machine monitor would finally save and restore the floating-point registers as it would have originally done, and enables the use of floating-point instructions.

Is this optimization secure with respect to the virtual machines providing isolation? You may assume non-speculative hardware processors. Explain why or why not.

- (b) [5 points] Ben allows virtual machines to dynamically manage their memory allocations. In particular, he adds special calls from the virtual machine into the virtual machine monitor that allow the VM to either allocate more pages of memory, or to give some pages of memory back to the shared global set of free pages. When allocating or freeing memory, the virtual machine monitor ensures that page tables are properly updated, so that at most one virtual machine has a page of memory mapped in its page tables at a time.

Does this design provide strong non-interference between virtual machines? Explain why or why not.

Problem 12. [10 points] **Fuzzing** (2 parts)

Consider the following C function being tested by a fuzzer that supplies random inputs in the 32-byte input array:

```
unsigned char input[32];

void f() {
    if (input[0] != input[1])
        return;

    if (input[2] < 128)
        return;

    if (input[3] > 32)
        crash();
}
```

- (a) [5 points] How many inputs, on average, will a purely random (not coverage-guided) fuzzer need to trigger the call to `crash()`?
- (b) [5 points] How many inputs, on average, will a coverage-guided fuzzer need to trigger the call to `crash()`?

Problem 13. [15 points] **Symbolic execution** (2 parts)

Consider the following C function being executed in a symbolic execution system with an arbitrary symbolic argument x :

```
void f(unsigned int x) {  
    if (x % 2 == 0) {  
        x = x + 1;  
    }  
  
    if (x > 4096) {  
        return;  
    }  
  
    if (x < 16) {  
        crash();  
    }  
}
```

(a) [10 points] Which of the following queries will a symbolic execution issue to its SAT solver? Check all that apply.

1. $((x \bmod 2) == 0)$
2. $\neg((x \bmod 2) == 0)$
3. $(x > 4096)$
4. $\neg(x + 1 > 4096)$
5. $(x + 1 > 4096) \wedge ((x \bmod 2) == 0)$
6. $(x + 1 > 4096) \wedge \neg((x \bmod 2) == 0)$
7. $(x < 16) \wedge (x > 4096) \wedge ((x \bmod 2) == 0)$
8. $(x < 16) \wedge \neg(x > 4096) \wedge \neg((x \bmod 2) == 0)$
9. $\neg(x + 1 < 16)$
10. $\neg(x < 16) \wedge (x > 4096) \wedge \neg((x \bmod 2) == 0)$

- (b) [5 points] How many execution paths will the symbolic execution engine explore in executing $f(x)$?

Problem 14. [15 points] **Sandbox Sharing** (3 parts)

Consider the following system where public keys identify users. Users own hardware authentication tokens running a Python host program. In this process is embedded a WebAssembly instance which owns the private key and signs messages on the user's behalf. The instance's source code is `auth.c` (implementation omitted):

```
static void uint8_t private_key[32];
void key_gen(uint8_t public_key[32]);
void hash(const uint8_t *message, uint16_t message_len,
          uint8_t hash_out[32]);
void sign(const uint8_t message_fixed[32], uint8_t signature[64]);
```

For each of the user's applications, the Python host embeds the application program in a separate sandboxed instance. `friends.c` is the source code for a distributed public key infrastructure application where users produce signed attestations of friendship (similar to that of lab 2):

```
void certify_friend(const uint8_t friend_public_key[32],
                   uint8_t signature[64]) {
    sign(friend_public_key, signature);
}
```

`payment.c` is the source code for an application where users sign transactions for a payment system:

```
void authorize_payment(const uint8_t *payment, uint16_t payment_len,
                      uint8_t signature[64]) {
    uint8_t hash_out[32];
    hash(payment, payment_len, hash_out);
    sign(hash_out, signature);
}
```

The programs above are compiled into `auth.wasm`, `friends.wasm`, and `payment.wasm` before they are instantiated, and the Python host allows the applications to call the `hash` and `sign` functions exported by `auth.wasm`.

Assume the following:

- `private_key` is securely confidential to `auth.wasm`.
- `private_key` is generated exactly once with `key_gen`.
- `hash` implements a collision-resistant hash function.
- `sign` implements a secure public-key signature scheme.

- (a) [5 points] How many bytes are copied from the `payment.wasm` sandbox to the `auth.wasm` sandbox memory by the Python host in a single call of `authorize_payment`? Explain your answer.
- (b) [5 points] Suppose that an attacker can make arbitrary calls to `authorize_payment`. Can the attacker certify a malicious public key as a friend? Explain why or why not.
- (c) [5 points] Suppose that an attacker can make arbitrary calls to `certify_friend`. Can the attacker authorize a malicious payment? Explain why or why not.

Problem 15. [5 points] **Differential Privacy** (1 part)

Let f be a function over sensitive data sets. The data sets D and D' both have n records. We will define the global sensitivity of f as:

$$S(f) = \max_{\forall \text{dist}(D, D')=1} |f(D) - f(D')|$$

where $\text{dist}(D, D') = 1$ if and only if D' can be obtained from D by changing *one* record in D .

Any possible record has a single value that is in between the minimum value a and maximum value b .

What is the global sensitivity when f is the median?

In order to get ϵ -differential privacy, what noise distribution should we add to the outcome $f(D)$?

Problem 16. [20 points] **Timing Side Channel** (4 parts)

Recall Lab 5 and the token comparison algorithm that was susceptible to timing attack. Bob now wants to use it to compare passwords.

Note that in this problem, you should ignore timing differences due to micro-architectural structures.

(a) [5 points] Bob thinks he has solved the timing issue. Assuming the real password is of length ≥ 1 , he hopes his algorithm respects the following properties:

- **Correctness:** The algorithm will return True if the password provided by the user is the exact same one as expected.
- **Security:** The algorithm will return False if the password is not the one expected.
- **Timing Independence:** The execution timing of the algorithm is independent of the real password.

Look at the following code snippet.

```
def check_password(p, real_p):
    if len(p) != len(real_p):
        return False
    result = True
    for i in range(len(p)):
        result = result and (p[i] == real_p[i])
    return result
```

Can an attacker recover any sensitive information regarding the real password using a timing attack? Can this have an impact on security?

(b) [5 points] Bob decides to write another version:

```
def check_password(p, real_p):
    max_idx = len(real_p) - 1
    result = True
    for i in range(len(p)):
        j = min(i, max_idx)
        c = real_p[j]
        result = result and (p[i] == c)
    return result
```

Can an attacker recover any sensitive information regarding the real password using a timing attack? Is this code secure as defined in the previous question?

(c) [5 points] In the next questions, assume that `check_password` is secure and satisfies timing independence as defined earlier. Bob would like to use it to perform some access control into his secure server. He hopes his scheme respects the following properties:

- **Correctness:** The algorithm will return `True` if the given username correspond to an existing user on the server and the associated password is the exact same one as provided by the user.
- **Security:** The algorithm will return `False` if the username is not registered on that machine or the password is not the one associated with the given username.
- **Timing Independence:** The execution timing of the algorithm should not help a potential attacker to recover *any* information regarding other users of the server.

He writes the following code:

```
def has_access(username, password):
    if username not in user_password_dict:
        return False
    else:
        return check_password(password, user_password_dict[username])
```

Can an attacker recover any sensitive information regarding the secure server using a timing attack? Can this have an impact on security?

(d) [5 points] Bob decides to write another version:

```
def has_access(username, password):
    cond = username not in user_password_dict
    empty_password = '00000000'
    if cond:
        real_p = empty_password
    else:
        real_p = user_password_dict[username]
    return check_password(password, real_p)
```

This code has a series of security and timing issues. Detail one of them.

