0368.4162 Foundations of Cryptography December 28, 2017

Problem Set 5 - Reference Solution

Due: January 11

1. Consider the GMW zero-knowledge proof system for 3COL when repeated sequentially $t = n \cdot |E|$ times. Let $G = (U, E)$ be a graph and let P^* be a (w.l.o.g deterministic) prover that manages the convince the verifier V of accepting with probability $n^{-O(1)}$.

   We will prove that we can efficiently extract a legal 3-coloring of G given oracle access to P^*. Here oracle access means that we can *rewind* P^*. Formally, we are given access to *the next message function of* P^* that given a transcript of all prover-verifier messages up to some point, generates the next prover message. In particular, any partial interaction in the first $i-1$ rounds, can be continued in different ways, by having the extractor choose different verifier messages as the $i+1$st message.

   Consider the random process of running t interactions with the prover (where at each one the verifier sends a random message). For $i \in [t]$, let p_i be a random variable that is the probability, over V's coins, that V accepts in the ith interaction, conditioned on the first $i-1$ interactions (this random variable becomes fixed once we fix the first $i-1$ interactions). Let G_i be the event that $p_i > 1 - \frac{1}{|E|}$.

   (a) (15 pts) Prove that the probability that in t interactions the prover convinces the verifier of accepting, but non of the events $G_1, \ldots, G_t$ occurred is bounded by $2^{-\Omega(n)}$.

   Solution: Let S_i be the event in which the prover convinces the verifier of accepting in the ith interaction. Then, by repeated conditioning,

$$\Pr\left[\bigwedge_{i\in[t]} S_i \bigwedge_{i\in[t]} \bar{G}_i\right] \leq \prod_{i\in[t]} \Pr\left[S_i \,\middle|\, \bigwedge_{j\leq i} \bar{G}_j \bigwedge_{j<i} S_j\right] \leq \left(1 - \frac{1}{|E|}\right)^{n|E|} \leq 2^{-n} \ .$$

   (b) (10 pts) Deduce that in t interactions the probability that for some i, the event G_i occurs is $n^{-O(1)}$.

   Solution: Using the previous item and the fact that the verifier is overall convinced with probability $n^{-O(1)}$, we have

$$\Pr\left[\bigvee_{i\in[t]} G_i\right] \geq \Pr\left[\bigwedge_{i\in[t]} S_i\right] - \Pr\left[\bigwedge_{i\in[t]} S_i \bigwedge_{i\in[t]} \bar{G}_i\right] \geq n^{-O(1)} - 2^{-n} \geq n^{-O(1)} \ .$$

   (c) (15 pts) Prove the existence of the required extractor.

   Solution: The extractor will sample t sequential interactions, and then attempt to extract from each one of them, by rewinding the prover, and asking it to reveal for every choice $e \in E$. We know that with probability $n^{-O(1)}$, there is an interaction $i \in [t]$, where the verifier gives a valid answer for more than $(1 - \frac{1}{|E|})|E| = |E| - 1$ edges, which means it gives a valid answer on all edges. By the binding of the commitments, we know that the resulting coloring is consistent and is valid. The extractor's success probability can be amplified to $1 - 2^{-n}$ by repeating the above $n^{O(1)}$ times.

2. Consider an auction with a seller S party and three participants A, B, C with inputs $a, b, c \in [2^n]$ representing they're bids. They run an MPC protocol (against malicious parties) for the function that gives S the identity and the bid of the highest bidder. Assume that b and c are chosen at random.

   (a) (15 pts) Prove that the probability that a corrupted A^* outputs b is negligible.

   Solution: Let SIM be the ideal world simulator for A^*. Let us denote the output of A^* in the real world by y and SIM's output in the ideal world by $\tilde{y}$. Then, there exists a negligible μ such that for all $n \in \mathbb{N}$ and any choice of inputs $a, b, c \in \{0,1\}^n$,

$$\Pr[y = b] \leq \Pr[\tilde{y} = b] + \mu(n) \ .$$

   Otherwise, we get a distinguisher between the ideal and real worlds distributions that simply test whether A's output y, in the real world, and $\tilde{y}$, in the ideal world, equal b.
   Thus, it suffices to show that for any SIM

$$\Pr[\tilde{y} = b \mid b \leftarrow \{0,1\}^n] \leq 2^{-n} \ .$$

   This holds since the view of SIM in the ideal world is completely independent of B's bid b.

   (b) (15 pts) Prove that the probability that a corrupted A^* wins with bid $1 + \max\{b, c\}$ is negligible.

   Solution: Again, let SIM be the ideal world simulator for A^*. Let us denote the output of the (honest) seller S in the real world by y and her output in the ideal world by $\tilde{y}$. As before, there exists a negligible μ such that for all $n \in \mathbb{N}$ and any choice of inputs $a, b, c \in \{0,1\}^n$,

$$\Pr[y = (A, \max\{b, c\} + 1)] \leq \Pr[\tilde{y} = (A, \max\{b, c\} + 1)] + \mu(n) \ .$$

   Otherwise, we get a distinguisher between the ideal and real worlds distributions that simply test whether the output y, in the real world, and $\tilde{y}$, in the ideal world, equal $(A, \max\{b, c\} + 1)$.
   Thus, it suffices to show that for any SIM

$$\Pr[\tilde{y} = (A, \max\{b, c\} + 1) \mid b, c \leftarrow \{0,1\}^n] \leq O(2^{-n}) \ .$$

   Indeed, $\tilde{y} = (A, \max\{b, c\} + 1)$ if and only if the simulator SIM in the ideal world sends $\max\{b, c\} + 1$ to the trusted party. Since the view of SIM in the ideal world is independent of b and c, this happens with probability $O(2^{-n})$. This holds since the view of SIM in the ideal world is completely independent of B's bid b.

3. In the following question, addition and multiplication are done modulo 2.

   (a) (15 pts) Consider the following m-party randomized function mapping m pairs of bits to m bits:

$$(a_1, b_1), \ldots, (a_m, b_m) \longmapsto c_1, \ldots, c_m \ ,$$

   where $c_1, \ldots, c_m$ are uniform in $\{0,1\}^m$ subject to $\sum_{i \in [m]} c_i = \left(\sum_{i \in [m]} a_i\right) \times \left(\sum_{i \in [m]} b_i\right)$.
   Describe a semi-honest protocol for computing the above function, assuming a semi-honest protocol for any two-party function.

   Solution: Note that

$$\left(\sum_{i \in [m]} a_i\right) \times \left(\sum_{i \in [m]} b_i\right) = \sum_{i \in [m]} a_i b_i + \sum_{1 \leq i < j \leq m} (a_i b_j + a_j b_i) \ .$$

Accordingly, to compute the above sum, each two parties $1 \leq i < j \leq m$, will execute a two party protocol in which i learns a random value r_{ij} and j learns $r_{ji} := a_i b_j + a_j b_i + r_{ij}$. As in class, we can have i samples r_{ij} herself and inputs (r_{ij}, a_i, b_i) whereas j inputs (a_j, b_j).

Eventually, each party i sets $c_i = a_i b_i + \sum_{j \neq i} r_{ij}$.

(b) (15 pts) Use the fact that $\{+, \times\}$ is a universal set of Boolean gates to describe a semi-honest protocol for any deterministic m-party function.

Solution: The solution is similar to what we've seen in class:

i. **Input Sharing:** each party i splits its input x_i into m random shares that sum up to x_i. It broadcasts the other $m-1$ parties their relevant shares.

ii. **Evaluating Multiplication:** to evaluate a multiplication gate $c = a \times b$ over the shares, the parties execute the protocol from the previous item.

iii. **Evaluating Addition:** to evaluate an addition gate $c = a + b$ over the shares, each sets the share for the output gate to $c_i = a_i + b_i$ by adding its own corresponding shares (this steps involves no interaction).

iv. **Output Reconstruction:** when reaching the output gate, the parties send all their shares for this gate to the reconstructing party, who adds them up to obtain the output.

4. (**Bonus** 10 pts) Show that any two-message $(1, 2)$-OT (that is semi-honestly secure) implies public-key encryption.

Solution: The construction is the following:

- The key generator $G(1^n)$, runs the OT receiver R with input $i = 1$. It then sets the public-key pk to be the message that R generates, and sk to be the randomness that R used.
- To encrypt a bit b, under pk, the encryptor runs the OT sender S, with input $(\sigma_1, \sigma_2) = (b, 0)$. The ciphertext ct is then set to be the sender's message.
- To decrypt ct, the decryptor runs the receiver R with input $i = 1$ and randomness sk, and the received message ct from the sender. The decrypted message is the output of R.

The correctness of the scheme follows directly from that of the OT protocol.

We will now sketch the proof of security, which follows by a simple hybrid argument. Let us denote by $R(i)$ the random variable that describes R's message given input i. Also, denote by $S(\sigma_1, \sigma_2, m)$ the random variable that describes S's message given input (σ_1, σ_2) and receiver message m. Then, for any message $b \in \{0, 1\}$,

$$pk, E_{pk}(b) \equiv R(1), S(b, 0, R(1)) \approx_c R(2), S(b, 0, R(2)) \approx_c R(2), S(0, 0, R(2)) \ .$$

Here the first indistinguishability follows from the receiver privacy; indeed, $R(i)$ can be simulated independently of i, and thus $R(1) \approx_c R(2)$. The second indistinguishability then follows by the sender privacy; Indeed, $R(2), S(\sigma_1, \sigma_2, 0, R(2))$ can be simulated from σ_2 and independently of σ_1, and the above indistinguishability follows as a special case.

Overall, we see that $pk, E_{pk}(b)$ is indistinguishable from a distribution that is independent of b.