Princeton, Spring 09 | Projection PCP

Lecture 9: Combinatorial Transformations on Projection Games: Right Degree Reduction

Lecturer: *Dana Moshkovitz* | Scribe: *Moritz Hardt*

1 Right Degree Reduction

Recall that we think of a projection game $\mathcal{G}$ as a bipartite graph on the vertex set (A, B) so that for every edge (a, b) there is a constraint $f_{a,b} \colon \Sigma_A \to \Sigma_B$.

In this lecture, we will see a combinatorial transformation of a projection game which reduces the right degree of the instance, i.e., the degree of vertices in B, to a qunatity that depends on the error ε alone, and not on the instance size n. At the same time, we maintain soundness and completeness.

Theorem 1 (Right degree reduction)
Let $\mathcal{G}$ be a projection game so that the average right degree is D. Then, for any $d > 0$, there is an efficient transformation of $\mathcal{G}$ into $\mathcal{G}'$ such that:

1. *The degree of all the right vertices in $\mathcal{G}'$ is d.*
2. $\mathrm{val}(\mathcal{G}) = 1 \implies \mathrm{val}(\mathcal{G}') = 1$ *(Completeness)*
3. $\mathrm{val}(\mathcal{G}) \le \epsilon \implies \mathrm{val}(\mathcal{G}) \le \epsilon + O(\frac{1}{\sqrt{d}})$ *(Soundness)*

Furthermore, the size of $\mathcal{G}'$ is bounded by $|\mathcal{G}'| \le D|\mathcal{G}|$.

Remarks:

1. We take $d = \mathrm{poly}(1/\epsilon)$ so that there is no dependence on n.
2. After right degree reduction, player B has a lot more information about player A's question: player B can always guess one of d possibilities.
3. The soundness tradeoff is essentially optimal. Indeed, for the reason mentioned above, the degree has to be at least $1/\epsilon$ if we're shooting for soundness ϵ.

1.1 Background on Expander graphs

For the sake of completeness we recall the notion of an expander graph and state the well-known expander mixing lemma.

Definition 1 (Expander graph) *A d-regular graph $H = (V, E)$ is an (n, d, λ)-expander, if the second largest eigenvalue of H's adjacency matrix is at most λ.*

We remark that there is an efficient algorithm that given as input size n, degree $d > 3$ and second eigenvalue $\lambda = \Omega(\frac{1}{\sqrt{d}})$ (for some specific function $\Omega(\frac{1}{\sqrt{d}})$), constructs an (n, d, λ)-expander in time poly(n, d). Graphs with $\lambda = O(\frac{1}{\sqrt{d}})$ are called *Ramanujan graphs.*

Lemma 2 (Expander Mixing Lemma)
Let $H = (V, E)$ be an (n, d, λ)-expander. Then, for every two sets $X, Y \subseteq V$, we have

$$\left| E(X,Y) - \frac{d}{n}|X||Y| \right| \le \lambda\sqrt{|X||Y|}. \tag{1}$$

Proof: Denoting by x, y the characteristic vectors of X, Y we observe that $E(X, Y) = x^T Ay$ where A is the adjacency matrix of H. On the other hand, we can write x, y as $x = \sum_i \alpha_i u_i, y = \sum_i \beta_i u_i$ where $\{u_i\}_{i=1}^n$ is an orthonormal eigenbasis of A. Let $d = \lambda_1 \ge \cdots \ge \lambda_n$ be the eigenvalues of A, where $\lambda = max\,\{|\lambda_2|, |\lambda_n|\}$.

We then see

$$\begin{aligned} E(X,Y) = x^T Ay &= \sum_{i,j} \alpha_i \beta_j \lambda_i \lambda_j \langle u_i, u_j \rangle \\ &= \sum_i \alpha_i \beta_i \lambda_i \\ &\le \alpha_1 \beta_1 \lambda_1 + \lambda \sum_i \alpha_i \beta_i \\ &= \frac{d}{n}|X||Y| + \lambda \sum_i \alpha_i \beta_i, \end{aligned}$$

where we used that $\lambda_1 = d$ and $u_1 = \mathbf{1}$ (the all one's vector). Taking absolute values we get that

$$\left| E(X,Y) - \frac{d}{n}|X||Y| \right| \le \lambda \sum_i |\alpha_i||\beta_i|.$$

It only remains to observe that

$$\sum_i |\alpha_i||\beta_i| \le \sqrt{\sum_i \alpha_i^2}\sqrt{\sum_i \beta_i^2} = \sqrt{|X||Y|}$$

via Cauchy-Schwarz inequality. □

1.2 Idea of the construction

Picture a single vertex b on the right hand side, let D_b denote its degree and $N(b)$ its neighborhood in A. We will make D_b copies of b, call those vertices $C(b)$. How do we connect those copies to $N(b)$? Suppose, we pick a *random* d-regular bipartite graph between $N(b)$ and $C(b)$. Every edge between $a \in N(b)$ and $b' \in C(b)$ will simply inherit the constraint $f_{a,b}$.

It is easy to see that this construction will satisfy completeness, since a truthful prover can simply give all vertices in $C(b)$ the same intended label.

To argue soundness, suppose a dishonest prover labeled a large fraction of $C(b)$ using several different labels. We can then extract a good (randomized) labeling to the old vertex

b by picking at random one these labels, denoted σ, with probability proportional to the fraction of vertices in $C(b)$ that have the label σ. Since the graph between $N(b)$ and $C(b)$ is random and regular, this randomized assignment will have the same *expected* value as in $\mathcal{G}$.

The actual construction will be deterministic and hence use expander graphs.

1.3 Construction

We will now construct $\mathcal{G}'$: As described earlier, we replace B by $B' = \{(b,i)\colon b \in B, i \in [D_b]\}$. We denote the copies of b by $C(b) = \{(b,i)\colon i \in [D_b]\}$. For each $b \in B$, let H_b be a bipartite $(D_b, d, \frac{c_0}{\sqrt{d}})$-expander for some constant c_0. Connect the two vertex sets $N(b)$ and $C(b)$ by the edges of H_b, i.e., if a is the j'th neighbor of b and (j,i) is an edge of H_b then $(a,(b,i))$ is an edge in the new graph. Each edge $(a,(b,i))$ for $i \in [D_b]$, inherits the constraint $f_{(a,b)}$.

It is clear by construction that the right degree of $\mathcal{G}'$ is d and the instance size is increased by at most a factor of D. The left degree is multiplied by d. If the graph were left-regular, then it remains left-regular.

Soundness and completeness are proven next.

Proof of Theorem 1

Completeness: Given a labeling $\ell\colon B \to \Sigma$ that satisfies all constraints in $\mathcal{G}$, consider the labeling $\ell'(b,i) = \ell(b)$. It will satisfy all constraints in $\mathcal{G}'$.

Soundness: The claim is proven contrapositively. Fix a labeling of A and B' that achieves in $\mathcal{G}'$ value greater than $\epsilon + \frac{c_1}{\sqrt{d}}$ where c_1 is a sufficiently large constant, say, $c_1 > c_0$. We will then show how to construct a labeling to A and B that achieves in $\mathcal{G}$ value greater than ϵ.

Let $\ell_{B'}\colon B' \to \Sigma$ denote the labeling of B', we define a new labeling $\ell_B\colon B \to \Sigma$ as follows: The value $\ell_B(b)$ is defined by randomly picking $i \in [D_b]$ and choosing label $\ell_{B'}(b,i)$.

To analyze the value of the labeling in $\mathcal{G}$, let us define the following partitions:

- $X_{b,\sigma}$ are those vertices in $N(b)$ which "vote" for σ. Formally,

$$X_{b,\sigma} = \{a \in N(b)\colon f_{a,b}(\ell_A(a)) = \sigma\},$$

  where ℓ_A is the labeling of the vertices in A.

- $Y_{b,\sigma}$ are those copies of b which are labeled by σ, i.e.,

$$Y_{b,\sigma} = \{i \in [D_b]\colon \ell_{B'}(b,i) = \sigma\}.$$

Notice that the expected value of ℓ_B is given by:

$$\frac{1}{|E|}\sum_{b\in B}\sum_{\sigma\in\Sigma}|X_{b,\sigma}| \cdot \frac{|Y_{b,\sigma}|}{D_b} = \frac{1}{|E'|}\sum_{b\in B}\sum_{\sigma\in\Sigma}|X_{b,\sigma}||Y_{b,\sigma}|\frac{d}{D_b}.$$

Hence, we can now lower bound the value of $\mathcal{G}$ using the Expander Mixing Lemma:

$$\begin{aligned}
\mathrm{val}(G) &\geq \frac{1}{|E'|}\sum_{b\in B}\sum_{\sigma\in\Sigma}|X_{b,\sigma}||Y_{b,\sigma}|\frac{d}{D_b} \\
&\geq \frac{1}{|E'|}\sum_{b\in B}\sum_{\sigma\in\Sigma}\left(E(X_{b,\sigma},Y_{b,\sigma}) - \lambda\sqrt{|X_{b,\sigma}|}\sqrt{|Y_{b,\sigma}|}\right) && \text{(by (1))} \\
&\geq \epsilon + \frac{c_1}{\sqrt{d}} - \frac{\lambda}{|E'|}\sum_{b\in B}\sum_{\sigma\in\Sigma}\sqrt{|X_{b,\sigma}||Y_{b,\sigma}|} \\
&\geq \epsilon + \frac{c_1}{\sqrt{d}} - \frac{\lambda}{|E'|}\sum_{b\in B}\sqrt{\sum_{\sigma\in\Sigma}|X_{b,\sigma}|}\sqrt{\sum_{\sigma\in\Sigma}|Y_{b,\sigma}|} && \text{(Cauchy-Schwarz)} \\
&\geq \epsilon + \frac{c_1}{\sqrt{d}} - \frac{c_0}{\sqrt{d}} > \epsilon
\end{aligned}$$