Corrected

COM 402 Final Exam

25.01.2021

1
Name: **Anonymous**
Sciper: **123456**

Do NOT open this document until instructed to do so.

Instructions

- This is a **closed book** exam. Books, notes and electronic devices are not allowed, except for up to two A4 pages of notes (this amounts for *one* double-sided A4 sheet or equivalent).
- **No question will be answered during the exam**, except possibly questions related to the understanding of the English language.

Multiple choice questions:

- There are 17 multiple choice questions, counting 1 point each.
- Only one answer per question is correct.
- Ticking the correct answer counts for 1 point.
- Ticking an incorrect answer, more than one answer, or no answer counts for 0 point.
- Make a mark **inside** the box corresponding to your answer.
- Use a black or blue pen to mark your final answers. **Pencils are not allowed**.
- If you ticked the wrong answer, use a white-out fluid or tape to erase the box completely. **Do not try to re-draw the box**.

Open text questions:

- There are 10 open text questions, counting 1 point each.
- Please write your answers in the corresponding text boxes.
- Do not write more than three lines. Any text outside of the boxes or after three lines **will be ignored**.
- Do not tick the '0', '0.5', '1' boxes of the top of the text boxes.

Question 1 A web application stores its data in a database on a server. For security reasons, the database is configured to encrypt all data before writing it to disks. Which attack is prevented by this configuration?

- ☐ Database administrators accessing all cleartext data of the users
- ☐ Users of the application accessing cleartext data of other users through an SQL injection
- ■ Server administrators accessing all cleartext data of the users
- ☐ Users of the application accessing cleartext data of other users through an insecure direct object reference

Question 2 Consider an attacker performing a grey-box attack against a machine learning-as-a-service (MLaaS) in which the machine learning model is non-linear such as neural networks. Which one of the following is *true*?

- ☐ The attacker must have information about the set or subset of the training data to steal the model.
- ☐ Membership inference attacks are not possible.
- ☐ Model stealing attacks are not possible.
- ■ If MLaaS is using output perturbation, the probability of a successful model stealing attack is reduced.

Question 3 Consider a setting where Alice and Bob use a Cloud service operating on data encrypted using the Paillier homomorphic cryptosystem. Let sk_A, sk_B be the secret-keys of Alice and Bob (respectively) and pk_A,pk_B be their public-keys (respectively). Also, let m_A and m_B be two secret values held by Alice and Bob (respectively) and $\mathsf{ct}_A = \mathsf{Enc}_{\mathsf{pk}_A}(m_A)$ and $\mathsf{ct}_B = \mathsf{Enc}_{\mathsf{pk}_B}(m_B)$.
If provided with pk_A, pk_B, ct_A and ct_B, which of the the following values can the Cloud compute?

- ☐ ct_{res} such that $\mathsf{dec}_{\mathsf{sk}_A}(\mathsf{ct}_{res}) = m_A \times m_B$.
- ☐ ct_{res} such that $\mathsf{dec}_{\mathsf{sk}_A}(\mathsf{ct}_{res}) = m_A + m_B$.
- ☐ ct_{res} such that $\mathsf{dec}_{\mathsf{sk}_A}(\mathsf{ct}_{res}) = m_B$.
- ■ None of these choices.

Question 4 Recall that adversarial examples are inputs to a machine learning model that an attacker has designed to cause the model to make a mistake. Which one of the following is *true* about adversarial examples?

- ☐ In high-dimensional spaces, it is possible to defend against all possible adversarial examples by detecting suspicious queries.
- ■ An attacker may use the transferability property of adversarial examples to attack different machine learning models with the same adversarial example.
- ☐ In all cases, adversarial examples have to be imperceptible by humans.
- ☐ It is possible to defend against all possible adversarial examples via adversarial training.

Question 5 The GA4GH Beacon Project gathers patients' genomic data in a database that can be queried by researchers. To limit privacy concerns, the researcher only receives "yes/no" answers. We assume here that the attacker is a honest-but-curious adversary and acts as a researcher sending queries to the database through the correct interface. Which of the following is *true*?

- [x] The probability of success for a re-identification attack consisting of j queries can be lowered by only answering "yes" to queries for which there are more than k patients satisfying the query terms.
- [ ] An attacker having no other side information about a subject except than his name can still re-identify or infer (with 100% certainty) the presence of this subject in the Beacon Project database.
- [ ] Limiting the number of queries that an attacker can send is enough to make re-identification and membership inference attacks impossible.
- [ ] It is possible to protect the database against re-identification attacks without altering its utility.

Question 6 Consider the following Python code. `validate_password` is a function that is vulnerable to a timing attack; an adversary can repeatedly call the function (for instance, as a remote procedure call) and analyze the timings to recover "`pwd`" without a full brute-force attack. The operator `^` denotes the bitwise exclusive OR and `|` the bitwise (non-exclusive) OR. `ord()` converts a char to a number, and `zip(a,b)` loops over the two sequences simultaneously.

```
1.  pwd="password123"

2.  def validate_password(user_input):
3.      return user_input == pwd             # timing attack!

4.  def fn1(a,b):
5.      a_original = a
6.      while len(a) < len(b):
7.          a += " "
8.      result = 0
9.      for x, y in zip(a, b):
10.         result = result | (ord(x) ^ ord(y))
11.     return (len(a_original) == len(b)) and (result == 0)

12.  def fn2(a,b):
13.     import time, random
14.     time.sleep(random.randint(1,100) / 10000)
15.     return (a == b)

16.  def fn3(a,b):
17.     result = 0
18.     for c in a:
19.         result = result ^ ord(c)
20.     for c in b:
21.         result = result ^ ord(c)
22.     return (len(a) == len(b)) && (result == 0)
```

To remove the timing attack in validate_password(), prevent the adversary to recover 'pwd', and keep the intended functionality, the *most secure option* is to:

- [ ] Replace line 3 by `return fn3(user_input, pwd)`
- [ ] Replace line 3 by `return user_input==sha256(pwd)`
- [x] Replace line 3 by `return fn1(user_input, pwd)`
- [ ] Replace line 3 by `return fn2(user_input, pwd)`

Question 7 The *Trojan* problem in access control relates to the case were a malicious program uses the access rights of a user to access protected data of this user and make it accessible to everybody. Which form of access control is designed to prevent this type of attacks?

- ☐ Discretionary access control
- ☐ Mandatory Access Control with a *no write-up* rule
- ■ Mandatory Access Control with a *no write-down* rule
- ☐ Role based access control

Question 8 Which of the following choices correctly completes the following statement?
The completeness, soundness and zero-knowledge properties of zero-knowledge-proof systems can be guaranteed in settings where...

- ■ *a potentially dishonest prover wants to prove a statement to an honest verifier.*
- ☐ *an honest prover wants to prove a statement to a potentially dishonest verifier.*
- ☐ *a potentially dishonest prover wants to prove a statement to a potentially dishonest verifier.*
- ☐ None of these choices.

Question 9 A Web Application Firewall (WAF) is a reverse proxy that tries to detect and block attacks to a web server by looking for patterns in the requests it receives. Which type of attack can not be detected by such a WAF?

- ☐ SQL injection attacks
- ☐ Buffer overflow attacks
- ☐ LDAP injection attacks
- ■ Insecure Direct Object References

Question 10 Which one of the following statements is *false*?

- ■ In Bitcoin, stealing a specific private key allows to rewrite the transaction history for that key.
- ☐ Bitcoin's security uses digital signatures.
- ☐ In permission-less blockchains the number of participants does not need to be predefined.
- ☐ Permission-less and open blockchain refer respectively to who can write and read the chain.

Question 11 Which of the following is *not* a primitive of asymmetric cryptography.

- ■ Message authentication codes
- ☐ Public and private keys
- ☐ Digital signature
- ☐ Interactive key exchange (e.g. Diffie-Hellman)

Question 12 Consider a machine M using a Trusted Platform Module (TPM). M possesses a large, secret list of pseudonyms L_1, *sealed* by the TPM. An initiator I contacts M and first *attests* which software is running on M; in this process, a TLS connection is established. Then, I sends a private list of pseudonyms L_2 over the confidential channel. M, using the TPM, computes $L_1 \cap L_2$ and returns the results over the TLS connection. Which statement is correct ?

- ☐ On the machine M side, the TLS connection is terminated by (that is, decrypted by) the Operating System, which forwards the data to and from the TPM.
- ☐ Changes to the OS (for instance, installing software updates) will alter the attestation process, but will *not* require re-sealing the list L_1.
- ■ Any change to the OS of M (for instance, installing software updates) requires re-sealing the list L_1.
- ☐ On the machine M side, the TLS connection is terminated by (that is, decrypted by) the TPM, which forwards L_1 to the OS to compute $L_1 \cap L_2$. The TPM encrypts and returns the answer to the initiator.

Question 13 What is a *zero day* exploit ?

- [x] An exploit for a vulnerability for which no patch exists yet
- [ ] An exploit that was published after the vulnerability it exploits was patched
- [ ] An exploit that is very simple to exploit
- [ ] An exploit that exists in the original version of a software

Question 14 Which is the best way to protect against SQL injections ?

- [ ] reject any input that contains SQL keywords (e.g. union, select)
- [ ] escape all occurrences of single and double quotes (' → \', " → \")
- [x] use prepared statements
- [ ] use only indirect object references

Question 15 Consider a machine M using a Trusted Platform Module (TPM) in the process of *attesting* which Operating System (OS) is running. First, M is turned on; once fully booted, M receives a network request from an external party to prove which OS has been loaded. This request is answered to by the TPM, which establishes a TLS channel and sends back information to the initiator of the request. Afterwards, the initiator of the request may continue interacting with the TPM for other purposes (e.g., sealing/unsealing, etc). Which statement is correct ?

- [ ] The very first steps of the booting process, that is, loading the BIOS and the Master Boot Record (MBR), may be safely omitted in the attestation process.
- [x] The Endorsement Key (EK) is used in this process.
- [ ] The establishment of the TLS channel is primarily for confidentiality purposes (for instance, so that the OS does not observe the subsequent communication).
- [ ] A challenge (or nonce) is not required to securely prove a given state of the Platform Configuration Registers (PCRs), but it is necessary for sealing/unsealing.

Question 16 In this scenario *Database 1* is private and *Database 2* is public, i.e., is fully accessible by any entity. *Database 1* can be queried by researchers with only queries of the form:

```
SELECT COUNT(*) FROM database 1 WHERE Disease=arg1 AND Age in [arg2,arg2+6]
```

A researcher chooses the two arguments `(arg1,arg2)` to perform a query and does not deviate from the protocol. We assume that *Database 1* is correctly protected against any type of injections and that a researcher knows that both databases contain the same individuals. Which of the following is *true*?

Database 1

Zipcode	Age	Sex	Disease
1183	29	M	COVID-19
1456	21	F	Diabetes
1183	25	M	COVID-19
1018	30	M	Flu

Database 2

Name	Zipcode	Age	Sex
Alban	1183	29	M
Michelle	1456	21	F
Louis	1183	25	M
Joseph	1018	30	M

- [ ] If an attacker performs the query:
  `SELECT COUNT(*) FROM Database 1 WHERE Disease=COVID-19 AND Age in [23,29]`
  on *Database 1*, he/she will know with only 50% probability that Louis has COVID-19.
- [ ] Generalizing `Sex` from *Database 1* would make re-identification attacks more difficult for a malicious researcher.
- [x] A researcher needs to perform at least 3 queries to *Database 1* in order to find the disease of each individual in *Database 2*.
- [ ] Homer's re-identification attack can be used, as such, in this case.

Question 17 In e-voting, when using a mixnet to shuffle the votes, assuming each server shuffles all the $N \geq 2$ votes and assuming honest voters, which of the following statement is *correct* ?

- ☐ To prevent linking the ballots to the voters, a threshold $T > 1$ (T defined as a cryptographic threshold used in secret-sharing) of honest mixnet servers is needed.
- ☐ A global network adversary observing both the inputs and the outputs of an honest mixnet server can correlate which input correspond to which output.
- ☐ If all mixnet servers are malicious and under the control of the adversary A, then A is still unable to link a ballot in the output to a particular voter.
- ■ To prevent linking the ballots to the voters, a single honest mixnet server is needed.

ASLR and buffer overflows

Question 18 Explain how Address Space Layout Randomization (ASLR) can prevent the exploitation of a stack-based buffer overflow.

☐0 ☐0.5 ■1

ASLR randomizes the address of the stack (among others). The attackers don't know the address of the code they are writing in the stack and thus can't make a return address point to it.

Certificate pinning

Question 19 Many mobile applications use certificate pinning. This protects against a certain type of attacks. Describe this attack in following terms:

- Who can do the attack,
- what type of attack is it,
- and how they can achieve it.

☐0 ☐0.5 ■1

Trusted CAs can make a (TLS) man-in-the-middle attack by creating a fake (but trusted) certificate for the attacked web server.

Rainbow tables

Question 20 Describe a way to construct a rainbow table that can crack passwords of different length, but that finds the short passwords faster that the long passwords.

☐0 ☐0.5 ■1

Method 1: in the last columns of the table, use reduction functions that generate short passwords. Use reduction functions that generate long passwords in the rest of the table. Short passwords will be found before long passwords.
Method 2: Use a reduction function that generates all lengths of passwords but that generates short passwords with a higher probability than expected. Short passwords will be found faster in average (and have a higher success rate).

Kerberos

Question 21 Describe the operations that a server must carry out to verify that a Kerberos ticket it received together with an authenticator is a valid ticket.

☐0 ☐0.5 ■1

Decrypt the ticket with its own key. Retrieve session key from the ticket to decrypt the authenticator. Verify that the id in the authenticator and ticket are the same. Verify that the timestamp in the tick is fresh (and not a replay).

Privacy Preserving Techniques

Question 22 In privacy evaluation of systems, which of *anonymity* and *pseudonymity* corresponds to the strongest privacy guarantee? Explain why.

☐0 ☐0.5 ■1

Anonymity provides the strongest privacy guarantee, because it forbids the existence of a per-identity pseudonym that can be used to link between the actions of an identity.

Federated Learning

Question 23 Consider a federated learning setting with 4 hospitals where the data of the hospitals remain in their own trusted servers, but a global logistic regression model is trained collaboratively. Assume that this is achieved by averaging the local models from hospitals after each iteration of the learning process to obtain the global model.
State one possible protection mechanism to protect hospitals' input data *during* training, explain how to use it and state its one advantage and one limitation.

☐0 ☐0.5 ■1

HE -> encrypt model that is exchanged in between hospitals. Advantage: easier to preserve privacy, gives higher guarantees on privacy. Limitation: computation complexity, scalability in terms of complexity of model/samples, OR
DP -> add noise to exchanged values. Advantage: easier to integrate. Limitation: harder to define the level of privacy that is achieved in practice, utility degradation

Blockchain

Question 24 Can somebody who lost their bitcoin-wallet secret-key realistically recover their bitcoins? If yes, explain how, if no explain why.

☐0 ☐0.5 ■1

No. The coins are lost as, under today's standard cryptographic assumptions, one cannot realistically break the bitcoin underlying PKI to recover the key or transfer the funds

Question 25 Propose the types of blockchains required [open/closed, Permissionless (public)/Permissioned (private)] by the following two use cases and explain why:

1. Medical records among hospitals
2. Personal fitness data in crowdsourcing

☐0 ☐0.5 ■1

1. Closed + Private : very sensitive and long term data + could implicate more than the owner. 2. Open + Public/Private : depends on the motivation but crowdsourcing makes it open.

Question 26 In blockchains, is there a best performing consensus algorithm? If yes, which one? If not, explain why.

☐0 ☐0.5 ■1

No there is not. Different need for scalability, energy consumption, use-case/business requirements, security and privacy, and throughput

Privacy and health

Question 27 An imaginary government decided that a citizen should be able to know a little bit more about its fellow citizens overall health status. It therefore decided to build a central database containing the name, age and disease of each citizen. For security, all elements are encrypted under a combination of the keys of the government and the hospitals by using a probabilistic homomorphic encryption scheme. The query:

```
SELECT COUNT(*) FROM db WHERE Disease=arg1 AND age=number ,
```

can be executed by any citizen. It is performed by relying on deterministic tagging, which enables re-encryption from probabilistic encryptions to deterministic encryptions. Explain **(1)** why the selection cannot be done directly on HE encrypted data, **(2)** which information can still be deduced from the deterministic tags by an entity that has access to the query and the deterministic ciphertexts and **(3)** which privacy risk (provide one example), even without deterministic tagging, is still present even if the data are encrypted and why? One line per answer.

☐0 ☐0.5 ■1

(1) the HE scheme is probabilistic, two encryptions of the same value are different;
(2) such an entity can observe the frequency or distribution of the selection criteria .....
in the database;
(3) re-identification attacks, centralization of the data.
.....

.....

Corrected

Corrected