CHALMERS UNIVERSITY OF TECHNOLOGY
Department of Computer Science and Engineering
Examination in Computer Security EDA263 (DIT641) for the International Master's Program in Computer Systems and Networks, Friday 18 March 2022, 14:00—18:00

Examiner: Associate professor Magnus Almgren, Ph.031-772 1702,
email: magnus.almgren@chalmers.se

Teacher available during exam: Magnus Almgren, Ph.031-772 1702
The teacher will aim to physically come twice to the exam: about 60—90 minutes after the start of the exam, and about 60--90 minutes before the end.

Language: Answers and solutions must be given in English.

Grades: will be posted before Monday 11 April 2022. The time for the exam review will be posted on canvas as an announcement when the results are available.

You are **not** allowed to use any means of aid.
However, according to general rules printed English language dictionaries are allowed.

Please write the answer to each question (question 1, question 2, etc) on a separate sheet of paper.

Grade: The grade is normally determined as follows:

30 p ≤ grade 3 < 38 p ≤ grade 4 < 46 p ≤ grade 5 (EDA263)

30 p ≤ pass < 46 p ≤ pass with distinction (DIT641)

1 Terminology (9p)

The book discusses deception and lists three types of attacks that can lead to this threat consequence: masquerade, falsification, and repudiation. Explain what each term means and give a practical example for each one.

page 32-33: For repudiation, we spent some time to explain the concept of "non-repudiation"

2 Security Principles (15p)

During the course, we have discussed several important security design principles. Please (1) explain the meaning of the following and (2) give an example of a security mechanism discussed in the course and what the particular design principle could mean practically for this mechanism.

a) Cost of security versus cost of failure and recovery (3p)
b) Fail-safe defaults (default permit vs default deny) (3p)
c) Separation of privilege (3p)
d) Least privilege (3p)
e) Isolation (separation) (3p)

a) p 47: many examples possible
b) p 40, p 315
c) p 41 and extension to Clark-Wilson (slides, online Ch 27)
d) p 41
e) p42 and slides on OS security

3 Passwords (6p)

In computer security, it is important to compare the opportunities and capabilities of the attacker versus the defender. Sometimes the relationship of attacking / defending the system is very asymmetric between the attacker / defender. For example, when it comes to system security the attacker basically just needs to find a single vulnerability to attack the system while the defender needs to patch all vulnerabilities.

a) Give an example from password authentication where a defense mechanism could be turned into an attack. Explain the type of attack and the potential goals of the defender (that introduced the mechanism) and the attacker (that uses the mechanism as an attack). (3p)
b) There is a case when it comes to authentication with passwords that involves a "slow" hash where the defender may have the upper hand. Why? Explain and give an example. (3p)

a) For example: blocking account after three wrong passwords. Defender wants to stop guesses; new type of attacker want to cause a DoS for some users
b) offline attacks: As the passwords are hashed, the attacker needs to go through the mechanism with a clear text password to produce the hash to see if it is equivalent with what is stored in the file. The attacker will need to try many, while the authentic user should know her password, meaning that by making the mechanism (slow hash) very very slow, the defender can make it more difficult for the attacker while not penalizing normal users.

4 Operating System Security (5p)

Describe the reference monitor.

a) What is it for? (1p)
b) Name and explain at least two important requirements for a reference monitor. (2p)
c) Draw a picture showing the context in how it would be used. (2p)

See slides and online chapter 27.3

(exam continued on the next page)

5 Database security (15 p)

In the course, we discussed threats and attacks against databases. Consider the statistical database in Table 1 (after the question).

A normal user may *not* query the database on the field "Name" and may only use formulas such as:

`count(C), sum(C,Aj), median(C,Aj), max(C,Aj), min(C,Aj), etc.`

`C` is the characteristic formula, such as *`(Sex=Male) AND (Department=Math).`*

The query set, `X(C)`, is the set of records matching the characteristic formula. $|X(C)|$ is the *number of records* in this matching set. `N` is the size of the database (number of rows or records). A_j is a specific attribute, such as *`Salary`*. According to the table below, these values then give: `max(C, Aj) = 72`

a) Explain how the *query size restriction technique* can be used to protect the statistical database from an inference attack. Describe it formally using `N` and $|X(C)|$ as defined above, and the constant `k`. (2p)

b) Give a formal definition of the the *tracker attack* and describe its use in your own words. (2p)

c) Demonstrate how the *tracker attack* could be used to find the exact salary of Professor Dodd, if the attacker knew that *Dodd is the only female CS Professor*. *This is the only external information the attacker has and you cannot make any other assumptions about the values in the database*. Use `C & Aj` formally in the answer you give and list the queries used. The database is protected with the technique from (a) with `k=2`. (6p)

d) We talked about Differential Privacy in the course. Explain briefly the ideas behind differential privacy. (2p)

e) Let's say the database was protected with differential privacy instead of query size restriction from (a). Would that change the outcome of the attack? Are there any concerns of using DP when we aim to protect the salaries?
(We expect a shorter answer, yes/no with two sentence reasonable explanation) (3p)

Table 1: Database (Q5)

Name	***Sex***	***Department***	***Position***	***Salary ($K)***
Adams	Male	CS	Prof	80
Baker	Male	Math	Prof	60
Cook	Female	Math	Prof	100
Dodd	Female	CS	Prof	60
Engel	Male	Stat	Prof	72
Flynn	Female	Stat	Prof	88
Grady	Male	CS	Admin	40
Hayes	Male	Math	Prof	72
Irons	Female	CS	Student	12
Jones	Male	Stat	Admin	80
Knapp	Female	Math	Prof	100
Lord	Male	CS	Student	12
Major	Female	CS	Admin	80

(exam continued on the next page)

a) slides & videos: $k \leq |X(C)| \leq N - k$; Explain terms; if too few records are touched, no results is returned; Need restrictions on both sides.
b) The tracker attack formalizes the way to extract information to defeat the query size restriction. We create a larger set passing the query restriction with the entry, and then the tracker that excludes the entry.

C is what we are after but forbidden, single record; Rewrite it as
C = C1 AND C2 so that C1 passes query restriction and
T = C1 AND $\sim C2$ (tracker); passes query restriction

Then the set C is the following
C = C1 - T

c) Need to define a number of sets; need at least 2 records

C: Sex=Female AND Position=Prof AND Department=CS → Dodd's record = 1
C1: Sex=Female AND Position=Prof → 4 records with Dodd
C2: Department=CS → 6 records with Dodd
T: C1 AND NOT C2 → 3 records with no Dodd

SUM(C, Salary) = SUM(C1, Salary)-SUM(T, Salary) → 348 – 288 => 60

d) It is difficult to assume what the adversary knows and can use for inference.

We are assuming the most powerful adversary possible that knows all values, except the one value that the attacker wants.

Formally we say we have two databases that just differ in this single value, and we then add controlled noise to the output of the statistical query to make sure that the results are almost the same. That is, one would get "approximately" the same answer running a statistical query regardless of the database. For that reason, one needs non-deterministic functions.

The noise can hide the single user, but it can also make the utility of the database less useful (too much noise). See e)

e) In principle, the salary of Dodd would be protected if we consider two databases (with or without Dodd) and adding noise to the salary query so the results from the two databases would look "similar". In practice, this may decrease the utility as the noise added need to compensate for the highest salary (100 is much larger than 12).

6 Security Models (10 p)

You are working for a law firm with the following eight clients:

New York Times, Bank of Scotland, Scandinavian Airlines, Bank of England, Air France, Los Angeles Times, American Airlines, Bank of Wales.

The law firm is using the *Chinese Wall Model*.

a) What is the primary goal of the Chinese Wall Model? We are not looking for terms from the CIA, but a related concept. (1p)

b) Draw a figure, showing how this (general) model would look in the specific example for this law firm. Show in the picture *the three levels information* is organized into and explain them with a concrete example. (3p)

c) Define the simple security rule formally in the following way: (2p)
Simple Security Rule: A subject S can read object O only if ...

d) Alice and Bob work for the law firm. State whether the following *read* accesses (performed in the order shown here) will be accepted or denied. Use your answer in (b+c) to explain your reasoning. Structure your answer in the following way:
Answer: x) Accepted/Denied, because … (2p)

1) Alice reads a document outlining which new offices will open in 2022 for Bank of Wales.
2) Alice reads a document outlining which new offices will open in 2022 for Bank of England.
3) Alice reads a document outlining which new offices will open in 2023 for Air France.
4) Bob reads a document outlining which new offices will open in 2024 for American Airlines.
5) Alice reads a document outlining which new offices will open in 2023 for Bank of England.
6) Bob reads a document outlining which new offices will open in 2023 for New York Times.
7) Alice reads a document outlining the yearly summary of earnings / losses for Air France.
8) Alice reads a document outlining the yearly summary of earnings / losses for Bank of Wales.
9) Bob reads a document outlining the yearly summary of earnings / losses for Bank of Wales.
10) Bob read a document outlining the yearly summary of earnings / losses for American Airlines.
11) Bob reads a document outlining the yearly summary of earnings / losses for Air France.
12) Alice reads a document outlining which new offices will open in 2018 for Bank of Wales.

e) Let's add the following write accesses to the list of accesses in (d):

8.1 Alice updates (writes) to the document about outlining the yearly summary of earnings / losses for Bank of Wales. (coming after 8 but before 9).

10.1 Bob updates (writes) a document outlining the yearly summary of earnings / losses for American Airlines.

Reflect on your answer in (d), by considering the "new" accesses in the interval (8)—(11). (2p)

a) page 27-16: Conflict of Interest
b) See fig 27.6 or examples in slides: make sure all three levels are shown based on example.
c) Bottom page 27-16
d)
1 Accepted because Allice has not yet accessed any information from this conclict class
2 Denied because Allice has already accessed information from this conclict class (1)
3 A, 4 A, 5 D, 6 A, 7 A, 8 A, 9 A, 10 A, 11 D, 12 A
e)
8.1 A; 10.1 A

See page 27-17 on indirect flow of information; and put this into context.