CS-523 Final Exam 2022

Most Repeated Errors

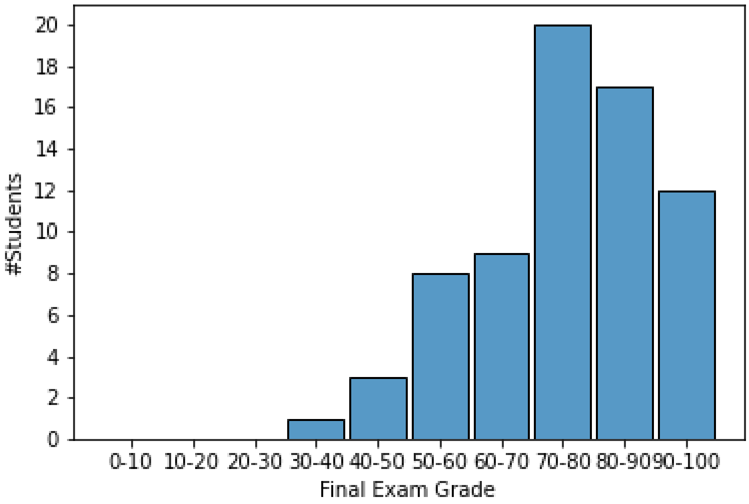

Figure: Distribution of Final Exam Grades Across Class

General:

- It is better to stick with just answering the question. We subtracted points for obviously wrong claims in the answers.
- We encountered some incomplete answers that failed to provide requested arguments or failed to provide a clear description of a design or attack. In these cases we awarded partial points only.

Question 1 - SMC

MRE 1: Leaking information. Some solutions revealed more information than required by the PSI functionality, yet did not acknowledge it, or failed to explain why this leakage is ok. For instance, if your solution leaked the number of items in each bin, or did not argue how this information was protected, we awarded partial points.

MRE 2: Assuming telepathy. Solutions that rely on symmetrical behaviour (ordering, binning, encoding, …) without communication between parties were awarded partial points only. Furthermore, the distribution of the server and client sets is not the same so problems are likely to occur when an arbitrary decision is made based on one party’s set. Example: If each party picks a binning algorithm such that there are the same number of phone numbers per bin, it is unlikely that the client’s bins will match the server’s.

Question 2 - Anonymous Communication and Differential Privacy

MRE1: Claiming that the mix network provides perfect sender anonymity. Collusion between the first and the last server means that they together can identify the set of possible senders, even if an honest server in the middle shuffles the messages. Answers that failed to clearly identify this limitation, e.g., by stating that the system provides perfect anonymity, were awarded partial points.

MRE2: Assume broken crypto and then attack the system. If you assumed that the encryption scheme was particularly bad (e.g., deterministic or lacking integrity checks) and used that to attack the system, we only awarded partial points.

MRE3: Incorrect/unrealistic computation of sensitivity. The sensitivity is derived from the difference between the setting where Bob sends to Alice, and where he doesn't. The sensitivity therefore does not depend on the total number of users or the actual number of messages sent. It also should not depend on the mean number of messages from Bob to Alice, sensitivity is a worst-case measure. We subtracted points for incorrect computation of sensitivities.

MRE4: Failing to take into account the effect of multiple rounds on the DP budget. There were two ways to analyse this problem:

- Consider all rounds together, and let the sum of messages be the output. You'd then use Laplace noise (with a suitable sensitivity that takes into account the 24 rounds) to compute *the total number of dummy messages to Alice to add for this day.* Your answer must then include how this total number of messages is distributed over the 24 rounds.
- Consider rounds one by one, and use sequential composition to determine the privacy budget. You need to use Laplace noise with a per-round sensitivity to compute *the number of dummy messages to Alice per round.*

Answers that failed to explain distribution of messages or composition, were awarded partial points.

MRE5: Add a fixed number of messages. Adding a fixed number of messages to pass the threshold does not work. First, this approach is not differentially private. Two, a strategic adversary would simply adjust the threshold to take into account the added messages.

MRE6: Depending on hidden information. A solution cannot depend on information such as "whether Bob is online" or "the number of messages to Alice". The hidden mix server in the middle does not have access to this information.

Question 3 - Privacy Engineering

MRE 1: Suggest a privacy-enhancing modification that reduces the functionality of the system/affects functional requirements: In Part 2, some answers suggested privacy-enhancing modifications that affected the functional requirements of the ranking system. We only awarded partial points to these answers because they did not take into account one of the very important privacy-by-design goals presented in class: To design systems that maximise privacy *without reducing the core functionality of the system or changing functional requirements.*

MRE 2: Suggest a privacy-enhancing modification to System 1 that does not address the centralisation of trust problem: In Part 2, some answers suggested modifications to System 1 that did not address the root cause of the privacy problem and thus did not reduce privacy risks.
The main problem of this system design is the centralisation of trust. Answers that suggested modifications, such as to locally encrypt data with a key *provided by a central authority*, that still required applicants to trust a central authority to act truthfully to preserve their privacy, were not awarded full points.
Only answers where, for instance, the key generation scheme included some kind of decentralisation and led to a reduction of trust assumptions were given full points.

Question 4 - Tracking

MRE 1: Proposing "using Tor browser/another tool" as the defence strategy and proceeding to outlining the generic functioning and drawbacks of the tool. Per the given template, we expected to see a defence strategy expressed in terms of approaches, not existing tools. Proposing to "use Tor", however, would get the full grade if the answer detailed what are the features of Tor that mitigate browser fingerprinting, how they work, and what are the drawbacks of those particular features. If the answer instead described the IP anonymization capabilities of Tor, which are not related to browser fingerprinting, or drawbacks of Tor that are not related to browser fingerprinting such as increased network latency, the answer did not receive full points.

MRE 2: Not detailing which concrete browser features should be modified in the strategy. If an answer said that the "browser fingerprint" should be homogenised or randomised without any concrete examples of what particular components of the fingerprint should be changed and how, it did not get the full grade.

Question 5 - Censorship resistance

MRE 1: Proposing solutions requiring additional software, if it is mentioned as a disadvantage in 1.b or 1.c

Part 2 required to propose a solution which does not have disadvantages listed in Part 1. Installing additional software is a legitimate disadvantage for Tor or Telex. However, solutions such as “covert channels” or CloudTransport require the installation of additional software as well. Some answers argued that such applications might be easier to use for “non tech-savvy people”, than Tor Browser. While it might be true, it is not an inherent disadvantage of Tor nor Telex. For example, one can design an app using Tor which is reduced to one button: “Get the info”. The real disadvantage is a necessity to install *something*, and it is not eliminated with replacing Tor protocol with another custom protocol. This is why such answers did not receive full points. Note, that if this disadvantage is not listed in Part 1, proposals containing additional software are correct.