ECE 750: Usable Security and Privacy - Practice	Marks obtained ↓
Date: July 8, 2025, Total questions: **8** Total points: **74**	
ID: Name:	Time: 40 minutes

Page:	2	3	4	5	6	7	8	10	11	Total
Points:	12	4	6	10	8	14	10	6	4	74
Score:										

Mock Instructions

Practice only - This "Exam" is only for practice, it will not impact your final grade in any way.

40 minutes - We will only be spending 40 minutes today taking the exam, the real exam is 2.5 hours.

Merged exam - I merged a couple USEC exams together to make this one. These are all real questions that appeared on finals. But the exam is a bit more password/authentication heavy than a real exam would be.

Original Instructions

No aids allowed. All you are allowed is a pen and pencil.

Use space provided. Answer the questions in the spaces provided. If you run out of room for an answer you can use the back of the page but clearly note in the original answer space that the back was used.

Point value in right-hand margin. The number of points each question is worth is listed in the right hand margin. The number of points and the space provided are roughly indicative of how long of an answer is expected.

Pencils and pens allowed. Both pens and pencils are allowed. Pencil marks need to be clearly present or erased. If it is unclear if a mark is or is not present then it is up to the discretion of the grader and this point cannot be contested later.

Scratch paper. At the end of the exam there are two blank pages marked "SCRATCH PAPER", you may tear these off and use them as scratch paper.

Solution: Solutions appear in boxes like this one. Most questions have multiple possible solutions. The solutions in the boxes are the solution I originally thought would be chosen by most people. Sometimes there is a POST MARKING COMMENTS box which was written after grading the 2024 exam to address answers I saw.

Online safety

1. When people talk about being "safe" online they mean safe against certain threats. Name two threats actors that general people might consider to be a threat online. I am looking for people, groups, or organizations that could be considered threats, I am not looking for technical threats. The next question will use the threats you name, so I recommend reading it as well before answering. (2)

   a: _____

   b: _____

   Solution: Answers should show some thought about who it is that threatens users online. Some possible answers include: advertisers, insurance companies, stalkers (via buying data broker data), attackers aiming to steal money by buying ads that trick users, nation states, divorce lawyers, and social media companies.

2. For both of the two threats you named in question 1, what technical protections other than encryption could be used to defend against the threat, and why are those protections effective for this specific threat? (6)

   Solution: Answers here should show thought about the types of attacks this attacker might use. An advertiser, for example is likely to want to collect web visitation data. So an ad blocker would be quite effective, as would any tool that blocks trackers. Opting out of tracking might also be effective since generally advertisers follow the law.

Passwords

3. Imagine an attacker was able to gain access to the shadow password file on a Linux server. The file contains usernames, hashed passwords, and their associated salts. The attacker wants to find the passwords of users like root.

   (a) We learned about online and offline password guessing attacks. Is the attacker most likely to use an online or an offline attack in the above situation? Briefly describe what about the situation makes your chosen attack the most appropriate for the attacker. (4)

   Solution: Offline. Because the attacker has the file content and can therefore directly make guesses against the hash+salts. In an offline attack they can create and test their guess locally which is the fastest option. In an online attack they would be limited by the network speeds, page reload time, and any guessing rate limits.

   Solution: POST MARKING COMMENTS:

   - Pre-computed hash tables (aka rainbow tables) will not work in this case because the passwords used a salt. The salt means that the attacker must compute hash(password,salt)==passwordHash for every possible password and every password entry in the file. There is no way to do a precomputation because the salt for every password is different.

(b) Imagine a user is assigned a 4 character password randomly selected from a possible set of 56 characters. The password is used on a website that uses lockout. If more than 10 passwords are attempted against 1 user name, then the system locks the account for 24 hours. Under what assumptions might this password be considered sufficiently secure? (4)

Solution: Under the scenario in Question 3 above, this situation can never be secure because once the shadow password file is lost, 4 character passwords are simply too insecure.

OR

If we assume a remote attacker who does not have local access, then the "don't care" region of password strength becomes important. If lockout is correctly implemented then the attacker is limited in their guessing to a relatively small number of passwords. As long as the resource being protected isn't too sensitive. A short password can be ok if it is truly random.

Online Experiment

4. Long random passwords are better for security because they are harder to guess. But long truly random passwords are bad for usability because they are challenging for people to remember.

   To better understand what length random password people can reliably remember Edith designs an online user study. She recruits people using an online survey platform like Amazon Mechanical Turk and requires them to have a high positive rating on the platform. She uses survey software to create an online experiment where users see information and answer questions via a webpage. Users in her study go through the following steps:

   1. Consent to be part of research.
   2. Survey software randomly decides if they will get a 6, 9 or 12 character password, and then creates and records a random password of that length containing at least one upper case, one lower case, one number, and one symbol. The user is shown the password and told they will need to enter this password later in the study.
   3. On the next page the user is asked to enter the password to make sure they remember it, if they enter the wrong password, they are shown the password again and reminded they need to remember it.
   4. The user then fills out the IUIPC scale. (A set of survey questions measuring privacy.)
   5. The user is asked to enter their password. No feedback is given about if it is correct or wrong.
   6. The user then fills out the SeBIS scale. (A set of survey questions measuring security behavior intention.)
   7. The user is asked to enter their password. No feedback is given about if it is correct or wrong.
   8. The user fills out a set of demographics questions including age, gender, native language, and current ocupation.
   9. The user is asked to enter their password. No feedback is given about if it is correct or wrong.
   10. They are thanked for taking part in the survey.

   To analyze the results, Edith counts how many times participants enter their password correctly on steps 5, 7, and 9 producing a score between 0 and 3. She then uses t-tests to see if there is a statistically significantly difference in the mean number of errors between each pair of password lengths.

   Answer the questions below about the study. Feel free to use the step numbers above in your answers.

   (a) List the **independent** variable(s) Edith is collecting in her study. (3)

   > **Solution:** She collects and uses: the condition (step 2, password length). She also collects demographics (step 8). She also collects the IUIPC scale (step 4), and SeBIS (step 6) though it is unclear how they are used.

   (b) List the **dependent** variable(s) Edith is collecting in her study. (3)

   > **Solution:** Password entry accuracy which is a sum of steps 5, 7, and 9. The password entry in step 3 is ignored in the analysis and so is normally not a dependent variable.

(c) The study has several internal validity problems. Describe one of them. (6)

Solution: The largest issue is that users are being asked to remember the password over a very short time and being frequently asked to recall it. The situation is quite different from how normal password recall is meant to happen. Online users who know that they do not need this password long-term are quite likely to write the password down rather than try and recall it from memory. These issues mean that Edith is unlikely to be able to accurately answer her question.

(d) Like all studies, this study has some external validity limitations. Describe one external validity limitation. (4)

Solution: The most obvious external validity issue is the users chosen. Similar to Assignment 2, the people being studied are all from an online survey platform. They take surveys every day and are generally quite familiar with the Internet. The study results are unlikely to directly align with how people less skilled with internet technology might behave.

Solution: POST MARKING COMMENTS:

- Limiting participation to those that have high scores on Mechanical Turk is a normal thing to do in studies. It is done to avoid having participants who are known to cheat or that use robots to fill in answers. It is an external validity issue, but having lots of participants that are known to not pay attention when taking surveys is a larger issue. We did not learn this point in class, so no marks were taken in regards to it.

Survey study

5. Amy wants to understand if people understand how secure their own passwords are. She decides to run a survey study where she prints the survey on sheets of paper and has volunteers ask people to take the survey at places around campus that have allot of people, like the line in front of Tim Horton's, cafeterias, and the bus station. The survey asks people the questions below. After collecting the data she analyzes the password complexity and compares it to how challenging people think their password would be to guess.

> For each of the following, circle your answer or fill in the blank.
>
> i) Which best describes you: a) student, b) staff, c) faculty
>
> ii) Which of the following does your password contain: a) upper case letters, b) lower case letters, c) numbers (0-9), d) symbols (&,%,.,-, etc.)
>
> iii) Is the first character of your password an uppercase letter? YES NO
>
> iv) Does your password contain English words? YES NO
>
> v) Is the last character of your password a symbol like &, $, and @. YES NO
>
> vi) How long is your password? _____
>
> vii) Do you think your password would be challenging to guess?
> Very challenging, Challenging, Somewhat challenging, Not challenging

(a) Select one of the question's from Amy's survey above that you feel could be phrased better. State what you think is wrong with the question and propose better wording or better options. (4)

Solution: Several have problems with them. The easiest is probably v "How long is your password?" because it is unclear what "long" means, a computer scientist will assume number of characters, but someone might answer in terms of how many years they have had the password. A better phrasing would be "How many characters are in your password? Including letters, numbers, and symbols."

(b) What are the **independent** variable(s) that Amy is collecting? (2)

Solution: The independent variables are questions i)-vi.

Solution: Marking Notes: Several answers included information that *could* have been collected but is not being collected according to the above paragraphs. I generally ignored these answers as neither right nor wrong and did not take points off for them.

(c) What are the **dependent** variable(s) that Amy is collecting? (2)

Solution: The dependent variable is Amy's comparison between the password complexity and the answer to how challenging they think their password is to guess (Q vii).
OR
The dependent variable is how challenging they think their password is to guess (Q vii).

(d) What **external validity** issues exist with this study? (4)

> **Solution:** Many. The largest issue is how people are being recruited. Recruiting from busy places on campus likely over sample for students. People who bring their own lunch or who drive to work are less likely to encounter one of Amy's volunteers. That means that the results will better represent students, and represent staff and faculty less well. Amy is also not recording what college/department survey takers are in. It is quite likely that someone from a computer-focused degree program may behave differently in terms of passwords than someone in a college/department that uses computers less.

(e) Does Amy's study have good **internal validity**? Why or why not? (4)

> **Solution:** No. The information collected in this survey is not enough to accurately compute entropy and it definitely is not enough to compute guessability. So this study will not definitively show if people are accurate in their assessments or not.
>
> This question is listed as a bonus mostly because *internal validity* was only lightly touched on in class and isn't on the slides at all. Internal validity is if the study design itself makes sense and is able to answer the question the researcher asked. An extreme example would be if a researcher wanted to measure peoples favorite colors so he asked them what their favorite music was. Knowing favorite music doesn't really answer what their favorite color is, so the study has poor internal validity.

Public/Private Key Cryptography

6. (a) Fill in the blanks in the following text describing Alice and Bob correctly communicating securely. Each blank needs to be filled in using one of the following words: (6)

- Encrypt, encrypted, encrypts
- Decrypt, decrypted, decrypts
- Sign, signed, signs, signature
- Public
- Private

Alice wants to send an encrypted and signed message to Bob who she has met in the past. When they last met in person they verified and then _____ each other's _____ keys using their respective _____ keys.

To send an email to Bob, Alice first _____ the message using her _____ key and then _____ the resulting message using Bob's _____ key. Alice then sends the message over an untrusted connection.

Bob receives the message and first _____ it using his _____ key. He then verifies that the message really was from Alice by verifying the _____ using Alice's _____ key.

Solution: Alice wants to send an encrypted and signed message to Bob who she has met in the past. When they last met in person they verified and then **signed** each other's **public** keys using their respective **private** keys.

To send an email to Bob, Alice first **signs** the message using her **private** key and then **encrypts** the resulting message using Bob's **public** key. Alice then sends the message over an untrusted connection.

Bob receives the message and first **decrypts** it using his **private** key. He then verifies that the message really was from Alice by verifying the **signature** using Alice's **public** key.

(b) Research, such as the *Why Johnny Can't Encrypt* paper, often mentions "confusing metaphores" in regards to encryption interfaces. Give a specific example of a metaphor related to encryption that is likely confusing for people. (4)

Solution: The most obvious answer is a key or a lock. Both are commonly used in encryption technology to refer to the act of encryption using the metaphor of a key and lock. Unfortunately physical keys work rather differently than encryption keys, particularly in regards to public key cryptography where there are two keys that both interact with the same lock (encrypted text) in different ways.

(c) Would Alice and Bob's security in the interaction above be improved by using a Certificate Authority (CA)? State 'Yes' or 'No' and then briefly explain what a CA would do to help Alice and Bob OR explain why their current security practices are equal to or better than what a CA would provide. (6)

Solution: No. Alice and Bob directly verified their public keys and signed them in the past. A CA's role is to attest that a particular public key is associated with a specific person. Since Alice and Bob already directly verified and signed each other's keys, they do not need a third party to do the verification. Therefore a CA would not improve the security of their interaction. If anything it would add the risk that the CA was corrupt which would increase the risk. CAs are run by many entities, including entities tied to governments. So when two countries are in conflict, sometimes the CA for one country will sign bad certificates enabling a digital attack.

Solution: POST MARKING COMMENTS:

- This question and the points below assumes you answered Question (a) correctly. If you did not, that was taken into account.
- A Certificate Authority *only* provides verification of identity by using its private key to sign the server's public key. It puts the signature in a "certificate" that contains information beyond just the key like the domain the key is for. It is used during a TLS setup by the browser, app, or email provider. But a CA itself does not provide TLS or in any way perform the encryption.
- TLS itself happens over an untrusted connection. A trusted connection would be something like a dedicated wire used only for this purpose or a courier who is paid to hand deliver a message.
- TLS and PGP use the same level of encryption. The keys are different, so there is a tiny improvement in terms of someone being able to break the encryption. But it is rare now

days that anyone outside of research tries to break encryption directly. It is much easier to do things like bribe a CA to sign the wrong key.

- Not covered in this course, so I do not expect you to know. But there are situations where applying encryption more than once can degrade the encryption and make the message less secure.

Encryption

7. The following is an excerpt from the *Why Johnny Can't Encrypt* paper which explains the steps a participant had to go through to setup their own key pair and send an encrypted message to 5 people.

> In order to [correctly send an encrypted email], a participant had to generate a key pair, get the team members' public keys, make their own public key available to the team members, type the (short) secret message into an email, sign the email using their private key, encrypt the email using the five team members' public keys, and send the result. In addition, we designed the test so that one of the team members had an RSA key while the others all had Diffie-Hellman/ DSS keys; thus, if a participant encrypted one copy of the message for all five team members (which was the expected interpretation of the task), they would encounter the mixed key types warning message. Participants were told that after accomplishing that initial task, they should wait to receive email from the campaign team members and follow any instructions they gave.

WhatsApp, a modern chat app, also does end-to-end encryption but the steps to send an encrypted message are essentially invisible to the user unless they explicitly look in the settings. Each device running WhatsApp automatically creates its own public/private key pair. The main WhatsApp server keeps a record of which public keys are associated with which users. When a user wants to chat with someone (a phone number) the central WhatsApp server provides the appropriate public key. Private keys never leave their devices.

(a) WhatsApp's approach requires the user to perform fewer steps and is therefor easier to use than the PGP approach tested in *Why Johnny Can't Encrypt*. Describe one situation where the user would be safe using PGP but not if they used WhatsApp. You can assume that they can use PGP correctly without making errors. (6)

Solution: The extreme example is if a nation state, or a hostile WhatsApp employee, got WhatsApp to put the wrong public key on their server. Such a situation would allow an attacker to listen in on all the messages and even inject messages. This attack though is unlikely to happen to average citizens. This attack cannot happen with PGP because there is no central repository that can be poisoned and even the repositories that exist also have signatures associated with the keys that cannot be spoofed.

A less extreme example involves metadata. WhatsApp is owned by Facebook. If Facebook ever wanted to use information about who is communicating for marketing purposes or to "better match content to your interests" the user could not protect themselves. PGP is run over email (or really anything that supports text) so the two people communicating can pick whatever medium they want to avoid observation. They could even chat over Reddit.

If you only consider the description above (not how WhatsApp actually works) then another attack is the phone number. If someone were to steal a phone number by say bribing a telecom provider employee, then they could pretend to be the person who's phone number they stole. WhatsApp server will accept their new public key if it is sent from a phone with the correct phone number.

Solution: Marking Notes: As stated in the question:

- WhatsApp's central server does not have private keys for anyone.

- PGP in Why Johnny Can't Encrypt also uses a key store in the internet. Users can exchange public keys directly, but similar to WhatsApp they normally do not and instead rely on the ones posted on the key store. The big difference is that in PGP those public keys are signed and PGP will try and verify the signatures using web of trust.
- I don't think that taking a person's sim (phone number) is enough to start getting their WhatsApp messages. But information about that isn't in the question and I am not certain so I did not mark down for the answer.
- Answers of the form: "imagine someone stole Alice's phone which had WhatsApp installed but not an email app" did not get full marks. The statement is true, but exhibits no understanding of the technologies involved.

(b) Name and briefly describe a type of study that could be used to better understand why most WhatsApp users are not manually verifying the keys they receive from the central server. You do not need to plan a whole study, a simple description is fine. (4)

Solution: Many answers possible. A qualitative study that looks at attitudes would likely be best. Users are likely not manually verifying because: 1) they don't know they can, 2) they don't know why they should, or 3) they want to verify but cannot figure out how to do so.

If we assume 1 or 2, then a qualitative study is best. Something like an interview. Or a survey that has open text boxes where people can write in their reasoning. A focus group would also work.

If we assume the problem is 3, then a study that looks at the user interface would be best. A cognitive walkthrough might work. As would a think-aloud. As long as they focused on seeing if users can find and correctly use the verification feature.

Solution: Marking Notes: I expected to see a type of study named or very clearly described in specific terms that make the exact type of study very clear. There is a huge difference between a cognitive walkthrough, a lab study, and a survey. Just saying "study" was not enough.