Semidefinite programming and computational aspects of entanglement **IHP Fall 2017**

Lecture 3: November 10

Lecturer: Anand Natarajan *Scribe: ?*

Yesterday, you saw computational hardness results for h_{Sep}. Today, we'll do it for ω^*.

3.1 Definitions

A game G is specified by the question and answer alphabets X, Y, A, B, the distribution π over questions, and the predicate function V. Given a game G, want to approximate ω^*. To compute a (c, s)-approximation to ω^* means to decide if

- $\omega^* \geq c$, or
- $\omega^* \leq s$,

promised that one of the two cases holds. Think of $c - s$ as the precision to which we're computing ω^*.

3.2 CSPs and Games

A CSP over alphabet Σ with clauses $C_1, \ldots, C_m : \Sigma^k \to \{0, 1\}$ is given by the function

$$\phi(x) = \frac{1}{m} \sum_{i=1}^{m} C_i(x_{j_{i,1}}, x_{j_{i,2}}, \ldots, x_{j_{i,k}}).$$

The value of a CSP is $\max_x \phi(x)$. Let's also recall the PCP theorem, which (roughly speaking) says that there exists a constant s such that $1, s$-approximations of our CSP are NP-hard.

CSPs are intimately connected to games. For every CSP ϕ, can construct a 2-player game called the *clause-variable game* $G_{cv}(\phi)$:

- Referee samples an index $i \in [m]$ of a clause at random, and an index $\ell \in [k]$.
- Alice receives i, and responds with an assignment $a_{i,1}, \ldots, a_{i,k}$.
- Bob receives $j_{i,\ell}$, and responds with b.
- Win iff $b = a_{i,\ell}$, and $C_i(a_{i,1}, \ldots, a_{i,\ell}) = 1$.

It's easy to see that if $\omega(G_{cv}(\phi)) = 1$ iff $\text{val}(\phi) = 1$: Alice and Bob both have to play according to the satisfying assignment. We will now show that a converse result. Suppose that $\omega(G_{cv}(\phi)) \geq 1 - \epsilon$. Then we will show a lower-bound on $\text{val}(\phi)$. Indeed, let b be the fixed assignment to all the variables used by Bob in the optimal strategy; then we claim that $\phi(b) \geq 1 - (k^2 + 1)\epsilon$. To see this, we compute the probability that b satisfies a random clause C_i. Let a_i be the assignment used by Alice to the variables in this clause.

$$\begin{aligned}
\Pr[C_i(b) = 1] &\geq \Pr\big[C_i(a_i) = 1 \wedge \forall \ell \in [k], a_{i,\ell} = b_{j_{i,\ell}}\big] \\
&\geq 1 - \Pr[C_i(a_i) \neq 1] - \sum_{\ell=1}^{k} \Pr\big[a_{i,\ell} \neq b_{j_{i,\ell}}\big] \\
&\geq 1 - \epsilon - k^2\epsilon.
\end{aligned}$$

This means that a $(1, 1 - \epsilon)$ approximation to ω becomes a $(1, 1 - (k^2 + 1)\epsilon))$ approximation to ϕ. Hence, if our starting CSP was PCP-able, we have shown NP hardness for constant-factor approximations to ω. By parallel repetition, one can drive down the constants to whatever you want.

An alternative way to view what we did: a general optimal strategy allows Alice to choose any assignment she wants for a clause. The above analysis shows that actually, she's forced to play with assignments that are consistent with Bob's assignment.

Question: does this automatically imply that ω^* is NP-hard to compute? After all quantum games are "harder," right?

3.3 Quantum games: soundness and entanglement

You can think of the C-V game as a kind of proof system. Alice and Bob are collectively trying to prove to the referee that they have a satisfying assignment to ϕ. Above, we showed that the proof system is *sound*: the referee cannot be convinced that $\text{val}(\phi)$ is high if there exists no such assignment. (Story about the students and the flat tire)

From this perspective, it seems like entanglement could screw up soundness, and these games could actually be easier! This was noticed by Cleve, Høyer, Toner, Watrous.

One example of this is the magic square game. Consider 9 bits in a table as

x_1	x_2	x_3
x_4	x_5	x_6
x_7	x_8	x_9

Alice is asked about a row and Bob is asked about a column. The players win if Alice's three bits have parity 0, Bob's three bits have parity 1, and they agree on the coordinate they have in common. No assignment of 9 bits satisfies this, so the optimal winning probability is < 1 (in fact it is $8/9$).

However there is an entangled strategy which achieves value one. The players assign the following operators to each bit. They are chosen so that operators in that same row or column commute, so can be simultaneously measured.

XI	IX	XX
IZ	ZI	ZZ
$-XZ$	$-ZX$	YY

They measure these on a shared state $|\text{EPR}\rangle^{\otimes 2}$. Observe that the product of the operators in any row is II and in any column is $-II$. This implies that the entangled value of the game is 1.

Also: 2-player XOR games are in P, whereas their classical counterparts correspond to MAX-CUT, and are NP-hard.

3.4 Kempe Kobayashi Matsumoto Toner Vidick

Turns out there is a trick to preserve soundness against entangled provers: use monogamy of entanglement. Let's play an extended clause-variable game, with 3 players. In the game, two players are chosen at random to play the roles of Alice and Bob from the C-V game, and the third is ignored.

It is clear that if $\text{val}(\phi) = 1$, then $\omega^*(G'_{cv}(\phi)) = 1$, for the same reason as before: just play according to a correct assignment. But the surprising thing is that this game is that this game is sound against entangled provers. More precise, if $\text{val}(\phi) \leq 1 - \delta$, then $\omega^*(G'(\phi)) \leq 1 - \text{poly}(\frac{\delta}{n})$.

We're going to sketch the proof of this. First, let's introduce notation to describe an entangled strategy for this game. WLOG we can assume that the shared state is pure, the strategy is symmetric (that is, the state is invariant under permutation, and players 1,2,3 all use the same operators when receiving the same questions), and the measurements applied are projective. Thus, a strategy can be specified by a tripartite state $|\Psi\rangle$ and PVMS $\{A_i^{a_1,\ldots,a_k}\}$ (corresponding to "Alice" queries), and $\{B_j^b\}$ (corresponding to "Bob" queries).

In the calculations, I'll also assume a tensor product structure, although this isn't necessary.

Now, for the proof. Write the probability of success of a strategy

$$p_{\text{success}} = \frac{1}{6}\left(\frac{1}{m}\sum_{i\in[m]}\frac{1}{k}\sum_{\ell\in[k]}\sum_{(a_1,\ldots,a_k)\vdash C_i}\left(\langle\Psi| A_i^{a_1,\ldots,a_k}\otimes B_\ell^{a_\ell}\otimes \text{Id}\,|\Psi\rangle + \text{all permutations}\right)\right).$$

If $p_{\text{success}} \geq 1 - \epsilon$, then each term in the sum is at least $1 - 6mk\epsilon$:

$$\sum_{(a_1,\ldots,a_k)\vdash C_i} \langle\Psi| A_i^{a_1,\ldots,a_k} \otimes B_\ell^{a_\ell} \otimes \mathrm{Id} \,|\Psi\rangle \geq 1 - 6mk\epsilon.$$

Now, note that A_i^a and B_ℓ^b are both projectors. So the above expression implies that for any a_ℓ,

$$\sum_{a\vdash C_i} \| A_i^a \otimes \mathrm{Id} \otimes \mathrm{Id} \,|\Psi\rangle - \mathrm{Id} \otimes B_\ell^{a_\ell} \otimes \mathrm{Id} \,|\Psi\rangle \,\|^2 := d(\sum_{a\vdash C_i} A_i^a \otimes \mathrm{Id} \otimes \mathrm{Id}, \mathrm{Id} \otimes B_\ell^{a_\ell} \otimes \mathrm{Id})^2 \leq 6mk\epsilon.$$

This function $d(\cdot,\cdot)$ is called the state-dependent distance. It is the natural thing to use here, since we cannot control what the measurement operators do on parts of the Hilbert space where the state has no support. You should interpret this as saying that if you measure Bob's part of the state and obtain outcome a_ℓ, when you measure Alice's part of the state you will find a SAT assignment to the clause that is consistent whp.

By performing this kind of analysis, you can also show in state-dep distance that

$$B_\ell^b \otimes \mathrm{Id} \otimes \mathrm{Id} \approx \mathrm{Id} \otimes B_\ell^b \otimes \mathrm{Id}, \tag{3.1}$$

and so on for all permutations. This means that if you measure Bob's operators on two different systems, you get the same answer with high probability.

Putting these together, we basically do what we did for the classical C-V game. We show that we can extract from Bob's strategy a classical assignment that satisfies a high fraction of all the clauses. Instead of writing down a fixed assignment, we give a prob. dist.

$$p(a_1,\ldots,a_n) = \| B_n^{a_n} \ldots B_2^{a_2} B_1^{a_1} \otimes \mathrm{Id} \otimes \mathrm{Id} \,|\Psi\rangle \,\|^2. \tag{3.2}$$

We claim that an assignment sampled from this distribution satisfies a high fraction of the clauses on average. The way to do this: prob that an assignment generated this way satisfies a clause is approx the prob that Alice's operator would satisfy that clause. But this is known to be high.

To see this, consider one particular clause. We use (3.1) to move the $B_i^{a_i}$ in (3.2) not in that clause from player 1 to player 2. We can then sum over them (using $\sum_a B_i^a = I$) to eliminate them. We are left with the product of B_i^a corresponding to a clause. Then we can convert this into an appropriate clause operator A_i^a.

2-player version The effect of monogamy of entanglement doesn't require 3 players. We can simulate it using "oracularization". In this version of the game, Alice is asked for a clause i with variables $(a_1,\ldots,a_k)$, and Bob for two variables: one from the clause (say j), and one at random (say l). But he doesn't know which is which (say they're always in lexicographic order). You can show that Bob's measurement operators for different vars have to approximately commute, which is all that is needed for the analysis above.

Let's suppose that Bob's measurement operators are $B_{l,j}^{b_l,b_j}$ corresponding to questions l,j and answers b_l,b_j. Assume WLOG these are symmetric under exchange of l,j. Define the "marginals" to be

$$C_l^b = \frac{1}{n} \sum_{j,b_j} B_{l,j}^{b_l,b_j}. \tag{3.3}$$

Then we can show that $C_l C_{l'} \approx C_{l'} C_l$ in the state-dependent distance.

3.5 Generalizations: rigidity and self-testing

The above analyses are instances of "rigidity" results for games: we show that any strategy that is close to the optimal strategy must be close to a target strategy—in our case, a classical strategy. However, the analysis given above does not yield constant-factor hardness. A successful approach to get this is to use techniques from the classical PCP theorem, which we review before continuing.

PCP theorem The theorem states that for NP-hard CSPs, there exists a constant $s < 1$ such that the $(1, s)$-CSP is NP-hard. Let's focus on 3SAT. Originally we saw that a $(1, 1-1/m)$-approximation is NP-hard (this is the Cook-Levin thm). Random guessing means that a $(1, 7/8)$ approximation is in P (this can also be derandomized). PCP states that a $(1, 0.99)$ approximation is NP-hard. (In fact this can be improved to $(1, 7/8 + \epsilon)$.

The high level idea is that we take a 3SAT instance ϕ and map it to a new one ϕ' such that if ϕ is satisfiable then ϕ' is too, and if ϕ is unsatisfiable then ϕ' has at least a constant fration of clauses unsatisfied by any assignment. This map is based on a "low-degree test" which verifies that a function is close to a low-degree polynomial.

Using this idea, [Ito-Vidick '12] and [Vidick'13] showed that a $(1 - \epsilon, 1/2 + \epsilon)$-approximation to 3-player XOR games is NP-hard. [Natarajan-Vidick '17] also showed that a $(1, 1/2)$-approximation to 2-player games (not XOR) is NP-hard. Basically, instead of asking players for the assignment directly, we query bits of a robust *encoding* of it. Specifically, one encodes an assignment a by finding a low-degree polynomial over a finite field g_a whose values at some fixed set of n points encode a_i. Then we can query g_a on other points and interpolate to recover the assignment.

That was rigidity for classical strategies. But this approach seems limited to proving NP-hardness. Can we go further, by exploiting the quantum power of *honest* players? Yes: use a quantum analogue of a CSP, called a local Hamiltonian.

This will go beyond NP to a class called QMA, which means the class of yes-no questions where "yes" instances can be verified with a poly-time quantum computer given a witness $|\psi\rangle$ that has $\text{poly}(n)$ qubits. Here n is the size of the input in bits.

Just as 3SAT is the canonical NP-complete problem, there is a natural analogue for QMA called the local Hamiltonian problem. The input is a Hamiltonian

$$H = \frac{1}{m} \sum_{i=1}^{m} H_i, \tag{3.4}$$

where each H_i acts on 3 qubits and satisfies $0 \leq H_i \leq I$. The problem is to determine whether $\lambda_{\min}(H)$ is $\leq a$ or $\geq b$ given the promise that one of these holds. If $b - a = 1/\text{poly}(n)$ then this is QMA-complete. It is an open question - known as the quantum PCP conjecture - that this is also QMA-hard when $b - a$ is constant.

What does this have to do with ω^*? If the honest provers use entanglement then perhaps they can prove that the ground-state energy of a Hamiltonian is small. Specifically our goal is to show that estimating ω^* is QMA-hard by mapping a Hamiltonian H to a game $\mathcal{G}(H)$ such that $\omega^*(\mathcal{G}(H)) \geq 1 - \epsilon \Leftrightarrow \lambda_{\min}(H) \leq \delta$.

In [Fitzsimons-Vidick '14] and [Ji '15] it was proved that a 1/poly approximation to ω^* for 4-player games is NP-hard.

Games qPCP conjecture. This conjectures that even a constant approximation to ω^* is QMA-hard.

The proof technique is based on self-testing. This means that if $\text{val}(|\psi\rangle, M) \geq 1 - \epsilon$ then we must have $|\psi\rangle \approx |\psi_{\text{ideal}}\rangle$, where this "$\approx$" needs to still allow for $|\psi\rangle, M$ to be jointly rotated, and for the possibility of ancilla systems that are approximately ignored by M.

Recent work by [Natarajan-Vidick '17] improved the EPR self testing to certify n EPR pairs using $\text{poly}\log(n)$ communication. This self-test is based on low-degree testing. Using this we can show that the Hamiltonian qPCP conjecture implies the games qPCP conjecture.