CS 880: Complexity of Counting Problems	1/31/2012

Lecture 3: Baby Dichotomy Theorem

Instructor: Jin-Yi Cai   Scribe: Tyson Williams

After finishing the proof of #P-harness that we started last time, we embark on proving our first dichotomy theorem for a very restricted case.

1 Polynomial Interpolation

In the previous lecture, we introduced the proof technique of polynomial interpolation and were using it to show that the problem Holant$([0,1,1,0])$ is #P-hard over 3-regular graphs. We finish the last bit of that proof now.

Proof cont. Last time, we had just finished constructing our Vandermonde system and begin talking about our recursive gadget construction of the N_i's, which is shown in Figure 9 of [2]. The signature of N_i is $[a_i, b_i, b_i, a_i]$. For the Vandermonde system to be of full rank, we need some polynomially many N_i with distinct ratios a_i/b_i. The entries in the signature of N_{i+1} can be expressed as a linear combination of the entries in N_i. Thus, there exists a matrix A such that

$$\begin{bmatrix} a_{i+1} \\ b_{i+1} \end{bmatrix} = A^{i+i} \begin{bmatrix} a_0 \\ b_0 \end{bmatrix}.$$

By some careful counting, one can determine that $A = \begin{bmatrix} 20 & 60 \\ 20 & 75 \end{bmatrix}$. With more calculation, one can also determine that the eigenvalues λ and μ of A are distinct. Therefore, we can diagonalize A as $A = M \begin{bmatrix} \lambda & 0 \\ 0 & \mu \end{bmatrix} M^{-1}$. The reason for wanting to express A in this form is that when we raise A to some power i, then $A^i = M \begin{bmatrix} \lambda^i & 0 \\ 0 & \mu^i \end{bmatrix} M^{-1}$.

Now by previous work, there are three conditions that suffice to show that the N_i's will have distinct ratios a_i/b_i. They are

1. $\begin{bmatrix} a_0 \\ b_0 \end{bmatrix}$ not a column eigenvector,
2. $\det A \neq 0$, and
3. λ/μ not a root of unity.

For higher dimensional interpolation (i.e. the matrix A is of a higher dimension than 2), the correct generalization of the first condition is that $\begin{bmatrix} a_0 \\ b_0 \end{bmatrix}$ is not orthogonal to any row eigenvector of A. For 2-by-2 matrices, these two statements are equivalent. The reason for the third condition is that otherwise our construction will eventually produce the same N_i's (up to a constant factor, which does not matter since it is cancelled in the ratio).

If we write our diagonalization of A as $AM = M \begin{bmatrix} \lambda & 0 \\ 0 & \mu \end{bmatrix}$, then we see that the columns of M are the column eigenvectors of A. If we write it as $M^{-1}A = \begin{bmatrix} \lambda & 0 \\ 0 & \mu \end{bmatrix} M^{-1}$, then we see that the rows of M^{-1} are the row eigenvectors of A. Let $\mathbf{u}$ and $\mathbf{v}$ be the column eigenvectors of A, thus $M = [\mathbf{u}, \mathbf{v}]$. Then we can write $\begin{bmatrix} a_0 \\ b_0 \end{bmatrix} = s\mathbf{u} + t\mathbf{v}$ where $s, t \neq 0$ since $\begin{bmatrix} a_0 \\ b_0 \end{bmatrix}$ really is not a column eigenvector of A. Then

$$\begin{bmatrix} a_i \\ b_i \end{bmatrix} = A^i \begin{bmatrix} a_0 \\ b_0 \end{bmatrix} = M \begin{bmatrix} \lambda^i & 0 \\ 0 & \mu^i \end{bmatrix} M^{-1} \begin{bmatrix} a_0 \\ b_0 \end{bmatrix} = M \begin{bmatrix} \lambda^i & 0 \\ 0 & \mu^i \end{bmatrix} \begin{bmatrix} s \\ t \end{bmatrix} = M \begin{bmatrix} \lambda^i s \\ \mu^i t \end{bmatrix}.$$

Now for $i < j$, a_i/b_i will be distinct from a_j/b_j if the determinant of $\begin{bmatrix} a_i & a_j \\ b_i & b_j \end{bmatrix}$ is nonzero. Well,

$$\begin{aligned} \det \begin{bmatrix} a_i & a_j \\ b_i & b_j \end{bmatrix} &= \det M \cdot \det \begin{bmatrix} \lambda^i s & \mu^i t \\ \lambda^j s & \mu^j t \end{bmatrix} \\ &= \det M \cdot (\lambda^i \mu^j st - \lambda^j \mu^i st) \\ &= st\lambda^i \mu^i \det M \cdot (\mu^{j-i} - \lambda^{j-i}) \\ &\neq 0 \end{aligned}$$

since none of these terms is zero. In particular $(\mu^{j-i} - \lambda^{j-i})$ is not zero, since otherwise μ/λ is a root of unity, which it is not. Therefore, interpolation succeeded and the proof is complete. □

2 Dichotomy Theorem

Our goal for this section is to prove the following dichotomy, which can be found in [1].[1]

Theorem 1 (Theorem 8.3 in [1])**.** *Every counting problem* Holant$([x_0, x_1, x_2] \mid [y_0, y_1, y_2, y_3])$, *where* $[x_0, x_1, x_2]$ *and* $[y_0, y_1, y_2, y_3]$ *are Boolean signatures, is either*

- *in* P,
- #P-*complete but solvable in* P *for planar graphs, or*
- #P-*complete even for planar graphs.*

[1]A more complete version titled "Holographic Reduction, Interpolation and Hardness" was handed out in class and is also available on the class website. In this version, it is Theorem 6.3.

To know which case each pair of signatures is in, see the table about the statement of the theorem (in either version of the paper). The notation Holant($[x_0, x_1, x_2] \mid [y_0, y_1, y_2, y_3]$) means that the input graphs are bipartite where all the vertices in one partite set must take the signature $[x_0, x_1, x_2]$ and the vertices in the other partite set must take the signature $[y_0, y_1, y_2, y_3]$. This forces the input graphs to be (2,3)-regular so as to match the arity of these two signatures.

The tractability for general graphs (that is, (2,3)-regular graphs that are not necessarily planar) is covered by holographic algorithms with Fibonacci signatures as well as some trivial cases. The cases tractable for planar graphs but #P-hard in general are covered by holographic algorithms with matchgates. To show #P-hardness for the remaining pairs of signatures, even when restricted to planar graphs, the gadget constructions must be planar.[2]

Consider the problem Holant($[0, 1, 1] \mid [1, 1, 0, 0]$). This is an example of a problem tractable for a trivial reason. The arity two signature on the left says that at least half of the bits must be 1. On the right side, the arity three signature says that at most one-third of the bits can be 1. Since there is no assignment that can satisfy these two conditions simultaneously, the Holant must be 0, thus trivially tractable.

There were many papers published before [1] in this line of work on counting problems. However, this paper was the one where we had the first inkling that dichotomy theorems were possible. When trying to prove a new dichotomy theorem, you do not always know which cases are tractable and which cases are hard. You may start by trying to prove that a particular case is hard. If the hardness proof will not go through, then you switch and try to prove that it is tractable. Eventually you succeed in proving one or the other.

Definition 1. *A symmetric signature* $f = [x_0, x_1, \ldots, x_n]$ *is* non-degenerate *if*

$$\operatorname{rank} \begin{bmatrix} x_0, x_1, \ldots, x_{n-1} \\ x_1, x_2, \ldots, x_n \end{bmatrix} = 2$$

and degenerate *if* $f = [s^0 t^n, s^1 t^{n-1}, \ldots, s^n t^0] = (t, s)^{\otimes n}$.

If Ω is a signature grid with a vertex v assigned a degenerate signature $f = (t, s)^{\otimes k}$ of arity k, then we modify Ω by replacing v with k vertices of degree 1, one attached to each of the edges of v, and assigned the unary signature (t, s) to form a signature gird Ω'. The reason for this modification is that the Holant on each of these signature grids is the same! Thus, if all signatures are degenerate, the Holant becomes a product of $|E|$ many terms (where E is the set of edges in the graph), each term computable in constant time. This shows that the complexity is derived from the non-degenerate signatures.

[2] The gadgets used to prove #P-hardness for the problems that are tractable for planar graphs also happen to be planar. One might expect that the gadgets would be (or even need to be) non-planar. However, this is not the case because we reduce from a problem that is known to be #P-hard for planar graphs.

3 Proving Hardness

Consider the non-degenerate signature $[y_0, y_1, y_2, y_3]$. We can express y_2 and y_3 as linear combinations of y_0 and y_1, a linear recurrence relation of order 2.[3] If the characteristic polynomial of the recurrence relation has distance eigenvalues, then there exists α_1, α_2, β_1, and β_2 such that $y_i = \alpha_1^{3-i}\alpha_2^i + \beta_1^{3-i}\beta_2^i$. If instead there is a double root α, then there exists $A \neq 0$ and B such that $y_i = Ai\alpha^{i-1} + B\alpha^i$ (if $A = 0$, then the signature is degenerate).

For the first case with $y_i = \alpha_1^{3-i}\alpha_2^i + \beta_1^{3-i}\beta_2^i$, we have

$$\begin{aligned}
[y_0, y_1, y_2, y_3]^T &= \begin{bmatrix}\alpha_1\\ \alpha_2\end{bmatrix}^{\otimes 3} + \begin{bmatrix}\beta_1\\ \beta_2\end{bmatrix}^{\otimes 3} \\
&= \begin{bmatrix}\alpha_1 & \beta_1\\ \alpha_2 & \beta_2\end{bmatrix}^{\otimes 3} \left(\begin{bmatrix}1\\ 0\end{bmatrix}^{\otimes 3} + \begin{bmatrix}0\\ 1\end{bmatrix}^{\otimes 3}\right) \\
&= \begin{bmatrix}\alpha_1 & \beta_1\\ \alpha_2 & \beta_2\end{bmatrix}^{\otimes 3} [1, 0, 0, 1]^T,
\end{aligned}$$

so $[y_0, y_1, y_2, y_3]$ is just a holographic mixture of $[1, 0, 0, 1]$, the equality function of arity 3. To keep the Holant value unchanged, we must do the inverse transformation to the signature on the other side, namely $[x_0, x_1, x_2]$. We do not know what this signature becomes under this transformation, we just write

$$[x_0, x_1, x_2]\left(\begin{bmatrix}\alpha_1 & \beta_1\\ \alpha_2 & \beta_2\end{bmatrix}^{-1}\right)^{\otimes 2} = [a_0, a_1, a_2].$$

Thus, our original problem is equivalent to Holant($[a_0, a_1, a_2] \mid [1, 0, 0, 1]$), which is a question about (weighted) graph homomorphisms from 3-regular graphs to a target graph with two vertices, connected by an edge, each with a self loop. Although we started with signature with Boolean weights, notice that the a_i's could be complex numbers!

Now, consider the problem Holant($[0, 1, 1] \mid [1, 0, 0, 1]$). Can you tell what problem this is? It is a well-known problem. The signature on the right is the equality signature, so we might as well think that assignments are happening on those vertices instead of the edges. On the left side is a vertex of degree two. It is "happy" if one or two of its neighbor vertices is selected but "unhappy" if neither of them are selected. This is the #VertexCover problem on 3-regular graphs where the vertex of degree two on the left side is really acting like an edge between two vertices. That is, the signature grid contains the edge-vertex indicent graph of the graph for which vertex covers are being counted. This problem is #P-hard even for planar graphs. Thus, our strategy will be to reduce #VertexCover to our original problem via polynomial interpolation.

[3] Actually, this is only true as long as y_0 or y_1 is not zero. However, if they are both zero, we can y_0 and y_1 and a (trivial) linear combination of y_2 and y_3, which just means that the recurrence relation is in the opposite direction.

From #VERTEXCOVER, we do the constant time holographic reduction using the basis $\begin{bmatrix} \alpha_1 & \beta_1 \\ \alpha_2 & \beta_2 \end{bmatrix}$. This gives us

$$\begin{bmatrix} \alpha_1 & \beta_1 \\ \alpha_2 & \beta_2 \end{bmatrix}^{\otimes 3} [1, 0, 0, 1]^T = [y_0, y_1, y_2, y_3]^T$$

and

$$[0, 1, 1] \left(\begin{bmatrix} \alpha_1 & \beta_1 \\ \alpha_2 & \beta_2 \end{bmatrix}^{-1} \right)^{\otimes 2} = [x, y, x] = g.$$

Let Ω be a signature grid with a (2,3)-regular graph and let V_g be the subset of vertices assigned g with $|V_g| = n$. Then

$$\text{Holant}_\Omega = \sum_{i+j+k=n} x^i y^j z^k c_{ijk},$$

where c_{ijk} is the sum of products of values of all vertices not in V_g with assignments that have exactly i vertices of V_g with the input (0,0), j vertices with (0,1) or (1,0), and k vertices with (1,1). Just as in our first example of polynomial interpolation, we have grouped the exponentially many terms in the Holant sum into polynomially many groups, each of which still have exponentially many terms. This time, the number of groups is $\binom{n+2}{2}$, since we want to partition n items into three parts.

Since we do not actually have the signature g, we replace it with $g_s = [x_s, y_s, z_s]$ to form a signature grid Ω_s where each g_s is constructed using our signatures $[x_0, x_1, x_2]$ and $[y_0, y_1, y_2, y_3]$. Then $\text{Holant}_{\Omega_s} = \sum_{i+j+k=n} x_s^i y_s^j z_s^k c_{ijk}$. Crucially, the c_{ijk} remain the same. Once again, these g_s will be constructed recursively, so that the entries in g_{s+1} can be expressed as a linear combination of the entries in g_s. So,

$$\begin{aligned}
\begin{bmatrix} x_s \\ y_s \\ z_s \end{bmatrix} &= A \begin{bmatrix} x_{s-1} \\ y_{s-1} \\ z_{s-1} \end{bmatrix} \\
&= T^{-1} \begin{bmatrix} \alpha^s & 0 & 0 \\ 0 & \beta^s & 0 \\ 0 & 0 & \gamma^s \end{bmatrix} T \begin{bmatrix} x_0 \\ y_0 \\ z_0 \end{bmatrix} \\
&= T^{-1} \begin{bmatrix} \alpha^s u \\ \beta^s v \\ \gamma^s w \end{bmatrix} \\
&= T^{-1} \begin{bmatrix} u & 0 & 0 \\ 0 & v & 0 \\ 0 & 0 & w \end{bmatrix} \begin{bmatrix} \alpha^s \\ \beta^s \\ \gamma^s \end{bmatrix} \\
&= B \begin{bmatrix} \alpha^s \\ \beta^s \\ \gamma^s \end{bmatrix},
\end{aligned}$$

where we assume that A has distinct eigenvalues and our initial signature $[x_0, y_0, z_0]$ is not orthogonal to any row eigenvectors of A, so $uvw \neq 0$ since the rows of T contain the row eigenvectors of A.

Now, we define a set of strings $\kappa = \{0^i 1^j 2^k \mid i + j + k = n\}$, so $|\kappa| = \binom{n+2}{2}$. Consider the equation

$$\begin{bmatrix} x_s \\ y_s \\ z_s \end{bmatrix}^{\otimes n} = B^{\otimes n} \begin{bmatrix} \alpha^s \\ \beta^s \\ \gamma^s \end{bmatrix}^{\otimes n}.$$

Since $\det B \neq 0$, $\det B^{\otimes n} = (\det B)^{n3^{(n-1)}} \neq 0$, $B^{\otimes n}$ is non-singular. Now, partition the columns according to κ and sum the columns in each partition to form a 3^n-by-$\binom{n+2}{2}$ matrix $\widehat{B^{\otimes n}}$. Because the column sums are over disjoint sets of columns, $\widehat{B^{\otimes n}}$ has full rank. Otherwise, there would also be a nontrivial combination of the columns of $B^{\otimes n}$. Now $\widehat{B^{\otimes n}}$ has at most $\binom{n+2}{2}$ distinct rows since any two rows with indices that give the same value in κ are the same. We remove all duplicate rows to form the $\binom{n+2}{2}$-by-$\binom{n+2}{2}$ matrix $\widetilde{B^{\otimes n}}$, which does not lower the rank. Since the rank of $\widetilde{B^{\otimes n}}$ is at least $\binom{n+2}{2}$, it is a non-singular matrix. Thus, our Vandermonde system will have full rank.

We pick up from here next time.

References

[1] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holographic algorithms by fibonacci gates and holographic reductions for hardness. In *FOCS*, pages 644–653. IEEE Computer Society, 2008.

[2] Jin-Yi Cai, Pinyan Lu, and Mingji Xia. Holographic algorithms by fibonacci gates. *Linear Algebra and its Applications*, 2011.