*Department of Electrical Engineering and Computer Science*

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

**6.566 Spring 2024**

# Quiz II Solutions

Mean 129.7 Standard deviation 31.9

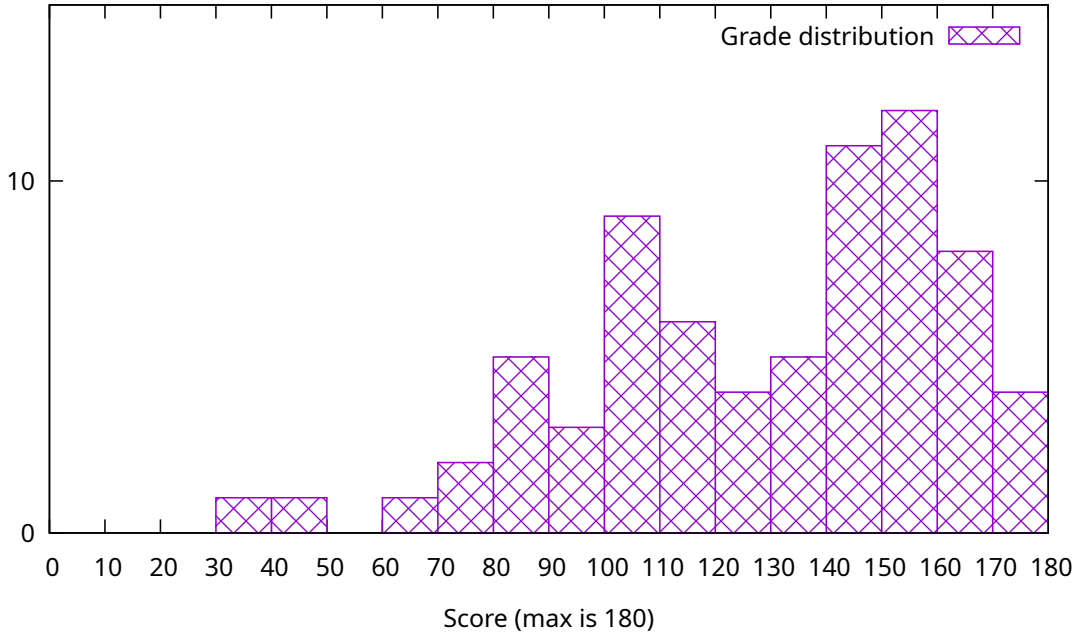

# I Guest lectures

According to Danny Weitzner, which of the following are major risks of key escrow?
**(Circle True or False for each choice; we subtract points for incorrect answers.)**

**1. [8 points]:**

**A. True / False** Key escrow undermines forward secrecy. **Answer:** True. Escrow means keys have to be derivable in the future (e.g., by a law enforcement agency), which is the opposite of forward secrecy.

**B. True / False** Key escrow means it will be possible to guess private keys through brute force. **Answer:** False. A key escrow system doesn't necessarily mean that the keys can be brute-forced.

**C. True / False** Key escrow introduces complexity that leads to bugs. **Answer:** True.

**D. True / False** Key escrow has high overheads, reducing performance. **Answer:** False. Key escrow doesn't necessarily have to impose high overheads.

According to Danny Weitzner, which of the following is the main challenge in appropriately estimating the risk of losses due to computer security breaches:
**(Circle the best choice; we subtract points for incorrect answers.)**

**2. [8 points]:**

**A.** Computer attacks are unpredictable.

**B.** Administrators can configure the same software in different ways, leading to different vulnerabilities.

**C.** There is not enough data about losses due to computer security breaches.

**D.** There are too many different problems that lead to significant losses.

**Answer:** C: not enough data. Danny did mention D in his talk, but it was not a challenge for estimating risk.

According to Russ Cox, which of the following would be examples of a supply chain attack against Google (not necessarily "open-source supply chain attack"):

**(Circle True or False for each choice; we subtract points for incorrect answers.)**

**3. [8 points]:**

**A. True / False** An adversary bribes a software developer at Google to use a known-buggy old version of an open-source library. **Answer:** False. The attack is by a software developer at Google, and does not involve the delivery of software to Google.

**B. True / False** An adversary gets a job at Google working on their compiler infrastructure, and introduces a backdoor into Google's version of the Clang C compiler. **Answer:** False. The attack is by a software developer at Google.

**C. True / False** An adversary bribes the maintainer of a popular open-source library to include an exploitable vulnerability, and this library is then downloaded and used by Google in its products. **Answer:** True. The attacker compromises the software before it's delivered to (and used at) Google.

**D. True / False** An adversary breaks into the web server hosting the Clang C compiler and replaces the latest version of the released source code with one that has a backdoor, which Google later downloads and uses in its systems. **Answer:** True. Again, the attacker compromises the software before it is delivered to Google.

**4. [8 points]:** According to Max Burkhardt, defenders are often thought to be at a disadvantage because attackers only need to find one way to break in, whereas defenders need to prevent all possible ways to break in. Why, then, did Max say that attackers are suddenly at a disadvantage after their initial break-in to gain access into some system?

**Answer:** Because the attackers have to do everything right to avoid detection after they've gained access to a system, whereas a defender just has to catch one of the attacker's mis-steps that cause their actions to be logged / detected.

# II TLS

Ben Bitdiddle is in charge of running a web server that supports TLS 1.3. He wants to log the plaintext HTTP requests sent to that web server, for debugging purposes. He has access to the TLS server certificate and the correspoding private key from the web server, but he does not want to make any changes to the web server machine or to the clients that are issuing these HTTP requests.

**5. [10 points]:** Ben sets up a second machine, with the TLS certificate and private key, and sends a copy of all network traffic to this second machine. Explain why Ben will not be able to log plaintext HTTP requests.

**Answer:** TLS provides forward secrecy, which means an adversary cannot decrypt past connections. This is the same setting as Ben's second machine, with access to the TLS certificate key, passively monitoring network traffic.

**6. [10 points]:** Explain what Ben needs to do in order to log plaintext HTTP requests without making any changes to the web server or the clients. Ben can change DNS records, though.

**Answer:** Set up another machine that will accept TLS connections, using the certificate and private key, and proxy the requests to the real web server. Log the request along the way, while proxying it.

Alyssa P. Hacker is building a mobile game that protects its network connections with TLS 1.3, and Alyssa is worried about latency in her game. While profiling the performance of her game, Alyssa notices that connections from her mobile device to her server use the P-256 EC Diffie-Hellman key agreement scheme, but that the X25519 EC Diffie-Hellman key agreement scheme is about twice as fast, both on the mobile device and on the server: P-256 takes about 60 microseconds, while X25519 takes 30 microseconds.

Alyssa changes her server to use X25519 instead of P-256. While the game still works, Alyssa is surprised to see that latency from her mobile device has increased substantially, by many milliseconds.

**7. [12 points]:** Explain why Alyssa observes higher latency with X25519 compared to P-256, and suggest what she should do to avoid this latency increase.

**Answer:** The TLS 1.3 client on the mobile device is guessing that the server supports P-256, and sends a P-256 key in the first packet; because the guess is wrong, the handshake ends up taking another network round-trip. Alyssa should change the TLS client in her game to guess X25519 instead.

# III Network performance

Suppose you want to transfer a file by sending it over a TLS 1.3 connection; that is, you run:

```
cat file | openssl s_client -connect othermachine:443
```

where othermachine is the hostname of the machine you are sending the file to, listening on port 443. For the following scenarios, estimate how long it will take, from the time you run the above command, until the destination machine gets the contents of your file. Assume your local DNS resolver already has the name `othermachine` cached, that the server is already running and ready to accept connections, and that the client knows the server's DH scheme for TLS 1.3 1-RTT mode. You can ignore TCP slow-start and window effects. Your TLS 1.3 client and server are using X25519 for DH key exchange (which has a throughput of 33,000 key exchange computations per second), RSA-2048 for signatures (which has a throughput of 2000 signs per second and 60000 verifications per second), and AES-128-GCM for authenticated encryption (which has a throughput of 1 GByte/second).

**(Circle the best choice; we subtract points for incorrect answers.)**

*Flip over to the back side of this page for the three scenarios you should analyze.*

**8. [14 points]:** The other machine is your friend in Australia. The round-trip latency to Australia is 300 milliseconds (150 milliseconds one-way), and the available bandwidth is 100 Mbit/sec. You are sending a 16-byte file.

**A.** 300 milliseconds

**B.** 450 milliseconds

**C.** 750 milliseconds

**D.** 1200 milliseconds

**Answer:** C. Dominated by network round-trips. One round-trip to establish TCP, one round-trip to exchange keys, assuming a correctly guessed DH scheme and a certificate that fits in a single network packet, and a one-way network delay to send the 16-byte encrypted file. So, about 300+300+150=750 milliseconds.

**9. [14 points]:** The other machine is directly connected by a 10-Gigabit ethernet link. The round-trip latency is 10 microseconds (5 microseconds one-way), and the available bandwidth is 10 Gbit/sec. You are sending a 1 KByte file.

**A.** 20 microseconds

**B.** 30 microseconds

**C.** 60 microseconds

**D.** 600 microseconds

**Answer:** D. Dominated by key-exchange costs. Two round-trips for TCP and TLS setup are just 20 microseconds total, but we also need ECDH, RSA signing, and RSA verification. The slowest by far is RSA signing: about 0.5 milliseconds. So the total will be a bit over 0.5 milliseconds.

**10. [14 points]:** The other machine is in another datacenter in the same city. The round-trip latency is 2 milliseconds (1 milliseconds one-way), and the available bandwidth is 1 Gbit/sec. You are sending 1 GByte.

**A.** 6 milliseconds

**B.** 100 milliseconds

**C.** 1 second

**D.** 8 seconds

**Answer:** D. Dominated by the gigabit link bandwidth limit: 8 seconds to send 1 GByte of data.

# IV Certificates

Ben Bitdiddle wants to get a TLS certificate for his web server, but his web server doesn't support directories that start with a dot, like the `/.well-known/` directory required by the ACME HTTP challenge. Ben wants to propose to Let's Encrypt that the ACME client should be able to specify the full pathname for the challenge URL when requesting a certificate from Let's Encrypt, instead of having Let's Encrypt hard-code the `/.well-known/acme-challenge/` directory as part of the standard.

**11. [10 points]:** Explain what security problem would arise if Ben's proposal were adopted.

**Answer:** Ben's modified HTTP challenge would not be secure for web sites that host user-supplied content, like Facebook, Reddit, etc, because an adversary could post the expected challenge response on the web site, and supply the URL where that challenge response appears to the Let's Encrypt ACME server.

# V Signal

**12. [12 points]:** Alyssa P. Hacker somehow manages to get the private keys for Ben's ephemeral prekeys (denoted $eprek^B$ in the "Analysis of the Signal Messaging Protocol" paper) from the Signal app on Ben's phone (and no other secret state). Suppose that Ben's friend Carol already knows Ben's identity key $ipk^B$ and starts a fresh conversaton with Ben. If Alyssa has control over Carol's network and Carol's connection to the Signal server, explain how Alyssa can decrypt Carol's message to Ben, or explain why she cannot do so.

**Answer:** She cannot do so. See Figure 3 of the Signal paper.

Answer 1: Alyssa doesn't have $ik^B$, whose corresponding $ipk^B$ was already distributed to Carol. She needs that in $DH(ik^B, epk^C)$.

Answer 2: Alyssa doesn't have $prek^B$. She needs that in $DH(prek^B, epk^C)$ and $DH(prek^B, ipk^C)$. She can't forge $prek^B$ because it needs to be signed with $ik^B$, which she doesn't have and can't forge.

In a separate scenario, unrelated to the above question, Carol keeps sending messages to Ben, Ben reads those messages, but does not reply. Alyssa P. Hacker monitors all encrypted messages sent between Carol and Ben. Alyssa breaks into Ben's phone and gets the current chaining key $ck$ for Ben's conversation with Carol.

**13. [10 points]:** Explain how Alyssa can decrypt Carol's messages to Ben from before Alyssa obtained the chaining key, or explain why she cannot do so.

**Answer:** Cannot: old message keys were derived from an earlier iteration of the chaining key, and cannot be re-derived.

**14. [10 points]:** Explain how Alyssa can decrypt Carol's messages to Ben sent after Alyssa obtained the chaining key, or explain why she cannot do so.

**Answer:** Compute the chaining key forward to obtain the message keys for future messages from Carol to Ben.

# VI Differential privacy

Suppose that you are tasked with adding a feature to Gradescope that reveals various statistics about exam scores in a class (where each score is between 0 and 1). For each of the following, explain whether or not it would achieve differential privacy in all cases, and why.

**15. [10 points]:** Publishing the list of student scores with Laplace($1/\varepsilon$) noise added to each score.

**Answer:** Not differentially private: the published list reveals exactly how many students are in the class, which would not be a possible answer if one of the students were to be removed from the database.

**16. [10 points]:** Publishing the sum of exam scores of students that received a score of 0.5 or higher, with Laplace($1/\varepsilon$) noise added to the sum.

**Answer:** Differentially private: the sensitivity is 1 even after some records have been filtered out.

*NOTE: The feedback question is on the back side of this page.*

# VII 6.566

We'd like to hear your opinions about 6.566. Any answer, except no answer, will receive full credit.

**17. [4 points]:** Out of the papers and guest lectures that we covered in the second part of the semester, listed below, circle the one that you think we should definitely remove next year (or circle the last choice if you think all papers and guest lectures should stay).

- Supply chain security: guest lecture by Russ Cox.
- Network security: TCP/IP.
- Secure channels: TLS 1.3.
- Certificates: Let's Encrypt.
- User authentication: U2F.
- Messaging security: Signal.
- Key transparency: CONIKS.
- Anonymity: Tor.
- Cybersecurity policy: guest lecture by Danny Weitzner.
- Security economics: Spamalytics paper.
- Differential privacy: PINQ.
- Information security in real life: guest lecture by Max Burkhardt.
- Do not remove any papers.

**Answer:** 21x Spamalytics (not a whole lecture worth of material). 17x PINQ. 14x None. 11x Danny's policy guest lecture. 6x CONIKS. 5x Signal (paper too hard to read; refactor lecture). 1x Max Burkhardt guest lecture. 1x TCP/IP security (attacks not so relevant). 1x Supply chain security.

1x Liked TLS 1.3. 1x Liked Let's Encrypt. 1x Liked Tor. 1x Liked Danny's policy guest lecture. 1x Liked Max Burkhardt guest lecture.

**18. [4 points]:** Were the lab recitations helpful? Should we keep them next year?

**Answer:** 48x yes. 13x did not attend. 10x not super helpful.

12x notes were helpful (even if did not attend). 8x especially lab5. 5x TAs should prepare code examples or slides, not just reading from lab page or Piazza; make the recitation more structured. 4x office hours and Piazza were just as helpful. 3x would be nice if they were recorded. 1x would have been more helpful after spending time on the lab. 1x recommend that students read the lab before attending recitation.

**19. [4 points]:** What else would you suggest we improve next year?

**Answer:** Labs and assignments: 9x more hints/guidance on lab5/webauthn (e.g., about encodings). 3x more conceptual assignments that are not labs / more like exam questions. 2x lab5 was too long; break up into sub-parts. 2x more labs. 2x better documentation. 2x lab4 AWS issues were frustrating; debug AWS setup beforehand. 1x technical problems in labs were frustrating. 1x exercises with real-life attacks. 1x recommend that everyone use AWS for labs. 1x lab5 was jarring change from previous labs in terms of not being step-by-step. 1x lab4 attacks felt unrealistic because they required disabling defenses. 1x cryptocurrency attacks. 1x build personal HTTPS web site. 1x make lab5 more verbose, give more skeleton code. 1x debugging guide for each lab.

Lectures: 7x answer reading questions in lecture, or post answers afterwards. 4x enjoyed guest lectures / more guest lectures. 3x more up-to-date and more detailed lecture notes. 2x more big-picture recaps to see how everything fits together. 1x liked guest lecture by max as wrap-up. 1x signal paper was hard to read. 1x live demos were hard to follow. 1x improve logical flow for first half of class to match the network section. 1x cryptocurrency / blockchain security. 1x more theory / math. 1x kerberos. 1x more real-world attacks like xz. 1x IoT security. 1x explain TLS from start to finish. 1x read paper after lecture rather than before.

Topics: 1x more hardware security. 1x more about firewalls. 1x how do antivirus tools work.

Logistics: 3x assign labs at same time as corresponding lectures. 2x more exam reviews; record review session. 2x align quiz questions with lecture content. 1x lecture Q-and-A from other students, for new papers. 1x weekend office hours should not be at 9am. 1x consistently turn on the mic for recordings. 1x more consistent recitation scheduling. 1x practice exam problems for new papers. 1x make it clear what the anonymous question site is, such as by posting on Piazza. 1x auto-apply lab extension days in gradescope.

# End of Quiz