MTAT.07.003 Cryptology II
Spring 2010 / Exercise session XI

1. Let π be a sigma protocol such that the challenge β is chosen uniformly from an n element set $\mathcal{B}$. Let $\mathcal{S}_0$ be a t_0-time algorithm that simulates perfectly the protocol transcript (α, β, γ) when both parties are honest:

$$\Pr\left[(\alpha,\beta,\gamma) \leftarrow \mathcal{S}_0\right] = \Pr\left[\mathcal{V}, \mathcal{P} \text{ create } (\alpha,\beta,\gamma)\right] \ .$$

   (a) Prove that the following black-box simulator

$$\begin{array}{l} \mathcal{S}^{\mathcal{V}_*} \\ \left\lceil \begin{array}{l} \text{For } i \in \{1,\ldots,k\} \text{ do} \\ \quad \left\lceil \begin{array}{l} (\hat{\alpha},\hat{\beta},\hat{\gamma}) \leftarrow \mathcal{S}_0 \\ \beta \leftarrow \mathcal{V}_*(\alpha) \\ \mathsf{if}\ \beta = \hat{\beta}\ \mathsf{then}\ \mathbf{return}\ \mathcal{V}_*(\hat{\gamma}) \end{array} \right. \\ \mathbf{return}\ \mathsf{Failure} \end{array} \right. \end{array}$$

   provides perfect simulation when it does not halt with $\perp$:

$$\Pr\left[\psi \leftarrow \mathcal{S}^{\mathcal{V}_*} | \neg \mathsf{Failure}\right] = \Pr\left[\psi \leftarrow \mathcal{V}_*^{\mathcal{P}}\right] \ .$$

   Establish also the corresponding failure probability $\Pr[\mathsf{Failure}]$ and compute the total running time of $\mathcal{S}^{\mathcal{V}_*}$ such that $\Pr[\mathsf{Failure}] \leq \varepsilon$. When are the running times of $\mathcal{V}_*$ and $\mathcal{S}^{\mathcal{V}_*}$ comparable?

   (b) There is a trade-off between simulation overhead and soundness, since the knowledge error can be expressed as $\kappa = \frac{1}{|\mathcal{B}|}$. Thus, by decreasing the set $\mathcal{B}$ we also increase acceptance probability for malicious provers. To compensate the effect, we must sequentially run several instances of π. Let κ be the desired knowledge error and let ε be the desired bound on simulation failure $\Pr[\mathsf{Failure}]$. What is the minimal number of rounds we need, if we require that $\mathcal{S}^{\mathcal{V}_*}$ can run only $\mathbf{poly}(\log_2(1/\varepsilon))$ times slower than $\mathcal{V}_*$?

2. One possibility to convert a sigma protocol into a zero knowledge proof is to use commitments to fix β value before $\mathcal{V}_*$ receives α value. As a result, we can use the following black-box simulation strategy

$$\begin{array}{ll} \mathcal{S}^{\mathcal{V}_*} & \mathcal{K}^{\mathcal{V}_*}(\mathsf{pk}, c) \\ \left\lceil \begin{array}{l} \mathsf{pk} \leftarrow \mathsf{Gen} \\ c \leftarrow \mathcal{V}_*(\mathsf{pk}) \\ \hat{\beta} \leftarrow \mathcal{K}^{\mathcal{V}_*}(\mathsf{pk}, c) \\ (\alpha,\hat{\beta},\gamma) \leftarrow \mathcal{S}_1(\hat{\beta}) \\ \mathbf{return}\ \mathcal{V}_*(\alpha,\gamma) \end{array} \right. & \left\lceil \begin{array}{l} \text{For } i \in \{1,\ldots,k\} \text{ do} \\ \quad \left\lceil \begin{array}{l} \alpha \leftarrow \mathcal{S}_2, d \leftarrow \mathcal{V}_*(\alpha) \\ \beta \leftarrow \mathsf{Open}_{\mathsf{pk}}(c,d) \\ \mathsf{if}\ \beta \neq \perp\ \mathsf{then}\ \mathbf{return}\ \beta \end{array} \right. \\ \mathbf{return}\ \perp \end{array} \right. \end{array}$$

where $\mathcal{S}_1$ simulates perfectly the protocol messages given $\hat{\beta}$ and $\mathcal{S}_2$ simulates perfectly the first message α.

(a) Prove that the simulation is perfect when there are no simulation failures. The simulation fails if one of two events happens: (i) the knowledge extraction fails $\mathcal{K}^{\mathcal{V}_*}(\mathsf{pk}, c) = \perp$ and the verifier $\mathcal{V}_*$ outputs a valid decommitment in the final run; (ii) the knowledge extraction succeeds $\mathcal{K}^{\mathcal{V}_*}(\mathsf{pk}, c) \neq \perp$ but $\hat{\beta} \neq \mathsf{Open}_{\mathsf{pk}}(c, d)$.

(b) Let $\varepsilon_1 = \varepsilon_1(\mathsf{pk}, c)$ be the probability that $\mathcal{V}_*$ manages to open c. Let ε be the desired failure probability for knowledge extraction procedure $\mathcal{K}^{\mathcal{V}_*}(\mathsf{pk}, c)$. Find out the corresponding value of k.

(c) Let Ω_{good} consist of all reachable pairs (pk, c) such that $\varepsilon_1(\mathsf{pk}, c) \geq 2\varepsilon$. What is the total simulation failure probability if the value of k is chosen according to the bound obtained in the part (b)? How many times $\mathcal{S}^{\mathcal{V}_*}$ is slower than $\mathcal{V}_*$?

(d) Compare the results with the previous exercise. Can we construct a simulator that is only $\mathbf{poly}(\log_2(1/\varepsilon))$ slower than $\mathcal{V}_*$?

3. Many challenge response protocols can be converted to zero-knowledge proofs if one can guarantee that the verifier knows the final answer. Witness-indistinguishable sigma protocols can be used for that purpose.

(a) As a first example, consider the zero-knowledge proof for quadratic non-residuosity presented in the lecture. Let the game $\mathcal{G}_0$ model the behaviour of honest verifier with $\beta = 0$ and let the game $\mathcal{G}_1$ model the behaviour of honest verifier with β. Write down these games under the assumption that v is quadratic residue. Simplify the games $\mathcal{G}_0$ and $\mathcal{G}_1$ until you have reached to the same game.

(b) More generally, let x be the public input and w the corresponding witness. Also, assume that the challenge α is computed as $\alpha(\beta; r)$ where $\beta \leftarrow_u \{0, 1\}$ and $r \leftarrow_u \mathcal{R}$. Finally, let

$$\text{POK}_\beta\left[\exists r : \alpha = \alpha(\beta, r)\right] \equiv \text{POK}_r\left[\alpha = \alpha(0; r)\right] \vee \text{POK}_r\left[\alpha = \alpha(1; r)\right]$$

be the corresponding disjunctive proof. Prove that the corresponding protocol transcript coincides for both potential witnesses $\beta \in \{0, 1\}$ if the input x is incorrectly formed:

$$\exists r : \alpha = \alpha(0; r) \wedge \exists r : \alpha = \alpha(1; r) \ .$$

Explain also why and unbounded prover can distinguish proves if this condition does not hold. What happens if for some incorrect input x only one disjunct can be satisfied?

4. The construction of sigma protocols for complex relations can be cumbersome and inefficient. Hence, it is often advantageous to use more simplistic alternatives for tightening challenge-response protocols. Let $\alpha = \alpha(\beta; r)$

be the challenge message where $\beta \xleftarrow{u} \mathcal{B}$ is the expected response and $r \xleftarrow{u} \mathcal{R}$ is the masking randomness. Then verifier can prove that he or she knows β by revealing β and r. Of course, the prover must commit his or her response before receiving β and r.

(a) As an example, consider the challenge-response protocol that proves the possession of secret key (ability to decrypt messages). Recall that the challenge is computed as $\alpha \leftarrow \mathsf{Enc}_{\mathsf{pk}}(\beta)$ and the response is computed as $\beta \leftarrow \mathsf{Dec}_{\mathsf{sk}}(\alpha)$. Construct the corresponding zero-knowledge proof using the idea explained above.

(b) Construct a simulator $\mathcal{S}^{\mathcal{V}_*}$ that first extracts β and then uses this value to pass the proof. Is there any conceptual difference compared to the construction given in Exercise 2? Can this simulator construction be used for any challenge response protocol?

(c) For clarity, assume that the cryptosystem in question is RSA with OAEP padding. Would you use the construction based on sigma protocol or the construction given above?

5. There is an important distinction between the strengthening of sigma protocols with commitments and with a coin-flipping protocol for $\beta \xleftarrow{u} \mathcal{B}$. Namely, the strengthening with coin-flipping protocol preserves proof of knowledge property, whereas the use of commitments does not. Analyse the soundness guarantees of both protocols under the assumption that the knowledge bound of a sigma protocol is κ, i.e., if a prover succeeds in the proof with probability strictly more than κ then by generating and analysing all protocol runs, we can extract the secret witness w.

(a) Let $(\mathsf{Gen}, \mathsf{Com}, \mathsf{Open})$ be a perfectly hiding commitment scheme such that randomness r used to compute commitments is uniformly sampled from a set $\mathcal{R}$. Then it is straightforward to prove that for any plausible value pk of public parameters sets

$$\mathcal{R}_{c,m} = \{r \in \mathcal{R} : (c, d) \leftarrow \mathsf{Com}_{\mathsf{pk}}(m; r)\}$$

have the same number of elements for any $c \in \mathcal{C}$ and $m \in \mathcal{M}$. Use this fact to prove that if a malicious prover succeeds with probability ε against the zero-knowledge protocol, then there exist possibly inefficient prover that succeeds with the same probability against the sigma proof. Why is this sufficient for soundness?

Hint: Consider first the simplest case when $|\mathcal{R}_{c,m}| = 1$.

(b) In brief, a coin-flipping protocol is ε-secure against malicious provers if there exists an efficient simulator Sim that given $\beta \xleftarrow{u} \mathcal{B}$ simulates the protocol for the prover $\mathcal{P}_*$ such that the revealed challenge after the coin-flipping protocol is β and the output distributions of $\mathcal{P}_*^{\mathcal{V}}$ and $\mathsf{Sim}^{\mathcal{P}_*}$ are statistically ε-close. Prove that any prover $\mathcal{P}_*$ against the zero-knowledge protocol can be converted to an adversary $\mathcal{P}_\sigma$

against the sigma protocol so that the advantage does not drop and the running time of $\mathcal{P}_\sigma$ is comparable to $\mathcal{P}_*$. How is the efficiency of $\mathcal{P}_\sigma$ connected to the efficiency of the simulator Sim.

Hint: Note that ε-closeness holds also for the provers $\mathcal{P}_{**}$ that run $\mathcal{P}_*$ and then output 1 if the transcript (α, β, γ) generated by $\mathcal{P}_*$ is valid and 0 otherwise.

6. Let SquareFree denote all square free integers, i.e., all integers that do not divisible by some prime square. Let ProdOfTwoPrimePowers denote all integers in the form $p^a q^b$ where $p, q \in \mathbb{P}$. Then the classical zero-knowledge proof $n \in$ RsaModulus consists of following steps.

   1. Prove that $n \in$ SquareFree.
   2. Prove that $n \in$ SquareFree is a product of two prime powers.

   In the following, we consider only the simplest proof that $n \in$ SquareFree.

   (a) Let $\phi(n) = |\mathbb{Z}_n^*|$ denote the Euler totient function. Prove that if $\gcd(n, \phi(n)) = 1$ then $n \in$ SquareFree. Describe the class of square free numbers such that $\gcd(n, \phi(n)) \neq 1$.

   (b) Let SquareFree* be the class of all square free numbers such that $\gcd(n, \phi(n)) = 1$. Construct a challenge-response protocol for proving $n \in$ SquareFree*.

   Hint: If $\gcd(n, \phi(n)) = 1$ and prover knows the factorisation, then a prover can efficiently find n-th roots.

   (c) Let $X_{\text{bad}} = \{x^n : x \in \mathbb{Z}_n^*\}$ be the set of n-th powers. Show that if $\gcd(n, \phi(n)) = d > 1$ then $|X_{\text{bad}}| = |\mathbb{Z}_n^*| / d$. Conclude that a malicious prover can win the challenge-response protocol with probability at most $\frac{1}{d}$.

   (d) Construct a corresponding zero-knowledge proof for $n \in$ SquareFree*.