CS 880: Complexity of Counting Problems 03/08/2012

Lecture 14:

Instructor: Jin-Yi Cai Scribe: Yanpei (Nick) Liu

A matrix $A \in \{0,1\}^{m \times n}$ is said to be decomposable if there exists I and J, $I, J \neq 0$, $I \subseteq [m], J \subseteq [n]$ such that $A[I, \bar{J}] = 0$ and $A[\bar{I}, J] = 0$ as shown in Fig. 1. Furthermore, at least one of them needs to be non-empty, i.e., either $I \neq [m]$ or $J \neq [n]$.

We can think of this as a bipartite graph on $[m] \times [n]$. Then the decomposability is equivalent to saying the underlying bipartite graph is disconnected. The connected components constitute blocks, which are made up of pairs $(I_1, J_1), (I_2, J_2), \ldots, (I_k, J_k)$ with no edges between I_i and J_j for all $i \neq j$. Note that some I_s might be empty, in which case the corresponding J_s must be singleton isolated points.

We can generalize this to matrices in $\mathbb{C}$. Replace the non-zero elements by 1. For symmetric $m \times m$ matrix, it is equivalent to the graph G on $[m]$ as adjacency matrix in the following ways. If $I_1 \cap J_1$ is non-empty, then $I_1 = J_1$. and the $|I_1|$ points of G correspond to a single connected component (of $2|I_1|$ points) of the graph H on $[m]$. If $I_1 \cap J_1$ is empty, then there is a bipartite component of size $|I_1| + |J_1|$, of the graph G on $[m]$, which corresponds to two connected components, one of size $|I_1|$ on the left and $|J_1|$ on the right, and another one of size $|J_1|$ on the left and $|I_1|$ on the right.

Given an $m \times m$ symmetric matrix A, recall that the graph homomorphism function is defined as:

$$Z_A(G) = \sum_{\sigma: V \to [m]} \prod_{(u,v) \in E} A_{\sigma(u), \sigma(v)}. \tag{1}$$

In the following subsequence lectures we will prove that for every A, Z_A is either in P or #P-hard. The #P hardness is proved by showing there exists a reduction. In the case it is tractable, we give an algorithm. We further prove that the tractability criterion on A is polynomial-time decidable.

We build our work on the ones by Dyer and Greenhill as well as Bulatov and Grohe. Recall the theorem by Dyer and Greenhill, which states as follows:

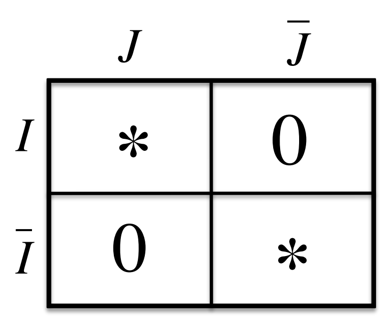

Figure 1:

Theorem 1 (Dyer and Greenhill)**.** *A symmetric* $\{0,1\}$ *matrix* A *defines a tractable graph homomorphism* $Z_A(\cdot)$ *if and only if* A *is rectangular or if and only if the connected components of* $G(A)$ *consists of complete bipartite graphs or cliques with loops or isolated points.*

Bulatov and Grohe then extends this result to all non-negative symmetric matrix A (every block of all 1's becomes rank 1 blocks). From their results, $Z_A(\cdot)$ is computable in P if each block of A has rank at most one, and #P-hard otherwise.

To be more precise, if G has connected component $G_1, G_2, \ldots, G_k$, then computing $Z_A(G)$ is the same as computing $\prod_i^k Z_A(G_i)$. If the problem on any G is tractable, then it is in particular tractable on connected subgraphs. We can decompose the underlying graph A as a direct sum of A_i, each being a connected component. Then for a connected graph G, we must have

$$Z_A(G) = \sum_{j=1}^{l} Z_{A_j}(G). \tag{2}$$

This implies if $Z_{A_j}(\cdot)$ is tractable for all j, then $Z_A(\cdot)$ is tractable. However we would like to have the converse as well, i.e., if one of $Z_{A_j}(\cdot)$ is hard, so is $Z_A(\cdot)$. That is the reduction from $Z_{A_j}(\cdot)$ to $Z_A(\cdot)$ for all j.

Let $wt_A(\sigma) = \prod_{\sigma\in\{u,v\}} A_{\sigma(u),\sigma(v)}$. Given a symmetric matrix $C \in \mathbb{C}^{m\times m}$ and let $\mathfrak{D} = [D^{[0]}, D^{[1]}, \ldots, D^{[N-1]}]$ be a sequence of diagonal matrices in $\mathbb{C}^{m\times m}$ for some $N \geq 1$ (we use $D_i^{[r]}$ to denote the $(i,i)^{th}$ entry of $D^{[r]}$. We define the following EVAL$(C, \mathfrak{D})$ problem

$$Z_{C,\mathfrak{D}}(G) = \sum_{\sigma: V\to[m]} wt_{C,\mathfrak{D}}(\sigma), \tag{3}$$

where $wt_{C,\mathfrak{D}}(\sigma)$ is defined as:

$$wt_{C,\mathfrak{D}}(\sigma) = \left(\prod_{(u,v)\in E(G)} C_{\sigma(u),\sigma(v)}\right)\left(\prod_{v\in V} D_{\sigma(v)}^{[\deg(v) \bmod N]}\right). \tag{4}$$

Suppose C is the bipartisation of an $m \times n$ matrix F, i.e.,

$$C = \begin{pmatrix} 0 & F \\ F^T & 0 \end{pmatrix}. \tag{5}$$

For any graph G and vertex u in G, we define $Z^{\rightarrow}_{C,\mathfrak{D}}(G,u)$ and $Z^{\leftarrow}_{C,\mathfrak{D}}(G,u)$ as follows. Let $\mathcal{X}_1$ denote the set of $\sigma : V \to [m+n]$ with $\sigma(u) \in [m]$ and $\mathcal{X}_2$ denote the set of σ with $\sigma(u) \in [m+1 : m+n]$, then we have

$$Z^{\rightarrow}_{C,\mathfrak{D}}(G,u) = \sum_{\sigma\in\mathcal{X}_1} wt_{C,\mathfrak{D}}(\sigma) \tag{6}$$

$$Z^{\leftarrow}_{C,\mathfrak{D}}(G,u) = \sum_{\sigma\in\mathcal{X}_2} wt_{C,\mathfrak{D}}(\sigma) \tag{7}$$

Then it follows that $Z_{C,\mathfrak{D}}(G) = Z_{\overrightarrow{C,\mathfrak{D}}}(G,u) + Z_{\overleftarrow{C,\mathfrak{D}}}(G,u)$.

Let A be an $m \times m$ symmetric matrix, we consider a EVALP(A) problem as follows. The input is a triple (G, w, i) where $G = (V, E)$ is an undirected graph, $w \in V$ is a vertex and $i \in [m]$. The output is:

$$Z_A(G, w, i) = \sum_{\substack{\sigma: V \to [m] \\ \sigma(w) = i}} wt_A(\sigma). \tag{8}$$

Then we have the First Pinning Lemma, stated as follows:

Lemma 1 (First Pinning Lemma)**.** *We have EVALP(A) $\equiv$ EVAL(A).*

It is easy to see that EVALP(A) $\geq$ EVAL(A). We need to prove the other direction also holds. We define the following equivalence relation $(i \sim j)$ over $[m]$:

$$i \sim j \quad \text{if } Z_A(G, w, i) = Z_A(G, w, j) \ \forall G, w \in V(G). \tag{9}$$

This equivalence relation divides the set $[m]$ into s equivalent classes $\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_s$ for some positive integer s. For any $t \neq t' \in [s]$, there exists a pair $P_{t,t'} = (G, w)$ that distinguish them, i.e.,

$$Z_A(G, w, i) \neq Z_A(G, w, i') \quad \forall i \in \mathcal{A}_t \quad \forall i' \in \mathcal{A}_{t'}. \tag{10}$$

Now for any subset $S \subseteq [s]$, we define

$$Z_A(G, w, S) = \sum_{\substack{\sigma: V \to [m] \\ \sigma(w) \in \bigcup_{t \in S} \mathcal{A}_t}} wt_A(\sigma). \tag{11}$$

To prove Lemma 1, we first present the following shrinking lemma:

Lemma 2 (Shrinking Lemma)**.** *If $S \subseteq [s]$ and $|S| \geq 2$, then there exists a partition $\{S_1, \ldots, S_k\}$ of S for some $k > 1$ and EVAL$(A, S_l) \leq$ EVAL(A, S) for all l.*

Proof. Let $t \neq t'$ be two integers in S. We let $P_{t,t'} = (G^*, w^*)$ For all $a, b \in S$ we define a equivalence relation $\sim^*$ over S:

$$a \sim^* b \quad \text{if } Z_A(G^*, w^*, i) = Z_A(G^*, w^*, j) \quad i \in \mathcal{A}_a, \ j \in \mathcal{A}_b. \tag{12}$$

The foregoing relation is independent of the choice of i and j and it gives us equivalence partition $\{S_1, S_2, \ldots, S_k\}$ of S. We let X_a denote $Z_A(G^*, w^*, i)$ where $i \in \mathcal{A}_a$ and X_a is independent of the choice of i. We define X_b similarly. Then $X_b \neq X_a$ if $a \nsim^* b$. Let G be an undirected graph and w be a vertex. For each $p : 0 \leq p \leq k-1$, we construct a sequence of graphs $G^{[p]}$ where $G^{[p]}$ is the disjoint union of G and p independent copies of $P_{t,t'}$, except the w in the input G and all w^* in the copies of $P_{t,t'}$ as one single vertex $w' \in V^{[p]}$. We then have the following collection of equations:

$$Z_A(G^{[p]}, w', S) = \sum_{l=1}^{k} Z_A(G, w, S_l) X_l^p, \tag{13}$$

where S_l's are equivalence classes. Since $X_b \neq X_a$ for all $a \not\sim^* b$, this is a Vandermonde system and we can solve it to get $Z_A(G, w, S_l)$ for all $l \in [k]$. This gives us a polynomial time reduction from $\mathrm{EVAL}(A, S_l)$ to $\mathrm{EVAL}(A, S)$ for every $l \in [k]$. □

We now prove Lemma 1.

Proof. We apply Lemma 2 to $S = [s]$. By Lemma 2, there exists a partition $\{S_1, \ldots, S_k\}$ of S, for some $k > 1$, such that $\mathrm{EVAL}(A, S_d) \leq \mathrm{EVAL}(A, S)$ for all $d \in [k]$. Consider some $t \in [s]$, without loss of generality, assume $t \in S_1$ (or we can always shrink S_1 using Lemma 2 until t is the only element in S_1). When this is true, we simply have

$$Z_A(G, w, i) = \frac{1}{|\mathcal{A}_t|} \cdot Z_A(G, w, \{t\}). \tag{14}$$

This completes the proof of Lemma 1. □