

For questions with **circular bubbles**, you may select exactly *one* choice on Gradescope.

- Unselected option
- Only one selected option

For questions with **square checkboxes**, you may select *one* or more choices on Gradescope.

- You can select
- multiple squares

For questions with a **large box**, you need to write a short answer in the corresponding text box on Gradescope.

You have 170 minutes. There are 10 questions of varying credit (250 points total).

The exam is open note. You can use an unlimited number of handwritten cheat sheets, but you must work alone.

Clarifications will be posted at <https://cs161.org/clarifications>.

Q1 MANDATORY – Honor Code

(7 points)

Read the honor code on the Gradescope answer sheet and type your name. Failure to do so will result in a grade of 0 for this exam.

Q2 True/false

(56 points)

Each true/false is worth 2 points.

Q2.1 TRUE or FALSE: You should always use HMAC instead of any other MAC because HMAC has stronger integrity and authentication guarantees than any other MAC.

- TRUE FALSE

Q2.2 TRUE or FALSE: A MiTM during the Diffie-Hellman Key Exchange can force both parties to derive a shared key (that the MiTM doesn't necessarily know) that is different than the one they would've derived otherwise.

- TRUE FALSE

Q2.3 TRUE or FALSE: A MiTM during the Diffie-Hellman Key Exchange can force both parties to unknowingly derive different keys that the MiTM knows.

- TRUE FALSE

Q2.4 TRUE or FALSE: A MiTM during the Diffie-Hellman Key Exchange can force both parties to derive a set of pre-determined keys that the MiTM knows.

- TRUE FALSE

Q2.5 TRUE or FALSE: CSRF tokens are an effective defense against CSRF attacks only if clients' browsers respect the same-origin policy.

- TRUE FALSE

Q2.6 TRUE or FALSE: An XSS vulnerability in a website cannot be exploited to gain control over a user's session if the session cookie has the HttpOnly flag set.

- TRUE FALSE

Q2.7 TRUE or FALSE: `https://secure.bank.com` is able to set the following cookie using the Set-Cookie header: `session=1234567; Domain=bank.com; HttpOnly`.

- TRUE FALSE

Q2.8 TRUE or FALSE: A user wants their web traffic to appear like it's coming from somewhere else with the lowest latency possible. This user should prefer a VPN instead of Tor.

- TRUE FALSE

Q2.9 TRUE or FALSE: In Bitcoin, once a transaction is successfully added to the blockchain, it can never be lost.

- TRUE FALSE

Q2.10 When you log in to Zoom, you make a POST Request to `https://zoom.us/berkeley/signin` with an email and password in the form data. The Response contains a session token cookie **without** the Secure flag set.

TRUE or FALSE: An on-path attacker could steal your session token by observing only this request.

TRUE FALSE

Q2.11 When you go to `https://berkeley.zoom.us/m/stanford`, you see an image of Stanford's lawn. The page source shows that the image is being loaded from `http://stanford.zoom.us/i/stanford.png`.

TRUE or FALSE: This a violation of the same-origin policy.

TRUE FALSE

Q2.12 You're using Tor with three intermediate nodes. Assume all nodes are handling a large amount of traffic.

TRUE or FALSE: Even if two of those nodes are compromised, your anonymity is still protected.

[*Clarification during exam:* This question was thrown out during the exam, and both True and False were accepted as valid answers. See solution for why.]

TRUE FALSE

Q2.13 Instead of using Tor, you forward your traffic through three intermediate proxies **unencrypted**. Using these proxies, you log into `https://twitter.com`

TRUE or FALSE: Assuming the entry proxy is honest, the middle and exit proxies cannot figure out your identity

TRUE FALSE

Q2.14 You decide to use a recursive resolver which uses DNSSEC. Your client uses standard DNS.

TRUE or FALSE: An on-path adversary cannot poison your client's cached DNS records.

TRUE FALSE

Q2.15 A recursive resolver supports DNSSEC. The resolver contacts three other nameservers to answer a certain query.

TRUE or FALSE: All three nameservers must support DNSSEC in order for DNSSEC to provide any guarantees.

TRUE FALSE

Q2.16 TRUE or FALSE: DHCP is secure against an on-path attacker.

TRUE FALSE

Q2.17 TRUE or FALSE: Using HTTPS is a good defense against clickjacking attacks.

TRUE FALSE

Q2.18 TRUE or FALSE: Spearphishing is more dangerous than standard phishing because it uses information about the victim.

- TRUE FALSE

Q2.19 TRUE or FALSE: If a website only allows HTTPS connections, it is secure from SQL injection attacks.

- TRUE FALSE

Q2.20 TRUE or FALSE: Parameterized SQL stops all SQL injection attacks.

- TRUE FALSE

Q2.21 Consider a website which inserts user input into a database using a SQL query. The information in the database is then used in subsequent internal SQL queries.

TRUE or FALSE: If the SQL query that accepts user input is parameterized, but the internal ones do not, then the website will be secure from SQL injection attacks.

- TRUE FALSE

Q2.22 TRUE or FALSE: Return-oriented programming (ROP) is not effective if non-executable pages (DEP or WX) are enabled.

- TRUE FALSE

Q2.23 TRUE or FALSE: Format string vulnerabilities are not effective if ASLR is enabled.

- TRUE FALSE

Suppose you find a stored XSS vulnerability on `https://berkeley.zoom.us/m/1234`.

Q2.24 TRUE or FALSE: Some cookies set by `https://berkeley.zoom.us/` could be **read** using your exploit.

- TRUE FALSE

Q2.25 TRUE or FALSE: Some cookies set by `https://berkeley.zoom.us/` could be **modified** using your exploit.

- TRUE FALSE

Q2.26 TRUE or FALSE: Some cookies set by `http://zoom.berkeley.edu/m/1234` could be **read** using your exploit.

- TRUE FALSE

Q2.27 TRUE or FALSE: Some cookies set by `https://berkeley.zoom.us/m/1234` could be **modified** using your exploit.

- TRUE FALSE

Q2.28 TRUE or FALSE: Some cookies set by `http://stanford.zoom.us/m/1234` could be **read** using your exploit.

TRUE

FALSE

This is the end of Q2. Proceed to Q3 on your answer sheet.

Q3 Password Storage

(28 points)

Bob is trying out different methods to securely store users' login passwords for his website.

Mallory is an attacker who can do some amount of *offline* computation before she steals the passwords file, and some amount of *online* computation after stealing the passwords file.

Technical details:

- Each user has a unique username, but several users may have the same password.
- Mallory knows the list of users registered on Bob's site.
- Bob has at most 500 users using his website with passwords between 8–12 letters.
- Mallory's dictionary contains all words that are less than 13 letters. [*Clarification during exam:* Mallory's dictionary contains all possible user passwords.]
- Mallory can do N online computations and $500N$ offline computations where N is the number of words in the dictionary.
- Slow hash functions take 500 computations per hash while fast hash functions require only 1 computation.¹

Notation:

- H_S and H_F , a slow and fast hash function
- Sign, a secure signing algorithm
- `uname` and `pwd`, a user's username and password
- k , a signing key known only by Bob

If Bob decides to use signatures in his scheme, assume he will verify them when processing a log-in.

Q3.1 (2 points) How many times could Mallory hash every word in the dictionary using H_S with **offline computation**?

- | | |
|---|---|
| <input type="radio"/> (A) She can't hash the whole dictionary | <input type="radio"/> (D) None of the above |
| <input type="radio"/> (B) 1 | <input type="radio"/> (E) — |
| <input type="radio"/> (C) 500 | <input type="radio"/> (F) — |

Q3.2 (2 points) How many times could Mallory hash every word in the dictionary using H_F with **online computation**?

- | | |
|---|---|
| <input type="radio"/> (G) She can't hash the whole dictionary | <input type="radio"/> (J) None of the above |
| <input type="radio"/> (H) 1 | <input type="radio"/> (K) — |
| <input type="radio"/> (I) 500 | <input type="radio"/> (L) — |

Q3.3 (2 points) How many times could Mallory hash every word in the dictionary using H_S with **online computation**?

¹Keep in mind this is much faster than a real-life slow hash function.

- (A) She can't hash the whole dictionary (D) None of the above
 (B) 1 (E) —
 (C) 500 (F) —

For each part below, indicate all of the things Mallory can do given the password storage scheme. Assume Mallory knows each scheme. **Unless otherwise specified, assume that she can use both offline and online computation**

Q3.4 (4 points) Each user's password is stored as $H_F(\text{pwd} \parallel \text{'Bob'})$.

- (G) Learn whether two users have the same password with only online computation (J) Learn every user's password
 (H) Learn a specific user's password (K) None of the above
 (I) Change a user's password without detection (L) —

Q3.5 (4 points) Each user's password is stored as the tuple $(H_S(\text{pwd} \parallel \text{'Bob'}), \text{Sign}(k, H_F(\text{pwd})))$.

- (A) Learn whether two users have the same password with only online computation (D) Learn every user's password
 (B) Learn a specific user's password (E) None of the above
 (C) Change a user's password without detection (F) —

Q3.6 (4 points) Each user's password is stored as the tuple $(H_F(\text{pwd} \parallel \text{uname}), \text{Sign}(k, \text{uname} \parallel H_F(\text{pwd})))$

- (G) Learn whether two users have the same password with only online computation (J) Learn every user's password
 (H) Learn a specific user's password (K) None of the above
 (I) Change a user's password without detection (L) —

Q3.7 (4 points) Each user's password is stored as $(H_S(\text{pwd} \parallel \text{uname}), \text{Sign}(k, H_S(\text{pwd})))$

[Clarification during exam: The expression was missing a leading parenthesis.]

- (A) Learn whether two users have the same password with only online computation (B) Learn a specific user's password
 (C) Change a user's password without detection

tion

(E) None of the above

(D) Learn every user's password

(F) —

Q3.8 (3 points) Describe a DoS attack Mallory can launch against Bob's server if he uses the scheme in Q3.7.

Q3.9 (3 points) Bob decides to add two-factor authentication to the scheme in Q3.7. Does this change your answer to Q3.7?

(A) Yes

(B) No

(C) —

(D) —

(E) —

(F) —


This is the end of Q3. Proceed to Q4 on your answer sheet.

Q4 Forwards, Backwards, Left, and Right


(16 points)

Consider the following properties. The solid part of each timeline denotes the time frame where messages remain confidential, even after Eve, an on-path eavesdropper, steals a key.


- *Forward secrecy*: If Eve steals a key, past messages remain confidential.




- *Backward secrecy*: If Eve steals a key, future messages remain confidential.



- *Weak forward secrecy*: If Eve stops recording messages, then steals a key, any messages Eve recorded before she stopped recording remain confidential.



- *Weak backward secrecy*: If Eve steals a key, then starts recording messages, any messages Eve record remain confidential.



Consider the following modified symmetric encryption schemes where Alice and Bob change their encryption key for each message they send. For each scheme, determine which of the given properties is ensured. Assume that all keys are 128 bits long, and no party will send more than one message in a row.

Q4.1 (4 points) Alice and Bob increment their shared key k by 1 for each new message, so $k' = k + 1$.

- | | |
|---|--|
| <input type="checkbox"/> (A) Forward secrecy | <input type="checkbox"/> (D) Weak backward secrecy |
| <input type="checkbox"/> (B) Backward secrecy | <input type="checkbox"/> (E) None of the above |
| <input type="checkbox"/> (C) Weak forward secrecy | <input type="checkbox"/> (F) — |

Q4.2 (4 points) Alice and Bob's current shared key is k . For each new message, the sender generates a small, 8-bit random number n and attaches it to the message before encryption. The next message will be encrypted under key $k' = k \oplus \text{PRG}(n)[: 128]$, where PRG is a secure PRG.

(G) Forward secrecy

(J) Weak backward secrecy

(H) Backward secrecy

(K) None of the above

(I) Weak forward secrecy

(L) —

Q4.3 (4 points) Alice and Bob's current shared key is k . For each new message, the sender generates a new symmetric key k' and attaches it to the message before encryption. The next message will be encrypted under k' .

(A) Forward secrecy

(D) Weak backward secrecy

(B) Backward secrecy

(E) None of the above

(C) Weak forward secrecy

(F) —

Q4.4 (4 points) For each new message, Alice and Bob conduct Diffie-Hellman key exchange to generate a new symmetric key.

(G) Forward secrecy

(J) Weak backward secrecy

(H) Backward secrecy

(K) None of the above

(I) Weak forward secrecy

(L) —

This is the end of Q4. Proceed to Q5 on your answer sheet.

Q5 *EvanBotOS*

(25 points)

EvanBot is building a new OS and wants to defend against buffer overflow attacks. Bot decides to use cryptography to secure values on the stack.

Assume any cryptography is executed separately and securely by the OS. This means that any cryptographic operations do not count as function calls on the program's stack, and the attacker cannot see the operations being executed. Also, unless otherwise stated, **any MACs or hashes generated are stored separately in the OS, not on the stack.**

Assume stack canaries are four random bytes (no null byte). Assume the OS has a secret key k that is unknown to any attacker.

For each part, mark which scheme is more secure (would defend against more buffer overflow attacks), or if both schemes would defend against the same set of attacks.

[Clarification during exam: For each scheme, unless otherwise specified all memory safety defenses are disabled.]

Q5.1 (3 points) Scheme A: When a function is called, push a random stack canary to the stack. Also, generate a MAC on the canary value using k . Before the function returns, in addition to checking that the canary is the same, also verify the canary with the MAC.

Scheme B: No cryptography, stack canaries are enabled, W^X and ASLR are disabled.

- | | | |
|------------------------------------|------------------------------------|-----------------------------|
| <input type="radio"/> (A) Scheme A | <input type="radio"/> (C) The same | <input type="radio"/> (E) — |
| <input type="radio"/> (B) Scheme B | <input type="radio"/> (D) — | <input type="radio"/> (F) — |

Q5.2 (3 points) Scheme A: When a function is called, encrypt a randomly-generated stack canary using k . Push the encrypted canary onto the stack. Before the function returns, decrypt the stack canary and verify that it is unchanged.

Scheme B: No cryptography, stack canaries are enabled, W^X and ASLR are disabled.

- | | | |
|------------------------------------|------------------------------------|-----------------------------|
| <input type="radio"/> (G) Scheme A | <input type="radio"/> (I) The same | <input type="radio"/> (K) — |
| <input type="radio"/> (H) Scheme B | <input type="radio"/> (J) — | <input type="radio"/> (L) — |

Q5.3 (3 points) Scheme A: When a program is first started, generate a signature on every page of the memory space using k . If the program tries to execute any instructions in memory, check that the page where the instruction is stored is correctly signed.

Scheme B: No cryptography, W^X is enabled, stack canaries and ASLR are disabled.

- | | | |
|------------------------------------|------------------------------------|-----------------------------|
| <input type="radio"/> (A) Scheme A | <input type="radio"/> (C) The same | <input type="radio"/> (E) — |
| <input type="radio"/> (B) Scheme B | <input type="radio"/> (D) — | <input type="radio"/> (F) — |

Q5.4 (3 points) Scheme A: When a function is called, using a cryptographic hash H , hash the RIP, *and push the value of the hash onto the stack*. Before the function returns, verify that the RIP still hashes to the same value.

Scheme B: When a function is called, generate a MAC on the RIP using k , *and push the value of the MAC onto the stack*. Before the function returns, verify the RIP with the MAC.

Assume that the hash and the MAC are the same length.

- (G) Scheme A (I) The same (K) —
 (H) Scheme B (J) — (L) —

Q5.5 (5 points) Consider Scheme A from the previous part. Briefly explain how you might create an exploit for Scheme A that overwrites the RIP. Assume you can debug only the vulnerable program with GDB, and you cannot access the OS-level cryptography operations.

- (A) — (B) — (C) — (D) — (E) — (F) —

Q5.6 (3 points) Scheme A: When a function is called, encrypt the RIP with a one-time pad, where the pad is a static value stored in the OS. (The pad value does not change when you rerun the program.) Before the function returns, decrypt the RIP and jump to that location.

Scheme B: No cryptography, stack canaries are enabled, W^X and ASLR are disabled.

- (G) Scheme A (I) The same (K) —
 (H) Scheme B (J) — (L) —

Q5.7 (5 points) Consider Scheme A from the previous part. In 2-3 sentences, explain how you might create an exploit for Scheme A that overwrites the RIP. Assume you can debug only the vulnerable program with GDB, and you cannot access the OS-level cryptography operations.

- (A) — (B) — (C) — (D) — (E) — (F) —

This is the end of Q5. Proceed to Q6 on your answer sheet.

Q6 DNS over TCP

(20 points)

Standard DNS uses UDP to send all queries and responses. Consider a modified DNS that instead uses TCP for all queries and responses.

Q6.1 (3 points) Which of the following does DNS over TCP guarantee against a man-in-the-middle attacker? Select all that apply.

- (A) Confidentiality (C) Authenticity (E) —
 (B) Integrity (D) None of the above (F) —

Q6.2 (3 points) Compared to standard DNS, does DNS over TCP defend against more attacks, fewer attacks, or the same amount of attacks against an on-path attacker?

- (G) More attacks (I) Fewer attacks (K) —
 (H) Same amount of attacks (J) — (L) —

Q6.3 (5 points) What fields does an off-path attacker **not know** and need to **guess** correctly to spoof a response in DNS over TCP? Assume source port randomization is enabled. Select all that apply.

- (A) TCP sequence numbers (C) Recursive resolver port (E) DNS NS records
 (B) Name server port (D) DNS A records (F) None of the above

Q6.4 (3 points) Is the Kaminsky attack possible on DNS over TCP? Assume source port randomization is disabled.

- (G) Yes, because the attacker only needs to guess the DNS Query ID
 (H) Yes, but we consider it infeasible for modern attackers
 (I) No, because the attacker cannot force the victim to generate a lot of DNS over TCP requests
 (J) No, because TCP has integrity guarantees
 (K) —
 (L) —
 (M) —

Q6.5 (3 points) Recall the DoS amplification attack using standard DNS packets. An off-path attacker spoofs many DNS queries with the victim's IP, and the victim is overwhelmed with DNS responses.

Does this attack still work on DNS over TCP?

- (A) Yes, the attack causes the victim to consume more bandwidth than the standard DNS attack

- (B) Yes, the attack causes the victim to consume less bandwidth than the standard DNS attack
- (C) No, because the DNS responses no longer provide enough amplification
- (D) No, because the attacker cannot force the server to send DNS responses to the victim
- (E) —
- (F) —

Q6.6 (3 points) What type of off-path DoS attack from lecture is DNS over TCP vulnerable to, but standard DNS not vulnerable to? Answer in five words or fewer.

Q7 I(T)C(P) You

(26 points)

EvanBot builds a new course feature that sends announcements to students over TCP. To receive announcements, a student initiates a TCP connection with the server. The server sends the announcements and terminates the connection.

Q7.1 (3 points) Assuming that no adversaries are present, which of the following does communication over a TCP connection guarantee? Select all that apply.

- (A) That both the server and client can detect if a particular announcement needs to be resent
- (B) That different announcements are delivered in the same order they were sent in
- (C) That announcements are delivered using the most efficient path through the internet
- (D) None of the above
- (E) —
- (F) —

Q7.2 (3 points) When only an on-path adversary is present, which of the following does communication over a TCP connection guarantee? Select all that apply.


- (G) That both the server and client can detect if a particular announcement needs to be resent
- (H) That different announcements are delivered in the same order they were sent in
- (I) That announcements are delivered using the most efficient path through the internet
- (J) None of the above
- (K) —
- (L) —

Q7.3 (3 points) Suppose that EvanBot instead sends announcements over UDP. Assuming that no adversaries are present, which of the following might happen? Select all that apply.

- (A) Students might not receive some announcements
- (B) Students might receive the announcements more quickly
- (C) The server might not detect some errors which it would have had it been using TCP
- (D) None of the above
- (E) —

EvanBot realizes that the server is sending messages to the student, but the student only responds with ACKs and never sends any messages after the initial handshake. They design a *Half TCP* protocol which provides TCP's properties for communications from the server to the student, but not for

communications from the student to the server. This is accomplished using a modified version of the standard three step handshake pictured below.



Q7.4 (5 points) Some sequence numbers are no longer necessary in *Half TCP*. Which fields **do not** need to be transmitted? Select all that apply.

- (G) The sequence number in the SYN packet
- (H) The sequence number in the SYN-ACK packet
- (I) The ACK number in the SYN-ACK packet
- (J) The sequence number in the ACK packet
- (K) The ACK number in the ACK packet
- (L) None of the above

Q7.5 (3 points) Which of these are consequences of moving from TCP to *Half TCP* for this application? Select all that apply.

- (A) The student will no longer receive announcements in the correct order
- (B) The server will not have to keep track of as much state
- (C) The student will not have to keep track of as much state
- (D) None of the above
- (E) —
- (F) —

The 161 staff likes security and decides to use TLS over *Half TCP*. Assume that the staff server has a valid certificate for their public key.

For each different adversary below, select all attacks which become *easier* when running TLS over *Half TCP* compared to normal TCP.

Q7.6 (3 points) Off-path adversary

- (G) RST Injection Attack
- (H) Interfere with a TLS handshake to learn the master key
- (I) Replay an encrypted command from a previous TLS connection

(J) None of the above

(K) —

(L) —

Q7.7 (3 points) On-path adversary

(A) RST Injection Attack

(B) Interfere with a TLS handshake to learn the master key

(C) Replay an encrypted command from a previous TLS connection

(D) None of the above

(E) —

(F) —

Q7.8 (3 points) Man-in-the-middle adversary

(G) RST Injection Attack

(H) Interfere with a TLS handshake to learn the master key

(I) Replay an encrypted command from a previous TLS connection

(J) None of the above

(K) —

(L) —

This is the end of Q7. Proceed to Q8 on your answer sheet.

Q8 Election Security

(23 points)

The 2020 elections are coming up, and the United States Government has tasked you with securing the nation's voting machines!

Assume election headquarters are in a top-secret, undisclosed site. All incoming network requests pass through a network-based intrusion detection system (NIDS), as well as a firewall. Outside users can only access the server with HTTPS.

Q8.1 (3 points) Which of these attacks are **always** preventable in this setup? Assume the attacker is on-path. Select all that apply.

- (A) RST Injection Attack
- (B) SQL Injection Attack
- (C) Reflected XSS Attack
- (D) None of the Above
- (E) —
- (F) —

Q8.2 (3 points) Which of these attacks are **always** preventable in this setup? Assume the attacker is on-path. Select all that apply.

- (G) SYN Flooding Attack
- (H) DNS Spoofing Attack
- (I) DDoS Attack
- (J) None of the Above
- (K) —
- (L) —

Q8.3 (3 points) An attacker injects malicious code on a server inside the election headquarters that changes all submitted votes to one candidate. Which detection system is best suited to defend against this attacker?

- (A) HIDS
- (B) NIDS
- (C) Firewall
- (D) —
- (E) —
- (F) —

Q8.4 (3 points) An attacker realizes that the ballot boxes are running a vulnerable version of Linux, and uses a previously-known buffer overflow exploit. Which detection method is best suited to defend against this attacker?

- (G) Anomaly-Based Detection
- (H) Signature-Based Detection
- (I) Specification-Based Detection
- (J) Behavioral-Based Detection
- (K) —
- (L) —

Q8.5 (5 points) Ben, a computer scientist at the top-secret site, has a HIDS installed on his work laptop. He decides to sign into his personal email account, claiming that HTTPS will protect the government from seeing his emails. Is he correct? Justify your answer in 1–2 sentences.

(A) Yes

(D) —

(B) No

(E) —

(C) —

Q8.6 (3 points) You've discovered that an attacker has managed to connect to a service running inside our network from IP Address 5.6.7.8 and is in the process of performing a DoS attack! Write a stateful firewall rule to block all traffic originating from the attacker. Our service is running on IP address 1.2.3.4 (port 443).

Q8.7 (3 points) You've received a tip that attackers have devised a plan to spoof ballot submissions. Here's the information that your source provides:

- 20 out of every 100 submissions are malicious.
- The cost to investigate an incorrectly flagged submission is \$5.
- The cost of letting a spoofed submission through is \$50.

You're offered two different intrusion detection systems. System A offers a false positive rate of 10% and a false negative rate of 25%. System B offers a false positive rate of 50% and a false negative rate of 5%. Which do you choose?

(A) System A

(D) Either system

(B) System B

(E) —

(C) Not enough information

(F) —

This is the end of Q8. Proceed to Q9 on your answer sheet.

Q9 Cookie Debugger**(37 points)**

EvanBot is adding a feature on the CS161 course website that lets students log in and view their grades. However, Bot forgot to remove a debugging feature—if anyone visits `cs161.org/debug`, the webpage will display all the cookies sent to the server.

Assume the `cs161.org/debug` page does not have any other functionality. Assume anyone can create an account on the website. Each subpart is independent.

Q9.1 (3 points) Which of the following URLs have the same origin as `http://cs161.org/debug` according to the same-origin policy?

- (A) `http://cs161.org/`
- (B) `http://cs161.org:8081/debug`
- (C) `https://cs161.org/debug`
- (D) None of the above
- (E) —
- (F) —

Q9.2 (5 points) Which of the following cookies would be displayed when visiting `https://cs161.org/debug`? Assume the client's origin is `https://cs161.org`.

- (G) Domain = `cs161.org`, Path = `/`, Secure
- (H) Domain = `cs161.org`, Path = `/`, HttpOnly
- (I) Domain = `debug.cs161.org`, Path = `/`, Secure, HttpOnly
- (J) Domain = `cs161.org`, Path = `/debug`
- (K) Domain = `cs161.org`, Path = `/`, SameSite=strict
- (L) None of the above

Q9.3 (3 points) Suppose you set a cookie `test=<script>alert("This exam is hard!")</script>` with valid attributes, and load `https://cs161.org/debug`. A pop-up that says `This exam is hard!` appears in your browser. Have you successfully found a server vulnerability?

[Clarification during exam: The pop-up had a typo in it.]

- (A) Yes, you found an XSS vulnerability
- (B) Yes, you found a CSRF vulnerability
- (C) No, because you have not changed any state on the server side
- (D) No, because the JavaScript does not run with the origin of `cs161.org`
- (E) —
- (F) —

Q9.4 (5 points) Consider a modification to the course website. Before rendering any page, the server searches for every pair of `<script>` and `</script>` tags and removes the tags *and everything between the tags*.

Can you still cause JavaScript to run in your browser using `<script>` tags? If yes, provide a cookie name and value (written as `name=value`) that would cause `alert(1)` to run. If no, briefly explain why.

- (G) Yes (H) No (I) — (J) — (K) — (L) —

Q9.5 (5 points) Consider a modification to the course website. Before rendering any page, the server renders the cookie name in an isolated environment and ensures that no scripts are run, and then does the same for the cookie value.

Assume that the website displays the cookie name and value with no added text in between. Can you still cause JavaScript to run in your browser using `<script>` tags? If yes, provide a cookie name and value (written as `name=value`) that would cause `alert(1)` to run. If no, briefly explain why.

- (A) Yes (B) No (C) — (D) — (E) — (F) —

Q9.6 (3 points) Is it possible to create a link to `cs161.org/debug` that will cause another user to run malicious JavaScript when they click on the link?

- (G) Yes, because you can place JavaScript in the HTTP GET parameters
- (H) Yes, because you can place JavaScript in the HTTP POST body
- (I) No, because there is nowhere to place the JavaScript
- (J) No, because the server is secure against this attack
- (K) —
- (L) —

Q9.7 (5 points) Suppose a victim visits the attacker-controlled `evil.cs161.org`. Write a JavaScript snippet that would cause the victim to run `alert(1)` in their browser with the origin of `cs161.org`. If you don't know the exact Javascript syntax, pseudo-code is acceptable.

Hint: `window.location = "google.com";` in JavaScript causes the user to load `google.com`.

Q9.8 (5 points) Which of the following malicious pages would be able to run your Javascript exploit against the user?

- | | |
|---|---|
| <input type="checkbox"/> (G) <code>http://very.evil.cs161.org/</code> | <input type="checkbox"/> (J) <code>http://cs161.org/evil</code> |
| <input type="checkbox"/> (H) <code>http://very-evil.cs161.org/</code> | <input type="checkbox"/> (K) <code>http://evil.com/</code> |
| <input type="checkbox"/> (I) <code>http://evil-cs161.org/</code> | <input type="checkbox"/> (L) None of the above |

Q9.9 (3 points) Consider a modification to the course website. The `cs161.org/debug` page only displays cookies if the request contains a valid session token. Does your Javascript exploit still work?

- (A) Yes, with no modifications
- (B) Yes, with minor modifications (changing 1-2 lines of code)
- (C) No
- (D) —
- (E) —
- (F) —

This is the end of Q9. Proceed to Q10 on your answer sheet.

Q10 Bitcoin**(12 points)**

Assume a simplified Bitcoin model, where each block contains the following fields:

- **minerID**: The public key of the node who mined this block. Recall that the person who mined a block is given a mining reward in Bitcoin. Assume that a miner can redeem this award by simply referencing the block ie. the initial award is *not* stored as a transaction.
- **prevHash**: The hash of the previous block
- **transactions**: The list of transactions. Recall each transaction contains references to its origin transactions, a list of recipients, and is signed using the private key of the coins' owner.
- **nonce**: A value such that the hash of the current block contains the correct number of zeros

Assume that the hash of a block is computed as:

$$\text{Hash}(\text{minerID} \parallel \text{prevHash} \parallel \text{transactions} \parallel \text{nonce})$$

Bob wants to save on computing power by omitting certain fields in a block from being part of the hash. For each modified block hashing scheme below, select all the things an adversary with a single standard CPU can do.

Assume that if the adversary can come up with a modified blockchain of the same length, the rest of the network will accept it. Furthermore, assume the adversary has not made any transactions thus far. **Any option that could result in an invalid state should not be selected.**

Q10.1 (4 points) Each block hash is computed as $\text{Hash}(\text{prevHash} \parallel \text{transactions} \parallel \text{nonce})$

- | | |
|---|---|
| <input type="checkbox"/> (A) Modify a block to gain Bitcoin | <input type="checkbox"/> (D) Can remove any transaction in an arbitrary block by <i>only</i> modifying that block |
| <input type="checkbox"/> (B) Given some amount of pre-computation, can consistently win proof of work | <input type="checkbox"/> (E) None of the above |
| <input type="checkbox"/> (C) Modify some transaction amounts | <input type="checkbox"/> (F) — |

Q10.2 (4 points) Each block hash is computed as $\text{Hash}(\text{minerID} \parallel \text{transactions} \parallel \text{nonce})$

- | | |
|---|---|
| <input type="checkbox"/> (G) Modify a block to gain Bitcoin | <input type="checkbox"/> (J) Can remove any transaction in an arbitrary block by <i>only</i> modifying that block |
| <input type="checkbox"/> (H) Given some amount of pre-computation, can consistently win proof of work | <input type="checkbox"/> (K) None of the above |
| <input type="checkbox"/> (I) Modify some transaction amounts | <input type="checkbox"/> (L) — |

Q10.3 (4 points) Each block hash is computed as $\text{Hash}(\text{minerID} \parallel \text{prevHash} \parallel \text{nonce})$

- | | |
|---|---|
| <input type="checkbox"/> (A) Modify a block to gain Bitcoin | <input type="checkbox"/> (D) Can remove any transaction in an arbitrary block by <i>only</i> modifying that block |
| <input type="checkbox"/> (B) Given some amount of pre-computation, can consistently win proof of work | <input type="checkbox"/> (E) None of the above |
| <input type="checkbox"/> (C) Modify some transaction amounts | <input type="checkbox"/> (F) — |

This is the end of Q10. Proceed to Q11 on your answer sheet.