6.841 Advanced Complexity Theory April, 23, 2007

Lecture 19

Lecturer: Madhu Sudan *Scribe: Kuen-Bang Hou and Ning Xie*

Today we will continue our study of PCP's and show an exponentially long PCP for SAT.

1 Review of last lecture: two views of PCP

First we review some results covered in last lecture: there are two different views to look at PCP systems. The first view is to treat PCP as a proof system having the following special property. There exists a probabilistic polynomial time proof verifier V for L such that:

- $x \in L \Rightarrow \exists \pi$ s.t. $\Pr_R[V^\pi(x,R)] = 1$
- $x \notin L \Rightarrow \forall \pi \ \Pr_R[V^\pi(x,R)] \leq 1-\epsilon$,

where $|\pi|$ should be small, $\epsilon > 0$ is a constant and #queries into π should be small (e.g. $O(1)$).

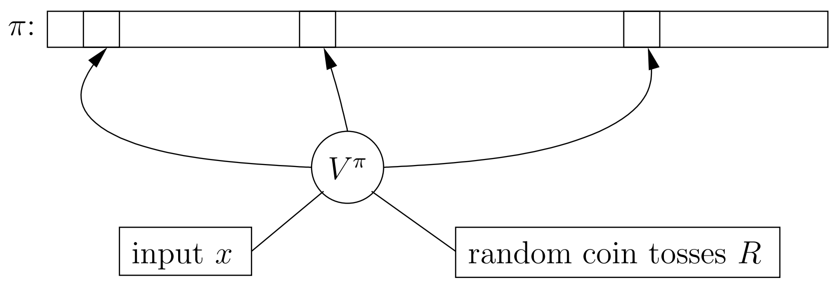

Figure 1: Relationship among the proof π, the input x, the random coin tosses R and the verifier V

The second view is to think PCP as a reduction from L to the problem of Generalized Graph k-coloring such that:

- $x \in L \Longrightarrow G_x$ is k-colorable;
- $x \notin L \Longrightarrow \forall k$-coloring of G_x, at least ϵ fraction of the edges are invalid coloring in G_x.

Note that if $|\pi|$ is polynomially bounded, these two views are equivalent. However, today we are going to see an exponentially long proof verifiable by $O(1)$ queries. This is trivial in View 2 but is highly non-trivial in View 1.

2 PCP system for Quadratic-SAT

2.1 Quadratic-SAT

Definition 1 (Quadratic-SAT) *Consider the following decision problem, which is a variant of* SAT*:*

Given : $x_1, \ldots, x_n \in \mathrm{GF}(2)$ *and a set of m degree-2 polynomials $P_1, \cdots, P_m$ in n variable;*
Question : *Does there exist an $\mathbf{a} = (a_1, \ldots, a_n) \in \mathrm{GF}(2)^n$ such that for all $j \in \{1, \ldots, m\}$ $P_j(\mathbf{a}) = 0$?*

It is easy to see that Quadratic-SAT $\in$ NP. It can also be checked that Quadratic-SAT is NP-hard[1]. Therefore it is NP-complete. In the following we will describe an exponentially long PCP for Quadratic-SAT. The key ideas will be arithmetization of SAT and exploiting some nice properties of linear functions (low-degree polynomials).

2.2 What is the proof

Let Q be a degree-2 polynomials in n variables over GF(2). Then $Q(x_1, \ldots, x_n) = \sum_{1 \leq i,j \leq n} q_{i,j} x_i x_j + q_0$. Since we are working over GF(2), $x^2 = x$ and this general form includes all the linear functions as well. Note that a quadratic polynomial over GF(2) is completely determined by the set of coefficients: $Q \equiv (\{q_{i,j}\}, q_0)$. It follows that the total number of degree-2 polynomials in n variables over GF(2) is $2^{O(n^2)}$. Now our proof π for the PCP system is simply the list of the evaluations of a satisfying assignment $\mathbf{a}$ at all quadratic polynomials. Therefore $|\pi| = 2^{O(n^2)}$.

2.3 What should be checked for the proof

There are two issues to address. First, π may not equal to $\{Q(\mathbf{a})\}$ for any $\mathbf{a}$. Second, $\mathbf{a}$ may not be a satisfying assignment.

- **Syntactic Question:** Does there exist an $\mathbf{a}$ such that $\pi[Q] = Q(\mathbf{a})$ for all Q?

  Since the number of quadratic polynomials is exponentially large and an invalid proof may be formed by flip only one bit from a valid proof, this is not possible to check in polynomial time. Instead, we relax the question to: Does there exist an $\mathbf{a}$ such that $\Pr_Q[\pi[Q] = Q(\mathbf{a})] \geq 1 - \delta$? Note that even after the relaxation there can be only one $\mathbf{a}$ that passes the check provided that δ is small enough.

- **Semantic Question:** Is $P_1(\mathbf{a}) = P_2(\mathbf{a}) = \cdots = P_m(\mathbf{a}) = 0$?

We will study the semantic question first since it is easier and then come back to handle the syntactic question later.

2.4 Semantic test

Now we assume that the proof π already passes the Syntactic test; i.e., we are given a table π which encodes the evaluation of some $\mathbf{a}$ at all the quadratic polynomials such that for at least $1 - \delta$ fraction of the points, $\pi[Q] = Q(\mathbf{a})$. We want to test if $\mathbf{a}$ is a satisfying assignment for Quadratic-SAT by probing the table only at a constant number of locations.

We start with the easiest case: Suppose that there is only one polynomial P_1 (i.e. $m = 1$). Note that we can not just read $\pi[P_1]$ since that point may be in the corrupted portion of the proof. Instead, we use the idea of random self-reducibility introduced before: Pick another polynomial Q at random, compute $\tilde{\pi}[P_1] \stackrel{\text{def}}{=} \pi[P_1 + Q] - \pi[Q]$ and check if it is 0. The key point here is that, for any fixed P_1, if Q is a random quadratic polynomial, then so is $P_1 + Q$. Therefore, $\Pr_Q[\pi[Q] \neq Q(\mathbf{a})] \leq \delta$ and $\Pr_Q[\pi[P_1 + Q] \neq P_1(\mathbf{a}) + Q(\mathbf{a})] \leq \delta$, applying union bound gives $\Pr_Q[\tilde{\pi}[P_1] \neq P_1(\mathbf{a})] \leq 2\delta$.

For general m, we can not repeat the above test for every P_i, since we are only allowed to query constant bits. We will use the idea of approximating OR gates by probabilistic low-degree polynomials in Razborov-Smolensky's proof of circuit lower bound for PARITY. Here our task is to check if $\bigwedge_{j=1}^{m}(P_j(\mathbf{a}) = 0)$.

1. pick $\alpha_1, \ldots, \alpha_m \in$ GF(2) uniformly at random;

2. check if $P_\alpha(x_1, \ldots, x_n) \stackrel{\text{def}}{=} \sum_{j=1}^{m} \alpha_j P_j(x_1, \ldots, x_n)$ evaluates to 0 at point $\mathbf{a}$.

[1] Note that if we map 1 to TRUE and map 0 to FALSE, then the AND gate and OR gate can be expressed by quadratic polynomials over GF(2) as $\text{AND}(x, y) = x \cdot y$ and $\text{OR}(x, y) = x + y + x \cdot y$. Consider the following reduction from 3SAT to Quadratic-SAT. Let $\psi \in$ 3SAT. For each clause $c_i = (x_i \vee y_i \vee z_i)$ in ψ (note that the complement of x is mapped to $(1 - x)$ and this will not increase the degree), build two polynomials P_{2i}, P_{2i+1} as $P_{2i} = x_i + y_i + x_i y_i + w_i$ and $P_{2i+1} = z_i + w_i + z_i w_i + 1$. This construction introduces at most a polynomial number of new variables and it is easily seen that $P_{2i} = P_{2i+1} = 0$ if and only if $c_i = (x_i \vee y_i \vee z_i) = $ TRUE.

Analysis: Note that P_α is a degree-2 polynomial in $x_1, \ldots, x_n$. It is easy to see that, if for all $j \in [m]$ $P_j(\mathbf{a}) = 0$, then $P_\alpha(\mathbf{a}) = 0$. On the other hand, if there exists a j such that $P_j(\mathbf{a}) \neq 0$, then P_α is a non-vanishing multilinear polynomial in $\alpha_1, \ldots, \alpha_m$. By Schwartz-Zippel Lemma, $\Pr_{\alpha_1,\ldots,\alpha_m}[P_\alpha(\mathbf{a}) \neq 0] \geq \frac{1}{2}$.

Now we combine these two ideas together and get the following Semantic Test:

Semantic Test: Is $P_1(\mathbf{a}) = P_2(\mathbf{a}) = \cdots = P_m(\mathbf{a}) = 0$?

- pick $\alpha_1, \ldots, \alpha_m \in \mathrm{GF}(2)$ uniformly at random
- set $P_\alpha = \sum_{j=1}^m \alpha_j P_j$
- pick Q randomly from the set of quadratic polynomials
- accept if $\pi[Q + P_\alpha] = \pi[Q]$

2.5 Syntactic test

Now we come back to the first test. Recall that our task is, given a proof π, to test if there exists an $\mathbf{a}$ such that $\Pr_Q[\pi[Q] = Q(\mathbf{a})] \leq \delta$. The idea is to look for structural properties of such a proof π and test for them. Observe that quadratic function evaluation satisfies the linearity property: $\pi[Q_1] + \pi[Q_2] = \pi[Q_1 + Q_2]$ for all Q_1 and Q_2. Conversely, for any $\tilde{\pi}$ satisfying that

$$\forall Q_1, Q_2, \tilde{\pi}[Q_1] + \tilde{\pi}[Q_2] = \tilde{\pi}[Q_1 + Q_2],$$

there exist $\{b_{i,j}\}_{i=1,j=1}^{n,n}$ and b_0 such that for all $Q = (\{q_{i,j}\}, q_0)$, $\tilde{\pi}[Q] = \sum q_{i,j} b_{i,j} + q_0 b_0$.

Therefore we break the Syntactic Test into two parts: First we test if the proof π passes the linearity test, then we check if

- $b_{i,j} = a_i a_j$ for all i and j, and
- $b_0 = 1$.

2.5.1 The First Part of the Syntactic Test

We've already used the idea $Q_1(a) + Q_2(a) = (Q_1 + Q_2)(a)$ in the previous test. Here we are going to use that idea again. Our linearity test is simply the following: Pick Q_1, Q_2 at random and check if $\pi[Q_1] + \pi[Q_2] = \pi[Q_1 + Q_2]$.

If the proof π passes the test, however, we can only conclude that $\Pr_{Q_1,Q_2}[\pi[Q_1] + \pi[Q_2] = \pi[Q_1 + Q_2]]$ is high, which is far from the statement that *for all* Q_1, Q_2, $\pi[Q_1] + \pi[Q_2] = \pi[Q_1 + Q_2]$. Fortunately, this property guarantees that there is another proof $\tilde{\pi}$ which is linear (i.e. for all Q_1, Q_2, $\tilde{\pi}[Q_1] + \tilde{\pi}[Q_2] = \tilde{\pi}[Q_1 + Q_2]$) and is very close to π. The existence of such $\tilde{\pi}$ is stated formally in the following remarkable theorem of Blum, Luby and Rubinfeld:

Theorem 2 (BLR Theorem) *If* $\Pr[\pi[Q_1] + \pi[Q_2] \neq \pi[Q_1 + Q_2]] \leq \delta$ *then*

1. $\exists \tilde{\pi}$ *such that* $\forall Q_1, Q_2 : \tilde{\pi}[Q_1] + \tilde{\pi}[Q_2] = \tilde{\pi}[Q_1 + Q_2]$ *and*
2. $\Pr_Q[\pi[Q] \neq \tilde{\pi}[Q]] \leq 2\delta$

provided that $\delta < 2/9$.

We will not prove this theorem here. Interested readers are referred to the two courses taught by Prof. Rubinfeld: 6.895 Randomness and Computation and 6.896 Sublinear Time Algorithms.

The proof is in [BLR90]. The constant 2/9 is important because we're only allowed to use a constant number of queries, and the soundness bound 2/9 is achievable by a constant number of queries.

2.5.2 The Second Part of the Syntactic Test

It remains to test the following: Are b's generated from some assignment $\mathbf{a}$ such $b_0 = 1$ and $b_{i,j} = a_i a_j$ if we write the value of $\tilde{\pi}[Q]$ as $\sum q_{i,j} b_{i,j} + q_0 b_0$? Keep in mind that we are only given a proof π which may disagree with $\tilde{\pi}$ on 2δ fraction of the points.

1. Is $b_0 = 1$? Define the polynomial $\bar{1}$ to be $\{\{0\}, 1\}$, then $\bar{1}(a_1, \ldots, a_n) = b_0$. Using the same idea we used before, we pick a random Q and test if $\pi[Q + \bar{1}] - \pi[Q] = 1$.

2. Is $b_{i,j} = a_i a_j$ for every i, j? This is equivalent to testing if $\mathbf{b} = \{b\}_{i,j} = \mathbf{a}\mathbf{a}^T$. First we state a technical claim.

   Claim 3 *Let* $\mathbf{b} \in \mathrm{GF}(2)^{n \times n}$ *and let* $\mathbf{a} \in \mathrm{GF}(2)^n$. *Let* $u, v \in \mathrm{GF}(2)^n$. *If* $\mathbf{b} \neq \mathbf{a}\mathbf{a}^T$, *then*

   $$\Pr_{u,v}[u^T \mathbf{b} v \neq u^T \mathbf{a}\mathbf{a}^T v] \geq 1/4.$$

   Proof Left as an exercise to the reader. ■

   This claim suggests the following test: Pick u and v at random and check if $u^T \mathbf{b} v = u^T \mathbf{a}\mathbf{a}^T v$. The left hand side can be read from $\pi[Q_{u,v}]$, where $Q_{u,v} = (q_{i,j}, q_0)$, $q_{i,j} = u_i v_j$ and $q_0 = 0$. But how do we compute $v^T a$ and $u^T a$? Instead, we ask the prover to provide the answers and we check them! Specifically, the prover provides an appendix π_{lin} where $\pi_{\text{lin}}[v] = v^T \mathbf{a}$ for each vector v. Note that this is just the evaluation of all linear functions at point $\mathbf{a}$. Now assuming the appendix is correct, the following test will work:

   (a) pick random vectors u, v and random quadratic polynomial Q

   (b) test if $\pi[Q + Q_{u,v}] - \pi[Q] = \pi_{\text{lin}}[u] \cdot \pi_{\text{lin}}[v]$

However, we have to make sure that the appendix is correct. As before, this can be done by picking u and v at random and testing if

$$\pi_{\text{lin}}[u] + \pi_{\text{lin}}[v] = \pi_{\text{lin}}[u + v].$$

2.6 Summary

Now we have completed the description a PCP system for Quadratic-SAT. The prover provides π and π_{lin}, which are evaluations of the set of quadratic polynomials and the set of linear functions at a satisfying assignment $\mathbf{a}$. The verifier performs the following tests and accepts only if all the tests pass.

1. Test 1: Linearity of π_{lin}
   - $\pi_{\text{lin}}[u] + \pi_{\text{lin}}[v] = \pi_{\text{lin}}[u + v]$?

2. Test 2: Quadraticity of π

   (a) $\pi[Q_1] + \pi[Q_2] = \pi[Q_1 + Q_2]$?

   (b) $\pi[Q + Q_{u,v}] - \pi[Q] = \pi_{\text{lin}}[u] \cdot \pi_{\text{lin}}[v]$?

   (c) $\pi[Q + \bar{1}] - \pi[Q] = 1$?

3. Semantic Test
   - pick $\{\alpha_1, \ldots, \alpha_m\}$ at random and set $P_\alpha = \sum_{j=1}^m \alpha_j P_j$
   - $Q \leftarrow$ random
   - $\pi[Q + P_\alpha] = \pi[Q]$?

Analysis:

- The verifier makes $14 = O(1)$ number of queries to the proof
- $(\exists \mathbf{a}$ s.t. $P_j(\mathbf{a}) = 0 \ \forall\ j) \Rightarrow \exists \pi, \pi_{\text{lin}}$ s.t. $\Pr[\text{Verifier accepts}] = 1$
- $\Pr[\text{Verifier accepts}] \geq 0.99 \Rightarrow \exists \mathbf{a}$ s.t. $P_1(\mathbf{a}) = \cdots = P_m(\mathbf{a}) = 0$

2.7 Exercise for the next lecture

We would like to have a reduction which has the following property: $\exists \tau > 0$ such that $\forall k$ Generalized k-coloring $\leq$ Generalized 3-coloring such that

- G_x is k-colorable $\rightarrow G'_x$ is 3-colorable
- G_x is ϵ-far from k-colorable $\rightarrow G'_x$ is $\epsilon \cdot \tau$-far from 3-colorable.

Note that this is different from the classical Garey-Johnson type reduction, in which τ is in general dependent on k.

References

[BLR90] M. Blum, M. Luby, and R. Rubinfeld. Self-testing/correcting with applications to numerical problems. In *STOC '90: Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 73–83, New York, NY, USA, 1990. ACM Press.