Statistical Zero Knowledge

CS6230: Topics in Information Security
Lecture 7 (Sep 28, 2021)
Lecture and notes: Prashant Nalini Vasudevan

In the previous lecture, we saw the definition of a notion of distance between probability distributions, and of statistical zero-knowledge proofs. We saw how the simple PZK protocol for Graph Isomorphism could be extended to a PZK protocol for the Statistical Closeness problem. More generally, the protocol works for SC^β, which is the problem of deciding whether the distributions sampled by two given circuits have statistical distance at most β, or 1.

The proof I suggested in class for the zero-knowledge property of this protocol for SC^β was flawed – see the previous lecture notes for details. The protocol, however, is indeed zero-knowledge, and for SC^β has completeness error $\beta/2$, soundness error $1/2$, and zero-knowledge error approximately β. While the completeness and soundness errors can be amplified by repetition, this only worsens the zero-knowledge error. While I am not aware of any general procedures to amplify zero-knowledge errors (though some may exist in the research literature), it turns out that there is an elegant way to get an SZK protocol with better errors for this particular problem.

In this lecture, we will how this is done. Then, we will move on to an extension of the other PZK protocol we saw – the one for Graph Non-isomorphism – to a protocol for a complement of the SC problem. We will see how this problem is, in fact, complete for the class SZK of problems that have SZK proofs. This completeness enables simple proofs of several closure properties of this class, and also of an interesting connection to cryptography.

For many of the proofs below, we will need the following facts about the statistical distance between joint distributions.

Fact 0.1. *For any distributions (X_0, Y_0) and (X_1, Y_1),*

$$\Delta((X_0, Y_0), (X_1, Y_1)) \geq \max\left[\Delta(X_0, X_1), \Delta(Y_0, Y_1)\right]$$

Fact 0.2. *For any distributions (X_0, Y_0) and (X_1, Y_1),*

$$\Delta((X_0, Y_0), (X_1, Y_1)) \leq \Delta(X_0, X_1) + \mathrm{E}_{x \leftarrow X_0}\left[\Delta(Y_0|_{X_0=x}, Y_1|_{X_1=x})\right]$$

In particular, if X_0 is independent of Y_0 and X_1 is independent of Y_1,

$$\Delta((X_0, Y_0), (X_1, Y_1)) \leq \Delta(X_0, X_1) + \Delta(Y_0, Y_1)$$

1 Polarising Statistical Distance

As indicated above, the SZK protocol for SC^β we saw earlier has zero-knowledge error approximately β. More specifically, the distance of the simulated distribution, conditioned on the simulator not failing, is β. Further, the simulator fails with at most some constant probability, which can be decreased by repetition to something exponentially small. Thus, given a security parameter 1^λ as additional input, it can simulate a distribution that has distance $\beta + 1/2^\lambda$ from the actual verifier's view.

However, in our definition of the SZK property, we required something substantially stronger – that the zero-knowledge error decrease as a negligible function of λ. So if β also happened to be a negligible function of λ, then this protocol would be actually SZK, but β is fixed by the problem description. We will see, however, that for any constant $\beta < 1$, the problem SC^β can be reduced to $\mathrm{SC}^{2^{-\lambda}}$ in about $\mathsf{poly}(\lambda)$ time.

Given such a reduction, we can get an actual SZK protocol for SC^β by first reducing it to $\mathrm{SC}^{2^{-\lambda}}$, and then running the earlier protocol for that. This would have zero-knowledge error $2^{-\lambda}$ and, as a useful byproduct, completeness error $2^{-\lambda}/2$. The soundness error can then be amplified by sequential repetition (strictly speaking, we would need to use the auxiliary input definition of SZK for the repetition to work, but we will ignore this for now).

The reduction follows immediately from what is sometimes called an XOR lemma. Fix any two distributions D_0 and D_1. For any bit b and $n \in \mathbb{N}$, denote by $D_b^{\oplus n}$ the distribution that is sampled as follows:

1. Sample uniformly random bits $b_1, \ldots, b_n \leftarrow \{0,1\}$ such that $b_1 \oplus \cdots \oplus b_n = b$
2. For each $i \in [n]$, sample $x_i \leftarrow D_{b_i}$
3. Output $(x_1, \ldots, x_n)$

Lemma 1.1. $\Delta(D_0^{\oplus n}, D_1^{\oplus n}) = \Delta(D_0, D_1)^n$

The proof of Lemma 1.1 is quite straighforward if you write out the definition of the statistical distance between the $D_b^{\oplus n}$'s in terms of the difference between probabilities and simplify each term. The intuition for the operation, however, is the following. Think of the process of picking a random bit b, and sampling $x \leftarrow D_b$. Now given just x, what is the probability that a distinguisher, even computationally unbounded, can guess b correctly? If D_0 and D_1 had disjoint supports, then b can be guessed perfectly. If D_0 and D_1 are close, however, this becomes harder, and by what we saw in the last lecture, the advantage any distinguisher has in guessing b is at most $\Delta(D_0, D_1)$. In the case of $D_b^{\oplus n}$, in order to guess b, really the best a distinguisher can do is try to guess all the b_i's and take their parity. Thus, its task becomes much harder, so its success probability can be expected to decrease exponentially with n. Even allowing for the fact that it can get some of the b_i's wrong but still guess b correctly maintains this behaviour.

Exercise 1. *Given Lemma 1.1, write down the reduction from* SC^{β} *to* $\mathrm{SC}^{2^{-\lambda}}$*. Prove Lemma 1.1.*

More generally, Lemma 1.1 is a process by which, given two distributions that are somewhat close, we can obtain distributions that are very close. Further, the rate at which the distance between the distributions decreases is a function of their initial distance. In particular, if D_0 and D_1 are disjoint, $D_0^{\oplus n}$ and $D_1^{\oplus n}$ remain disjoint. There is also a simple process that increases the distance between distributions. Denote by $D_b^{\otimes n}$ the distribution obtained by sampled $x_1, \ldots, x_n \leftarrow D_b$, and outputting all of them.

Lemma 1.2. $1 - 2e^{-n\Delta(D_0,D_1)^2/2} \leq \Delta(D_0^{\otimes n}, D_1^{\otimes n}) \leq n \cdot \Delta(D_0, D_1)$

The intuition here is again to look at statistical distance as the best advantage of a distinguisher in recovering b given a sample from D_b. If we want the give the distinguisher more information about b, the simplest thing to do is to just give it many samples from D_b, which is what the above operation does. The lower-bound in the lemma can be proven using Chernoff bounds, and the upper bound using Fact 0.2.

Exercise 2. *Prove Lemma 1.2.*

So now we have two operations, one of which brings distributions closer (at different rates), and another that pushes them apart. Interestingly, a combination of these two operations can be used to "polarise" distributions – that is, bring them closer if they are already somewhat close, and push them apart if they are already somewhat far. This is captured by the following theorem that was proven by Sahai and Vadhan [SV03]

Theorem 1.3 (Polarisation Lemma)**.** *For any constants $\alpha, \beta \in [0,1]$ such that $\alpha^2 > \beta$, there is an algorithm that, given a bit b, security parameter λ, and sample access to distributions D_0 and D_1, runs in time $\mathsf{poly}(\lambda)$ and outputs a sample from distribution D'_b, such that these distributions D'_0 and D'_1 satisfy the following properties:*

- *If $\Delta(D_0, D_1) > \alpha$, then $\Delta(D'_0, D'_1) > 1 - 2^{-\lambda}$*
- *If $\Delta(D_0, D_1) < \beta$, then $\Delta(D'_0, D'_1) < 2^{-\lambda}$*

Whether the condition $\alpha^2 > \beta$ can be removed from the above theorem is still unknown.

2 Generalest Statistical Closeness

We will see now how the honest-verifier PZK protocol for Graph Non-Isomorphism (GNI) can be extended to work for an even broader generalisation of the Statistical Closeness problem. This generalisation will also be significant as a complete problem for the class SZK itself. A natural way to extend the SC^{β}

problem is to allow the distributions in the NO case to be not necessarily disjoint, but just far. For any $\beta < \alpha \in [0,1]$, define the following promise problem:

$$\mathrm{SC}_Y^{\beta,\alpha} = \{(C_0, C_1) \mid \Delta(D_{C_0}, D_{C_1}) \leq \beta\}$$
$$\mathrm{SC}_N^{\beta,\alpha} = \{(C_0, C_1) \mid \Delta(D_{C_0}, D_{C_1}) \geq \alpha\}$$

Note that the protocol for SC^{β} is not sound for $\mathrm{SC}^{\beta,\alpha}$ if $\alpha < 1$, as then the ranges of the cicuits C_0 and C_1 could be identical even if their output distributions are very far. We will see later how to get a protocol for $\mathrm{SC}^{\beta,\alpha}$. For now, we focus on the complement of this problem called the Statistical Difference problem, which we will denote by $\mathrm{SD}^{\alpha,\beta}$. That is, $\mathrm{SD}_Y^{\alpha,\beta} = \mathrm{SC}_N^{\beta,\alpha}$ and $\mathrm{SD}_N^{\alpha,\beta} = \mathrm{SC}_Y^{\beta,\alpha}$.

In an interactive proof for SD, the prover would like to prove that the output distributions of two given circuits are far apart. Just as SC represented a generalisation of the principles we used to get a protocol for GI, SD does the same for its complement GNI. Recall the HVPZK protocol (P, V) for GNI that, given graphs (G_0, G_1), works as follows:

1. V samples a uniformly random relabelling permutation $R : [n] \to [n]$, and a uniformly random bit $b \leftarrow \{0,1\}$. It sends $R(G_b)$ to the prover.
2. If G_0 and G_1 are isomorphic, P sets $b' = \perp$. Else, P sets the bit $b' = 0$ if G' is isomorphic to G_0, and $b' = 1$ otherwise. It sends b' to V.
3. V accepts if $b = b'$, and rejects otherwise.

To get soundness, this protocol uses the fact that $\Delta(R(G_0), R(G_1)) = 0$ if G_0 and G_1 are isomorphic. Completeness and zero-knowledge come from the fact that $\Delta(R(G_0), R(G_1)) = 1$ if the graphs are not isomorphic. A natural extension of the principles here gives us the following protocol (P, V) for $\mathrm{SD}^{\alpha,\beta}$ on input (C_0, C_1) (supposing the circuits both map $\{0,1\}^m$ to $\{0,1\}^n$):

1. V samples a uniformly random $r \leftarrow \{0,1\}^m$ and bit b. It computes $y \leftarrow C_b(r)$ and sends it to the prover.
2. If $\Delta(D_{C_0}, D_{C_1}) \leq \alpha$, P sets $b' = \perp$. Else, it sets $b' = 0$ if $D_{C_0}(y) > D_{C_1}(y)$, and $b' = 1$ otherwise. It sends b' to V.
3. V accepts if $b = b'$, and rejects otherwise.

Completeness. The comparison of the probability masses $D_{C_0}(y)$ and $D_{C_1}(y)$ above comes from the fact that this is the ideal distinguisher that has the greatest advantage in distinguishing between D_{C_0} and D_{C_1} given sample y. That is, let $S \subseteq \{0,1\}^n$ be the set of y's for which $D_{C_0}(y) > D_{C_1}(y)$. Then, we have the following relation:

$$\Delta(D_{C_0}, D_{C_1}) = \sum_{y \in S} D_{C_0}(y) - D_{C_1}(y) = D_{C_0}(S) - D_{C_1}(S)$$

Proving this is left as an exercise. Now, the probability that the prover guesses b correctly is given by:

$$\Pr[b' = b] = \frac{1}{2} \cdot \Pr[b' = 0 \mid b = 0] + \frac{1}{2} \cdot \Pr[b' = 1 \mid b = 1] \tag{1}$$
$$= \frac{1}{2} + \frac{\Pr[b' = 0 \mid b = 0] - \Pr[b' = 0 \mid b = 1]}{2} \tag{2}$$

The prover sets $b' = 0$ iff $y \in S$. So the probability that $b' = 0$ given $b = 0$ is exactly the probability that $y \in S$ when it is sampled from D_{C_0}, which is $D_{C_0}(S)$. Similarly, the probability that $b' = 0$ when $b = 1$ is $D_{C_1}(S)$. Thus, the probability that $b' = b$ is exactly $1/2 + \Delta(D_{C_0}, D_{C_1})/2 \geq (1 + \alpha)/2$.

Soundness. In order to guess b correctly, the prover's task is precisely to distinguish between D_{C_0} and D_{C_1} given just one sample. By the fact that statistical distance bounds the advantage of any adversary in doing this, we have that $(\Pr[b' = 0 \mid b = 0] - \Pr[b' = 0 \mid b = 1])$ is at most $\Delta(D_{C_0}, D_{C_1})$. It then follows from Eq. (1) that the probability that any prover can make the verifier accept is at most $1/2 + \Delta(D_{C_0}, D_{C_1})/2 \leq (1 + \beta)/2$, which is smaller than $(1 + \alpha)/2$.

Honest-Verifier Zero-Knowledge. The simulator S for the honest verifier V takes advantage of the fact that it knows V's secret randomness, which the prover in an actual execution of the protocol does not. On input (C_0, C_1), it works as follows:

1. S samples a uniformly random $r \leftarrow \{0,1\}^m$ and bit b.
2. It computes $y \leftarrow C_b(r)$ and outputs $((r,b), y, b)$.

The only difference between the simulated distribution and the actual view of V, which is $((r,b), y, b')$, is that sometimes in the actual view, b' is not the same as b. In fact, as argued in the proof of completeness, $b' \neq b$ with probability exactly $(1-\alpha)/2$. This is the zero-knowledge error.

Exercise 3. *Prove the above bound on the zero-knowledge error formally using Fact 0.2.*

Reducing the errors. Again, for constants α and β, the above zero-knowledge error is not good enough, and cannot be generically amplified. However, we can take the same approach as we did in the case of SC^{β}, and first use the polarisation lemma (Theorem 1.3) to first reduce $\mathrm{SD}^{\alpha,\beta}$ to $\mathrm{SD}^{1-2^{-\lambda},2^{-\lambda}}$ and then run the above protocol. This is only possible, however, if $\alpha^2 > \beta$.

References

[SV03] Amit Sahai and Salil P. Vadhan. A complete problem for statistical zero knowledge. *J. ACM*, 50(2):196–249, 2003.