Week 7 Exercises (ECE 598 DA)

Exercise (Completeness and Soundness in Graph Isomorphism Protocol):

In the graph-isomorphism zero-knowledge protocol, the verifier sends a random challenge $b \in \{0, 1\}$, and the prover must reveal σ (if $b = 0$) or $\sigma \circ \pi^{-1}$ (if $b = 1$). Prove the following:

1. If $G_0 \cong G_1$ via isomorphism π, the verifier always accepts (completeness).
2. If $G_0 \ncong G_1$, then any prover can convince the verifier with probability at most 1/2 per round (soundness).

Exercise (Simulator Construction for Honest-Verifier Zero Knowledge): Construct a simulator for the honest-verifier case of the Graph Isomorphism protocol and explain why the simulated transcript is indistinguishable from a real one.

Exercise (Fiat–Shamir Transformation): Describe how the Fiat–Shamir heuristic converts an interactive ZK protocol into a non-interactive one. Apply it to the Schnorr identification protocol and explain what security assumption is required.

Exercise (Zero-Knowledge in the Ali Baba Cave Analogy): In the Ali Baba cave story, if Peggy does *not* know the secret word, what is the probability that she can correctly respond to Victor's challenge in one round? What happens after 10 independent rounds?

Exercise (Completeness, Soundness, Zero-Knowledge): Explain how each property (completeness, soundness, and zero-knowledge) is manifested in the Ali Baba cave example.

Exercise (ZK Proof for the Discrete Logarithm Relation): Design a zero-knowledge protocol proving knowledge of x such that $y = g^x \pmod p$. Formally verify completeness, soundness, and zero-knowledge.

Exercise (Challenge-Response Repetition): Why does repeating a 2-message ZK protocol (like Graph Isomorphism) multiple times reduce the soundness error exponentially? Quantify how many rounds are needed for a cheating probability below 2^{-80}.

Exercise (Simulation Soundness): Explain why, in computational zero-knowledge, the existence of a polynomial-time simulator implies that any efficient adversary cannot extract additional knowledge from transcripts beyond the fact that the statement is true.

Exercise (ZK for Machine Learning Models): Discuss how a zero-knowledge proof could be used to prove that a machine learning model's prediction $f(x)$ is correct without revealing the model parameters or the input x. Hint: think about zk-SNARKs or zkML protocols.