Advanced Topics on Privacy Enhancing Technologies
CS-523
Anonymous Communication Exercises

1 Cryptographers' Dinner

Consider a DC-networks scenario with a total of n cryptographers. Out of these, k cryptographers dislike each other and thus are guaranteed not to collude. The cryptographers decide to have a shared key setup and arrange themselves as a graph (a cryptographer is a node in the graph, edges between nodes indicate a shared key). Since a complete graph is expensive due to the large number of keys, the cryptographers form a trusted root clique structure. The structure is as follows: the k cryptographers form a root clique and share keys among themselves. All the other cryptographers create shared keys with each of the root clique members.

1. If all the members outside the root clique decide to collude, how does that affect the anonymity of the root clique?
2. If $k - 1$ cryptographers finally resolve their differences and decide to collude, how does that affect the anonymity of the other nodes (assuming the other nodes do not collude with the $k - 1$ nodes)?
3. What is the total number of shared keys required in the original setup?

The cryptographers decide to redesign their network. They form a structure as follows: they arrange themselves in cliques of k members. Within each clique, members create shared keys with every other member. There are l cliques in total. All the cliques are then arranged in a ring structure. Every clique selects a node as a leader, and leaders of each clique share keys with their immediate neighbors in the ring.

4. Within a clique, if m members decide to collude, how does that affect the anonymity of the clique?
5. If two cliques decide to collude, how does that affect the anonymity of the entire setup?

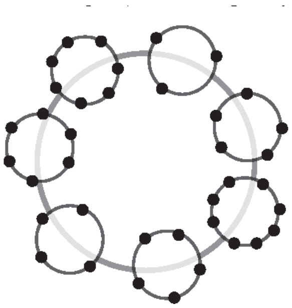

Figure 1: Topology for the second part of Question 2.

6. What is the total number of shared keys required in this setup?

2 Statistical Disclosure Attack

Alice uses a mix with a total of N participants. The mix works in rounds: in each round, the mix waits for $K < N$ users to send their messages. A message can be sent to any participant in the network.

Alice uses the mix multiple times. She is wondering how likely an adversary is to figure out that she is talking to Bob, one of the people she communicates with using the mix. The adversary that Alice is concerned about can observe network traffic: He knows who participates in each round, who sends messages, and who receives the messages.

Consider the following probabilistic model of each round in which Alice participates:

- The probability that the receiver i gets a message from Alice:

$$P[i \leftarrow \mathit{Alice}] = a_i$$

- The probability that the receiver i gets a message from a sender $j \neq \mathit{Alice}$:

$$P[i \leftarrow j] = u_{ij}$$

Over T rounds in which Alice participates, the adversary keeps track of all receivers' statistics $\bar{O}_i = \frac{1}{T}\sum_{t=1}^{T} o_i^{(t)}$, where $o_i^{(t)} \geq 0$ indicates the number of messages that i received in round t. Because all rounds are independent, these observations can be thought as i.i.d samples from a random variable $O_i = \mathbf{1}_{i \leftarrow \mathit{Alice}} + \sum_{j \neq \mathit{Alice}} \mathbf{1}_{i \leftarrow j}$.

1. In terms of the probabilistic model given above, what is the expected number of messages that the receiver i gets in one round, $\mathbb{E}[O_i]$? In T rounds?

2. Assume that Alice participated in enough rounds so that the Law of Large Numbers (LLN) applies to the average statistics: $\bar{O}_i \approx \mathbb{E}[O_i]$. How can the adversary estimate a_i from a given model of u_{ij}?

3. Suppose that Alice communicates *only* with Bob. The threshold of the mix is $K = 30$, and the total number of participants is $N = 290$. The adversary assumes that all users other than Alice send messages uniformly to any other user: $u_{ij} = \frac{1}{N}$. What is the expected number of messages that Bob and any other recipient receives in one round?

   Over $T = 1000$ rounds, the adversary observes that $\bar{O}_{Bob} = 0.9$. Adversary also suspects that Alice might be talking to Carol. Carol's statistic is $\bar{O}_{Carol} = 0.1$. What are the adversary's estimates for Alice's probability to send messages to Bob and to Carol?

3 Getting Through a Crowd

The Crowds system may have a high latency depending on its parameters, and some extensions like "always down or up" (ADU) aim to improve this latency. In this exercise, we study the latency and privacy of these extensions.

1. A good measure for the latency of messages in crowds is the number of IP hops l. Compute the expected value and variance of the number of hops in crowds.

In an ADU with parameters (e, l, h, m), the sender chooses a random number $u \overset{\$}{\leftarrow} [1, m]$ before sending the message. If $i \in [1, e] \cup [m-e, m]$ the sender directly sends the message. If the number is less than a lower bound $e < u \leq l$ or higher than an upper bound $h \leq u < m - e$ then the sender proceeds with AD or AU respectively. If the number is in the middle range $l < u \leq h$ then the sender chooses the direction randomly from AD and AU and forwards the mode and u with the message to the next node.

2. Compare the privacy of ADU with AD and AU.

3. Compare the privacy of ADU with crowds.

An alternative to ADU, is the "random always down or up" (RADU) algorithm is which does not send the direction (AD or AU) with the message. Each node on the path receives the message with a random number u. Similar to ADU, each node uses the ADU algorithm to decide the direction but with fresh randomness.

4. Compare the latency of ADU with RADU.

5. Compare the privacy of ADU with RADU.