Foundations and Frontiers of Probabilistic Proofs (Summer 2021)
Worksheet A.7: Intro to IOPs
Date: 2021.08.04

A PCP is (malicious-verifier) zero knowledge if there exists a polynomial-time simulator S such that, for every instance $\mathbb{x} \in \mathcal{L}$ and polynomial-time malicious verifier $\tilde{V}$, $S(\mathbb{x})$ outputs a view that is distributed identically to $\mathrm{view}_{\tilde{V}}(\tilde{V}^{P(\mathbb{x})})$. In this worksheet we consider a generalization of PCPs called *interactive PCPs* (IPCPs), where the prover supplies a PCP oracle (potentially of superpolynomial size) and then conducts a standard interactive proof (with polynomial-size messages). The view of a verifier in an IPCP consists of its randomness, the answers to its queries to the PCP, the prover messages it receives during the IP.

Problem. (Zero-knowledge sumcheck) We prove that $\#\mathrm{SAT} \in \mathsf{PZKIPCP}$, that is, $\#\mathrm{SAT}$ has an IPCP with perfect zero knowledge against polynomial-time malicious verifiers.

We provide the IPCP simulator with access to an oracle $\mathcal{Q}_{d,n}$ that samples partial sums of a random low-degree multivariate polynomial. $\mathcal{Q}_{d,n}$ takes as input a list $(q_1, \alpha_1, \ldots, q_t, \alpha_t, q^*)$ where $q_i \in \mathbb{F}^{j_i}$ for $j_i \leq n$, $\alpha_i \in \mathbb{F}$ and $q^* \in \mathbb{F}^{j^*}$ for $j^* \leq n$, and outputs a field element $\beta \in \mathbb{F}$ with the following distribution:

$$\Pr[\beta \leftarrow \mathcal{Q}_{d,n}(q_1, \alpha_1, \ldots, q_t, \alpha_t, q^*)] = \Pr_{Q \leftarrow \mathbb{F}^{\leq d}[X_1,\ldots,X_n]}[Q(q^*) = \beta \mid \forall i \in [t],\, Q(q_i) = \alpha_i]$$

where $\mathbb{F}^{\leq d}[X_1, \ldots, X_n]$ is the set of n-variate polynomials of individual degree at most d; $\mathcal{Q}$ outputs $\perp$ if the above conditional probability is undefined. Above $Q(q)$ is defined for $q \in \mathbb{F}^j$ with $j < n$ by "summing out" the remaining indices over $\{0, 1\}$, i.e.

$$Q(q) := \sum_{b_{j+1},\ldots,b_n \in \{0,1\}} Q(q, b_{j+1}, \ldots, b_n) \ .$$

In particular, $Q(\perp) := \sum_{b_1,\ldots,b_n \in \{0,1\}} Q(b_1, \ldots, b_n)$.

(The oracle $\mathcal{Q}$ can in fact be efficiently implemented, but we do not discuss this technique here.)

The prover and verifier receive as input a boolean k-CNF formula φ with n variables and m clauses, and a claimed number of satisfying assigments a. They agree on a field $\mathbb{F}$ whose size is a prime larger than 2^n, and also each compute the arithmetization $\widehat{\varphi}$ of φ, which has individual degree $d = \mathsf{poly}(n, m)$. They then interact as follows:

1. The prover samples $R \in \mathbb{F}^{\leq d}[X_1, \ldots, X_n]$ uniformly at random and sends it to the verifier, along with the value $z := \sum_{b_1,\ldots,b_n \in \{0,1\}} R(b_1, \ldots, b_n)$.
2. The verifier sends uniformly random $\rho \in \mathbb{F}$ to the prover.
3. The prover and the verifier engage in the standard sumcheck protocol with respect to the polynomial $\rho\widehat{\varphi} + R$ and claimed sum $\rho a + z$. (Here the verifier makes a single query to R.)
4. The verifier checks that R is δ-close to low-degree (e.g. using a line-vs-point test).

1. Show that this protocol is complete and sound.
2. Show that this protocol achieves perfect zero knowledge.

(Hint 1: the simulator can be "straightline", i.e., does not rewind the malicious verifier.)
(Hint 2: first consider the case where the malicious verifier does not query R before sending ρ.)