# CS-523 Midterm Most Repeated Errors

Spring 2022

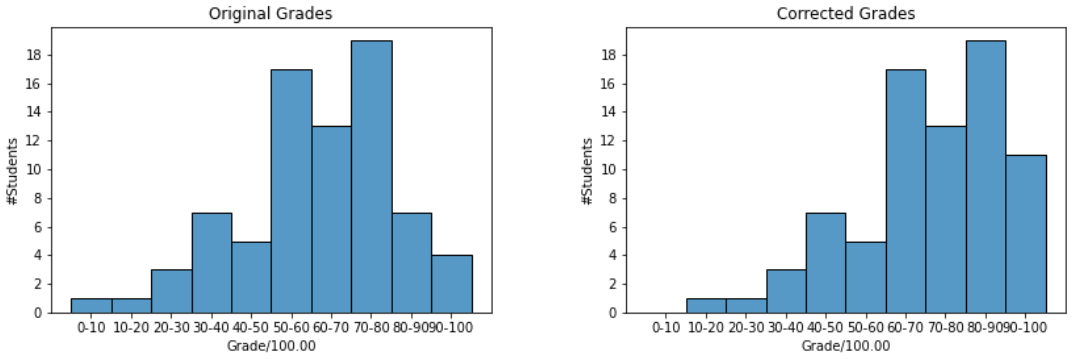


**Question 1: ABCs (Mathilde)**

- **Providing threats/answers that are way outside the security scope of the question.** You are partially evaluated on your ability to understand the problem at hand, which threats are relevant to the system proposed in the question and which are not. Of course, leaving a laptop unattended or publishing a password on the internet will lead to privacy problems, but these are not specific to the system in the question, and ZeroPass cannot help solve them. Thus, they are not relevant and cannot be used to assess your understanding of the technologies, and did not receive full points.
- **Relying on unjustified assumptions.** Whenever you assume that the adversary has some knowledge or capabilities (the adversary knows the email address / the adversary can create a credential), it is important to explain/justify how the adversary may have acquired such knowledge or capability. In order for your attack to be valid and realistic (and receive full points), each assumption has to be justified.
- **Forgetting about the unforgeability of credentials when claiming non-impersonation.** Many answers that claimed non-impersonation discussed the security of the email check, but not the one of credentials. Unforgeability of credentials is pivotal to ensure non-impersonation, and this property motivates the use of ABCs. Without it, the adversary could forge a credential with the email address of its choice and authenticate as another user, breaking non-impersonation.

- **Allowing the user to choose its username(s) without further checks during the issuance step of your protocol.** If you let the identifying attributes be user-defined, you are exposing your system to trivial impersonation, where the adversary inputs the username of the user they want to impersonate. Instead, the identifying attributes need to be checked in some way before they are signed, or they should be defined by the server.
- **Not specifying the kind of randomness.** Sampling at random means to sample from a space according to a probability distribution. Which kind of probability distribution and the space where samples are taken from can lead to very different samples and need to be specified. This is particularly important for cryptographic and privacy problems, where the

power of the adversary depends on this sampling process. Thus, even if it may seem obvious, the answer must explain what is meant by random. For this question, it was important to specify that you sample uniformly from a space large enough in order to avoid collisions.

**Question 2: Data Publishing (Theresa/Carmela)**

- **Paraphrasing slides without specifying the link to the question context**. This applies for both privacy concerns (e.g., saying membership privacy) or for defence mechanisms (e.g., input/output perturbation to obtain differential privacy). If the answer does not provide any relation to the question it is not possible to evaluate that the concepts have been understood. These answers have been given 0 points.
- **Not answering particular parts of the questions.** Many students did not answer parts of the questions. Most commonly, students did not provide a threat model under which the concern could materialize; did not provide a discussion on the details that needed to be addressed to implement defenses (e.g., how to compute sensitivity or epsilon); or did not justify why the defenses actually address the concerns. Not answering part of a question results in points reduction.
- **List inferences about population/group patterns or about events as privacy concerns.** Many students refer to an adversary that learns about group characteristics, e.g. "the adversary can learn that *users in Vaud* travel more on the weekend", as a privacy concern. This is not a privacy violation. It is a valid statistical insight drawn from the data by an honest analyst. It does not violate the privacy of any individual.
  Other answers mentioned "the adversary can learn that a large event happened". Again, this is not a privacy breach for any user, but a typical learning from data analysis.
  Privacy concerns refer to facts learned about individual *users*.
- **Describing an attack instead of a privacy concern and a threat model.** The question did not ask how one could extract information, but what privacy concerns would arise from the publication of histograms -- what could users fear others learned about them; and a threat model -- which attacker could materialize the privacy concerns: what are their capabilities and background knowledge. Answers that just provided an attack got reduced points, even 0 if none of the information asked by the question could be clearly identified.
- **Computing sensitivity wrongly for output perturbation.** The sensitivity refers to the impact of one user on the published data. Data is published per canton, per week, per day. Thus, we care about the influence a user can have on one day. On one day, a user can have at most 24 1-hour trips. The sensitivity is 24.
  We also gave points to:
  - Students that then grouped the sensitivity of the week (24*7) and accounted for this when discussing how to configure epsilon (see point on epsilon below).
  - Students that explicitly assumed that users don't make more than X 1h-trip per day, and thus sensitivity is X. We reduced points for other incorrect calculations.
- **Not explaining how epsilon is distributed among days.** While the designer chooses an overall epsilon for the system, it does not have to be exactly the epsilon that is used in the Laplacian parametrization. The output of the Laplacian may be distributed among different operations inside the system. In this case, there is a certain amount of noise that

has to be added per day to satisfy the overall epsilon. How to choose the Laplacian parametrization to obtain the desired epsilon was part of the expected answer. Not discussing it resulted in points reduction.

- **Using mechanisms that assume the existence of a database for Part 3**. Some answers to part 3 propose to use k-anonymity, or output perturbation. To use those mechanisms the data has to *already be collected.* But the question specifies that Paula has to create a "method for the app to collect the data from users and produce the statistics in a privacy-preserving way." Thus, those methods are not valid.
  Besides not answering the question, if Paula has access to a database, it means Swüber has access to the data, and one of the requirements of the question is broken.
  Correct solutions would be the use of input perturbation (e.g., randomized response) or anonymous reporting (e.g. based on ABCs and anonymous communications).

**Question 3: Anonymous Communication (Klim)**

- **Stating that increasing $p_f$ improves anonymity towards EPFL (Part 1).** The first step in the Crowds protocol is sending the message to a randomly selected crowd member (including the author). Since this selection is done with probability equal to 1/N for each user, this step ensures that all participants have the same probability to be the source of the message from the point of view of the server (including the sender, who can also be chosen as first recipient). This is regardless of $p_f$. Increasing $p_f$ affects privacy with respect to internal adversaries, i.e. malicious users.
- **Stating that the modification to the Crowds protocol reduces the anonymity set without any proof or explanation (Part 4).** The statement that a modified Crowds network *of the same size* as a normal Crowds network results in a smaller anonymity set with respect to EPFL without a justification did not result in any points. This statement on itself is not enough for us to evaluate your understanding of the system.
- **Arguing that the connected part of the network is smaller after the modification in Part 4 (communicate only in Bluetooth range).** First of all, if the network graph is not connected, it is not one network anymore, you are dealing with two or more separate networks. If you state it like there are separate networks, you must consider the properties of each network separately. Second, if you assume that these "micronetworks" consist of the users concentrated in certain areas, then this situation is not different from the part 3. In fact, the conclusion that the networks will be *smaller* seems implausible because eliminating the total-connectivity requirement from the part 3 should increase rather than decrease the number of possible reachable nodes.
- **Considering anonymity in the presence of malicious users (Part 4).** It is clearly stated in the question that we ask about "anonymity towards the server", therefore no concerns about malicious *users* should be considered. Answers that evaluated privacy with respect to malicious users received no points.

**Question 4: Location Privacy (Bogdan)**

- **Not identifying the key differences between geo-indistinguishability and distortion privacy.** Both notions can be seen as measures of privacy against an adversary that aims to

infer the real location behind an obfuscated output. However, they measure privacy in different ways. Distortion privacy measures average success of a predefined inference adversary, thus incorporating the adversary's *strategy* and *background knowledge*. Geo-indistinguishability does not explicitly mention an adversary, but it limits the worst-case leakage of the mechanism regardless of the concrete adversary's strategy of background knowledge. Failing to identify *at least* the fact that distortion privacy requires an explicit model of adversary's background knowledge, whereas geo-indistinguishability does not, resulted in a partial grade.

- **Stating that bad geo-indistinguishability necessarily implies bad distortion privacy.** In the worst case, bad geo-indistinguishability necessarily leads to easy inference of true location. Distortion privacy as defined in the question, however, is computed for a given model of an adversary. This adversary does not need to be able to perform such easy inference. For instance, they could have poor background knowledge or they could use a poor inference strategy. As an edge case, consider an adversary that outputs random guesses. Thus, it is possible to have bad geo-indistinguishability and good distortion privacy. Stating otherwise resulted in a reduced grade for Part 2.

- **"Alice should use distortion privacy if an adversary is strong and has significant background knowledge."** This particular formulation did not get the full grade for Part 3. This is because Alice can and should only use distortion privacy when she *can obtain a model of adversary's background knowledge.* Even if a considered adversary is strong, it might be hard to know what exactly they know.