Question 1 *Attack on TCP* ()
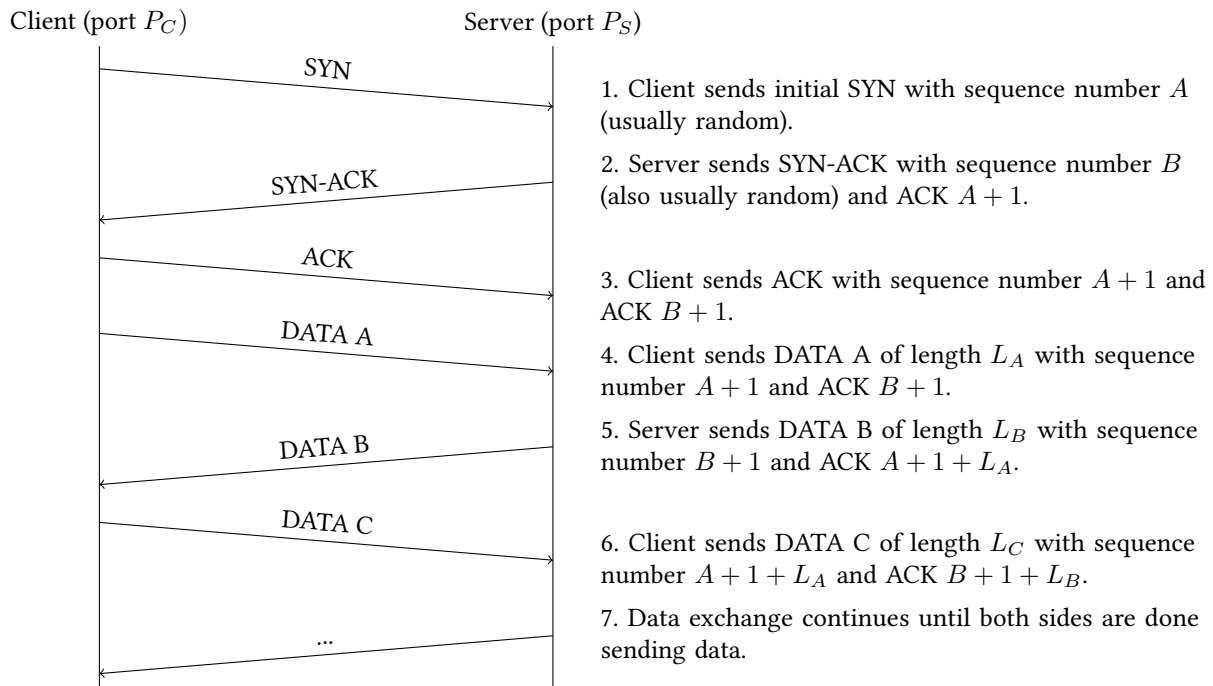
Suppose that a client connects to a server, and then performs the following TCP handshake and initial data transfer:

Client (port P_C)   Server (port P_S)

SYN →

1. Client sends initial SYN with sequence number A (usually random).

← SYN-ACK

2. Server sends SYN-ACK with sequence number B (also usually random) and ACK $A + 1$.

ACK →

3. Client sends ACK with sequence number $A + 1$ and ACK $B + 1$.

DATA A →

4. Client sends DATA A of length L_A with sequence number $A + 1$ and ACK $B + 1$.

← DATA B

5. Server sends DATA B of length L_B with sequence number $B + 1$ and ACK $A + 1 + L_A$.

DATA C →

6. Client sends DATA C of length L_C with sequence number $A + 1 + L_A$ and ACK $B + 1 + L_B$.

...

7. Data exchange continues until both sides are done sending data.

Q1.1 Assume that the next transmission in this connection will be DATA D from the server to the client. What will this packet look like?

Sequence number:	_____	ACK:	_____
Source port:	P_S	Destination port:	P_C
Length:	L_D	Flags:	ACK

Q1.2 You should be familiar with the concept and capabilities of a *man-in-the-middle* as an attacker who **can observe** and **can modify** traffic. There are two other types of relevant attackers in this scenario:

1. *On-path* attacker: **can observe** traffic but **cannot modify** it.

2. *Off-path* attacker: **cannot observe** traffic and **cannot modify** it.

Carol is an *on-path* attacker. Can Carol do anything malicious to the connection? If so, what can she do?

Q1.3 David is an *off-path* attacker. Can David do anything malicious to the connection? If so, what can he do?

Q1.4 The client starts getting responses from the server that don't make any sense. Inferring that David is attempting to hijack the connection, the client then immediately sends the server a **RST** packet, which terminates the ongoing connection. David wants to impersonate the client by establishing a new connection. How would he go about doing this?

Question 2 *TLS protocol details* ()

Depicted below is a typical instance of a TLS handshake.

Client   Server

1. ClientHello

                                                1. Client sends 256-bit random number R_b and supported ciphers

2. ServerHello

                                                2. Server sends 256-bit random number R_s and chosen cipher

3. Certificate

                                                3. Server sends certificate

4. ServerKeyExchange

                                                4. DH: Server sends $\{g, p, g^a \bmod p\}_{K_{\text{server}}^{-1}}$

5. ServerHelloDone

                                                5. Server signals end of handshake

6. ClientKeyExchange

                                                6. DH: Client sends $g^b \bmod p$

                                                   RSA: Client sends $\{PS\}_{K_{\text{server}}}$

Client and server derive cipher keys C_b, C_s and integrity keys I_b, I_s from R_b, R_s, PS

7. ChangeCipherSpec, Finished

                                                7. Client sends MAC(dialog, I_b)

8. ChangeCipherSpec, Finished

                                                8. Server sends MAC(dialog, I_s)

9. Application Data

                                                9. Client data takes the form $\{M_1, \text{MAC}(M_1, I_b)\}_{C_b}$

10. Application Data

                                              10. Server data takes the form $\{M_2, \text{MAC}(M_2, I_s)\}_{C_s}$

Figure 1: TLS 1.2 Key Exchange

Q2.1 What is the purpose of the *client random* and *server random* fields?

Q2.2 ClientHello and ServerHello are not encrypted or authenticated. Explain why a man-in-the-middle cannot exploit this. (Consider both the Diffie-Hellman and RSA case.)

Q2.3 Note that in the TLS protocol presented above, there are two cipher keys C_b and C_s. One key is used only by the client, and the other is used only by the server. Likewise, there are two integrity keys I_b and I_s. Alice proposes that both the server and the client should simply share one cipher key C and one integrity key I. Why might this be a bad idea?

Q2.4 The protocol given above is a simplified form of what actually happens. After step 8 (ChangeCipherSpec), the protocol as described above is still vulnerable. What is the vulnerability and how could you fix this?

Question 3   *TLS threats*   ()

An attacker is trying to attack the company Boogle and its users. Assume that users always visit Boogle's website with an HTTPS connection, using ephemeral Diffie-Hellman. You should also assume that Boogle does not use certificate pinning. The attacker may have one of three possible goals:

1. Impersonate the Boogle web server to a user

2. Discover some of the plaintext of data sent during a past connection between a user and Boogle's website

3. Replay data that a user previously sent to the Boogle server over a prior HTTPS connection

For each of the following scenarios, describe if and how the attacker can achieve each goal.

Q3.1 The attacker obtains a copy of Boogle's certificate.

Q3.2 The attacker obtains the private key of a certificate authority trusted by users of Boogle.

Q3.3 The attacker obtains the private key corresponding to an old certificate used by Boogle's server during a past connection between a victim and Boogle's server. Assume that this old certificate has been revoked and is no longer valid. Note that the attacker does not have the private key corresponding to current certificate.