Lecture 9: MIP=NEXP

Topics in Pseudorandomness and Complexity Theory (Spring 2017)
Rutgers University
Swastik Kopparty
Scribe: Jinyoung Park

1 Introduction

Definition 1. ***NP*** *is a class of language L such that there exists a poly time verifier V with*

- *For any $x \in L$, $|x| = n$, there exists a proof $\Pi \in \{0,1\}^{poly(n)}$ such that $V(x, \Pi) = ACCEPT$*
- *For $x \notin L$, for any proof $\Pi \in \{0,1\}^{poly(n)}$, $V(x, \Pi) = REJECT$*

Definition 2. ***IP*** *is a class of language L such that there is a randomized poly time verifier V such that*

- *For any $x \in L$, there exists a prover P such that $\Pr[V(x) = ACCEPT] = 1$*
- *For $x \notin L$, for any prover P, $\Pr[V(x) = REJECT] \geq 0.9$*

Definition 3. ***MIP*** *is a class of language L such that there is a randomized poly time verifier V such that*

- *For any $x \in L$, there exists a proof $\Pi \in \{0,1\}^{exp(n)}$ such that $\Pr[V^{\Pi}(x) = ACCEPT] = 1$*
- *For $x \notin L$, for all proof $\Pi \in \{0,1\}^{exp(n)}$, $\Pr[V^{\Pi}(x) = REJECT] \geq 0.9$*

The question is, which one is more powerful? In other words, which class of language is bigger? Today's goal is to prove the following theorem, based on the paper by Babai, Fortnow, and Lund in 90's.

Theorem 4. ***MIP=NEXP****(Non-deterministic Exponential Time)*

Note that one direction is obvious: $MIP \subseteq NEXP$.

(Sketch of proof) Suppose there is a language which has such a verifier in MIP. Then there is a non-deterministic exponential time algorithm to check whether x is in the language or not. Indeed, what this machine does is that it first guesses what the proof Π is, then confirms that for all outcomes of the randomness it always accepts. If this happens, then it means that x is in the language. □

The other direction is difficult, and we will do it today.

2 Plan

Three ingredients of today's proof are as the following.

1. Arithmetization
2. Low degree testing
3. Combinatorial Nullstellensatz

Before we start the proof, let's describe what $NEXP$ problems are.

Definition 5. ***3 SAT*** *is a boolean formula ϕ in 3 CNF which is satisfiable.*

An example of 3 CNF : $(x_1 \vee x_2 \vee \bar{x_3}) \wedge (\bar{x_1} \vee x_4 \vee x_7) \wedge (...)$

We know that 3 SAT is an NP complete problem. Similarly, we describe NEXP.

Definition 6. ***SUCCINT-3CSP*** *is a boolean formula $F : \{0,1\}^m \times \{0,1\}^{3l} \times \{0,1\}^3 \to \{0,1\}$ which represents a function H on 2^l variables $x_1, \cdots, x_{2^l}$,*

$$H(x_1, \cdots, x_{2^l}) = \wedge_{z \in \{0,1\}^m, i,j,k \in \{0,1\}^l} F(z, i, j, k, x_i, x_j, x_k)$$

such that H is satisfiable.

(i.e., $\exists x_1, \cdots, x_{2^l} \in \{0,1\}$ such that $H(x_1, \cdots, x_{2^l}) = 1$)

Fact 7. *SUCCINT-3CSP is NEXP complete.*

What is the standard NEXP algorithm for deciding if $F \in$ SUCCINT-3CSP?

1. Guess the setting $x_1, \cdots, x_{2^l}$($\to$ The NEXP proof/witness)
2. In exponential time, verify that $H(x_1, \cdots, x_{2^l}) = 1$

So our goal is to write some auxiliary information along with $x_1, \cdots, x_{2^l}$ so that the verification can be done in randomized poly time. This auxiliary information will be very closely related to the method that we used for hardness amplification.

3 Proof

Consider what should be written down in the "YES" case. If x is in the language, then there is an assignment $x_1, \cdots, x_{2^l}$. Let's view those variables $x_1, \cdots, x_{2^l}$ as a function $A : \{0,1\}^l \to \{0,1\}$ such that $A(i) = x_i$.

The key fact is that the function F, which is a boolean formula, can be written as a low degree polynomial. This corresponds to the first ingredient in our plan, "arithmetization".

3.1 Arithmetization

Definition 8. ***Arithmetization of F*** *is an arithmetic formula Q with $+, \times$, such that $Q(y) = F(y)$ whenever y is a 0,1 string.*

For this: remove $\vee$ by replacing with $\neg$ and $\wedge$, and then switch $b \wedge b'$ to $b \times b'$ and $\neg b$ to $1 - b$.

Note that for the resulting polynomial Q, $\deg(Q) \leq \text{size}(F)$.

Having done this, we want to check that for all $z \in \{0,1\}^m, i, j, k \in \{0,1\}^l$,

$$Q(z, i, j, k, A(i), A(j), A(k)) = 1$$

3.2 Low degree testing

Pick a prime p of size $> \text{poly}(m, l, \text{size}(F))$. We work with the field $\mathbb{F}_p$, and we view $\{0,1\} \subset \mathbb{F}_p$.

Given $A : \{0,1\}^l \to \{0,1\}$, let $\tilde{A}(R_1, \cdots, R_l)$ be the unique polynomial of degree at most 1 in each variable such that $\tilde{A}|_{\{0,1\}^l} = A$. Note that this is the low degree extension that we did before.

Write down $\tilde{A}(r)$ for all $r \in \mathbb{F}_p^l$

Now we need to check:

1. $\tilde{A}$ is $\{0,1\}$-valued on $\{0,1\}^l$
2. $\forall z \in \{0,1\}^m$, $\forall i, j, k \in \{0,1\}^l$, $Q(z, i, j, k, \tilde{A}(i), \tilde{A}(j), \tilde{A}(k)) = 1$

Note that Q is a low degree polynomial, so we can evaluate it in any way.

In the case $F \in L$, here are some observations:

1. $\tilde{A}$ is of degree $\leq l \ll |\mathbb{F}_p|$
2. $S(Z, I, J, K) \triangleq Q(Z, I, J, K, \tilde{A}(I), \tilde{A}(J), \tilde{A}(K))$ is of degree $\leq \text{size}(F) \cdot l \ll |\mathbb{F}_p|$

That is, $\tilde{A}$ and S are both "low degree," and once we show the low-degreeness of $\tilde{A}$, then by the definition of S, it will automatically imply the low-degreeness of S.

3.3 Combinatorial Nullstellensatz

First, we want to check 1. $\tilde{A}$ is $\{0,1\}$-valued on $\{0,1\}^l$. Note that this is equivalent to $\tilde{A}(I)(\tilde{A}(I) - 1) = 0$ on $\{0,1\}^l$, and we need the following algebraic tool.

Theorem 9. *(Combinatorial Nullstellensatz)*

Suppose deg $P(Y_1, \cdots, Y_l) = d$. Then $P(Y_1, \cdots, Y_l) = 0$ on $\{0,1\}^l$ iff there exist polynomials $P_1, \cdots, P_l$ such that

1. $deg(P_i) \leq d$

2. $P(Y_1, \cdots, Y_l) = \sum_{i=1}^{l}(Y_i^2 - Y_i)P_i(Y_1, \cdots, Y_l)$

Proof. Take $P(Y_1, \cdots, Y_l)$, and divide it by $Y_1^2 - Y_1, Y_2^2 - Y_2, \cdots, Y_l^2 - Y_l$ (that is, any Y_i^2 can be replaced by Y_i).

This gives a polynomial $P'(Y_1, \cdots, Y_l)$ such that

$$P(Y_1, \cdots, Y_l) = P'(Y_1, \cdots, Y_l) + \sum P_i(Y_1, \cdots, Y_l)(Y_i^2 - Y_i)$$

and P' has individual degree ≤ 1

Now, we want to show that P' is zero.

Suppose $P' \neq 0$. Then for some $y \in \{0,1\}^l$, $P'(y) \neq 0$. Let $P'(Y) = \sum_{S \in \mathcal{F}} a_S \Pi_{i \in S} Y_i$ with $a_S \neq 0$.

Take a minimal monomial $S \in \mathcal{F}$.

Set $y_i = 1$ for $i \in S$, and $y_i = 0$ outside. Then $P'(y) \neq 0$, so $P \neq 0$.

Thus, $P = 0$ on $\{0,1\}^l \Leftrightarrow P' = 0$. □

Now, by the above theorem we know that:

$\tilde{A}(\tilde{A} - 1)$ vanishes on $\{0,1\}^l$ if and only if there exist $P_1, \cdots, P_l$ such that $\tilde{A}(\tilde{A} - 1) = \sum_{j=1}^{l}(I_j^2 - I_j)P_j(I_1, \cdots, I_l)$

So far what we have discussed is what to do in the case of good proof. So we started from the case that the formula is in the language, and then discussed the low-degreeness of $\tilde{A}$, and etc. Later we need to show that if the formula is not satisfiable then this test cannot pass with good probability.

PCP part 2 : Write down evaluations of $P_j(i)$ for all $i \in \mathbb{F}_p^l$.

Verification procedure: part 1+2

1. Check that evaluation table of $\tilde{A}$ is low degree
2. Check that evaluation table of P_j is low degree
3. Check that at a random point $r \in \mathbb{F}_p^l$, $\tilde{A}(r)(\tilde{A}(r) - 1) = \sum_j P_j(r)(r_j^2 - r_j)$

In the above procedure, if 1,2 are true, then this is a good test; in item 3, for a random point if you see a difference, then you can reject, because two different low degree polynomial cannot agree on too many points. If this test passes, then that means that $\tilde{A}(r)(\tilde{A}(r) - 1)$ can be written down of the form $\sum_j P_j(r)(r_j^2 - r_j)$ and this implies that $\tilde{A}(r)(\tilde{A}(r) - 1)$ vanishes on $\{0,1\}^l$.

So far we have not seen how to check the low-degreeness.

3.4 Low degree testing theorem

Theorem 10. *(Low degree testing theorem) Let $\mathbb{F}_p$ be a field, and $d < p/2, m$*

There is a testing algorithm T such that given access to any $f : \mathbb{F}_p^m \to \mathbb{F}_p$ it makes $\leq$ poly(d) queries,

1. *If f is deg $\leq d$, then T ACCEPTS always.*
2. *If f is ϵ-far from deg $\leq d$, then T REJECTS with probability 0.9*

Note that the above generalizes the linearity testing that we did before, and we will prove it later.

3.5 Verification procedure

The verification procedure is as the following.

Proof format: This is the format that the prover gives to the verifier. We expect $\tilde{a}$ to be in $\tilde{A}$, but it may not be.

Let $\tilde{a} : \mathbb{F}_p^l \to \mathbb{F}_p$, $b_1, b_2, \cdots, b_l : \mathbb{F}_p^l \to \mathbb{F}_p$, $S_1, \cdots, S_{m+3l} : \mathbb{F}_p^{m+3l} \to \mathbb{F}_p$.

Verification Procedure

1. Test that $\tilde{a}, b_1, \cdots, b_l, S_1, \cdots, S_{m+3l}$ are all of degree $< p/1000$.

   If any low degree test rejects, then the verifier rejects.
2. Pick random $r \in \mathbb{F}_p^l$ and check that $\tilde{a}(r) \cdot (\tilde{a}(r) - 1) = \sum_j b_j(r) \cdot (r_j^2 - r_j)$
3. Construct arithmetization Q of F
4. Pick random $z \in \mathbb{F}_p^m$, $i, j, k \in \mathbb{F}_p^l$ and check that

   $Q(z, i, j, k, \tilde{a}(i), \tilde{a}(j), \tilde{a}(k)) - 1 =$
   $\sum_{u=1}^m (z_u^2 - z_u) \cdot S_u(z, i, j, k) + \sum_{v=1}^l (i_v^2 - i_v) S_{m+v}(z, i, j, k) + \sum_{v=1}^l (j_v^2 - j_v) S_{m+v}(z, i, j, k) + \sum_{v=1}^l (k_v^2 - k_v) S_{m+v}(z, i, j, k)$

   If any of check fails, verifier rejects, else verifier accepts.

In the case that there is a satisfying assignment, there is a proof that makes this test passes w.p. 1. So it is easy to see that:

If $F \in$ SUCCINT 3-CSP, then there exists a proof that makes v accept with probability 1.

Suppose v accepts with probability > 0.99. (The reason that we can assume this is: once we make sure that v rejects with probability 0.01, then we can amplify the rejection probability by repeating the test independently with the same proof, for the independent ramdomness.)

Since step 1 passes with good probability, there exist polynomials $\bar{\tilde{a}}, \bar{b}_1, \cdots, \bar{b}_l, \bar{S}_1, \cdots, \bar{S}_{m+3l}$ of low degree, ϵ-close to $\tilde{a}, b_1, \cdots, b_l, S_1, \cdots, S_{m+3l}$.

Since step 2 passes with probability ≥ 0.99, $\bar{\tilde{a}}$ is close to $\tilde{a}$ and $\bar{b}_j$ is close to b_j. So by union bound,

$\bar{\tilde{a}}(r) \cdot (\bar{\tilde{a}}(r) - 1) = \sum (r_j^2 - r_j) \cdot \bar{b}_j(r)$ with probability $\geq 0.99 - \epsilon(l+1) \cdots (*)$

Now, $\bar{\tilde{a}}$ and $\bar{b}_j$ are low degree polynomials so $(*)$ passes with probability 1.(both sides are of degree $=o(p)$)

This implies $\bar{\tilde{a}}(R)(\bar{\tilde{a}}(R) - 1)$ vanishes on $\{0,1\}^l$.

Claim 11. *$\bar{\tilde{a}}|_{\{0,1\}^l}$ is the satisfying assignment to the original H.*

Proof. Since step 4 passes w.p. ≥ 0.99,

$Q(z, i, j, k, \bar{\tilde{a}}(i), \bar{\tilde{a}}(j), \bar{\tilde{a}}(k)) - 1 =$

$\sum_{u=1}^{m} (z_u^2 - z_u) \cdot \bar{S}_u(z, i, j, k) + \sum_{v=1}^{l} (i_v^2 - i_v) \bar{S}_{m+v}(z, i, j, k) + \sum_{v=1}^{l} (j_v^2 - j_v) \bar{S}_{m+v}(z, i, j, k) + \sum_{v=1}^{l} (k_v^2 - k_v) \bar{S}_{m+v}(z, i, j, k) \cdots (**)$

with probability $\geq 0.99 - (m + 3l + 3)\epsilon$.

Both sides are degree $o(p)$, and so $(**)$ holds w.p. $= 1$. So this equality is actually equality of polynomials. Therefore,

$Q(Z, I, J, K, \bar{\tilde{a}}(I), \bar{\tilde{a}}(J), \bar{\tilde{a}}(K)) - 1$ vanishes on $\{0,1\}^{m+3l}$.

On boolean inputs, it means that

$F(z, i, j, k, \bar{\tilde{a}}(i), \bar{\tilde{a}}(j), \bar{\tilde{a}}(k)) = 1$, $\forall z, i, j, k$ □

4 Related facts to the low degree testing theorem

Fact 12. *Let $d < p/2$. Suppose $f : \mathbb{F}_p^m \to \mathbb{F}_p$ is such that $\forall$ lines l, $f|_l$ has degree $\leq d$.*

Then deg $(f) \leq d$.

Proof. (Incomplete) Take f, write it as a polynomial of degree $\leq p - 1$ in each variable. (We can always do this.)

$f(X_1, \cdots, X_m) = \sum_{i_1, \cdots, i_m} a_{i_1 \cdots i_m} X_1^{i_1} \cdots X_m^{i_m}$

Suppose there exists some monomial with degree $> d$. We want to find a line l such that $f|_l$ has degree $> d$.

If there is just one monomial of degree $> d$, then there is a line on which this monomial preserves its degree, so we are done. The issue is the case that f has more than one monomial of degree $> d$. In such a case the monomials might cancel each other on a line, and actually the statement is not true if $d \geq p/2$ or p is not prime.

The following proof strategy is similar to that of BLR linearity test.

Suppose $\deg(f) = e > d$ but $\deg(f) < p$. Then look at $f(A + TB) = f_0(A, B) + f_1(A, B)T + \cdots + f_e(A, B)T^e$, where each f_i has degree $\leq p - 1$ and $e > d$.

Since f_e is non-zero polynomial, and it's of degree $\leq p-1$, there is a substitution of values A and B that make f_e nonzero. This makes $f(A + TB)$ a polynomial of degree e, and this is a presentation of the function restricted on a line. □

We will do this proof again next time.