Lecture 14: Critical Theorem vs. Finite-Field Method vs. Modular Coloring

24 January 2020

Notetaker: Mike Gottstein

Example: The hollow Catalan arrangement. The hollow Catalan arrangement is $\mathscr{A}[\{\pm1\}K_n]$, with gain group $\mathbb{Z}^+$. We calculate the characteristic polynomial by counting the number of proper colorations of the hollow Catalan gain graph $\{\pm1\}K_n$ modulo m, i.e., proper group colorings of the gain graph in the group $\mathbb{Z}_m^+$. We have to choose m carefully: no circle can have gain that is a multiple of m. That preserves the balanced chromatic polynomial, by Theorem **??**. Since n is the largest possible gain of a circle in $\{\pm1\}K_n$, the obvious thing to do is to choose $m > n$.

The calculation is similar to that for the Catalan gain graph, but we have to allow for the fact that vertices may have the same color. Thus, we consider a partition of V into k parts, of which there are $S(n,k)$ (the Stirling number of the second kind). The partition consists of the sets of vertices having the same color; each part has one color, and every part has a different color from every other. The number of ways to color the k parts is the same as for the Catalan gain graph $\{0,\pm1\}K_k$ with k vertices, as in Lecture 13; it is $m(m-k-1)_{k-1}$. We have to multiply this by $S(n,k)$ for the number of k-partitions of V and sum over all possible numbers of parts. We get this:

Proposition 1. *For the hollow Catalan arrangement $\mathscr{C}_n^\circ$ with $n \geq 1$:*

(1) *The characteristic polynomial is*

$$p_{\mathscr{C}_n}(\lambda) = \lambda \sum_{k=1}^{n} S(n,k)(\lambda - k - 1)_{k-1}.$$

(2) *The number of regions is*

$$\sum_{k=1}^{n} S(n,k)(2k)_{k-1} = \sum_{k=1}^{n} S(n,k)k!C_k.$$

(3) *The number of bounded regions is*

$$\sum_{k=1}^{n} S(n,k)(2k-2)_{k-1} = \sum_{k=1}^{n} S(n,k)k!C_{2k-2}.$$

The "finite field" method. The foundation of the finite field method is a theorem of Crapo and Rota. Let's consider an arrangement $\mathscr{A}$ in $\mathbb{A}^n(\mathbb{F}_q)$.

Theorem 2 (Critical Theorem). *The number of points of $\mathbb{A}^n(\mathbb{F}_q)$ not in $\bigcup \mathscr{A}$ is $p_{\mathscr{A}}(q)$.*

Proof. For $x \in \mathscr{L}(\mathscr{A})$, define $f(x) := \#x = q^{\dim(x)}$ and $g(x) := \#(x \setminus \bigcup_{y>x} y)$. Then $f(x) = \sum_{y \geq x} g(y)$, so by Möbius inversion $g(x) = \sum_{y \geq x} f(y)\mu(x,y)$. This equals $\sum_{y \geq x} q^{\dim(y)}\mu(x,y)$. Setting $x = \hat{0}$, we have $\sum_{x \in \mathscr{L}(\mathscr{A})} q^{\dim(y)}\mu(\hat{0},y) = p_{\mathscr{A}}(q)$. □

Now suppose we have an integral arrangement $\mathscr{A}$ in $\mathbb{A}^n(\mathbb{R})$, What is its characteristic polynomial?

Let $\mathscr{A}_p = \mathscr{A} \mod p$ for a prime p, so $\mathscr{A}_p$ is an arrangement in $\mathbb{A}^n(\mathbb{F}_p)$. If $\mathscr{L}(\mathscr{A}_p) \cong \mathscr{L}(\mathscr{A})$, then $p_{\mathscr{A}_p}(\lambda) = p_{\mathscr{A}}(\lambda)$. Now take a prime power, $q = p^e$; we can think of $\mathscr{A}_p$ as generating an arrangement $\mathscr{A}_q$ (with the same defining equations) in $\mathbb{A}^n(\mathbb{F}_q)$ and we are assured that $\mathscr{L}(\mathscr{A}_p) \cong \mathscr{L}(\mathscr{A}_q)$. There are infinitely many q's for each p, so we could try to calculate $\#(\mathbb{A}^n(\mathbb{F}_{p^e}) \setminus \bigcup \mathscr{A}_{p^e})$ for all $e \geq 1$, therefore getting a formula for $p_{\mathscr{A}}(\lambda)$. (Tip: This is not what people do. But they could.)

The crucial requirement is that $\mathscr{L}(\mathscr{A}_p) \cong \mathscr{L}(\mathscr{A})$. So, when is it true? Let's be more precise about the arrangement. Say $\mathscr{A} = \{h_{\alpha_i, c_i} : i = 1, \dots, l\}$, with l hyperplanes of the form $h_{\alpha,c} = \{x \in \mathbb{A}^n(\mathbb{R}) : \alpha \cdot x = c\}$.

Since dependence of the hyperplanes corresponds to dependence of the defining equations, look at the matrix $U = \begin{pmatrix} \alpha_1 & \alpha_2 & \dots & \alpha_l \end{pmatrix} \in M_{n \times l}(\mathbb{Z})$ and the $1 \times l$ matrix $c = \begin{pmatrix} c_1 & c_2 & \dots & c_l \end{pmatrix}$. All the hyperplane equations are represented by the matrix equation $U^T x = c^T$. The solution set of this system of equations is $\bigcap \mathscr{A}$. Now projectivize to $\mathscr{A}_{\mathbb{P}}$ and let $U' := \begin{pmatrix} U \\ c \end{pmatrix}$. For any subarrangement $\mathscr{S} \subseteq \mathscr{A}_{\mathbb{P}}$, the rank is equal to the largest order of a nonsingular square submatrix of $U'_{\mathscr{S}}$, where the subscript means only taking the columns corresponding to hyperplanes in $\mathscr{S}$. If every such submatrix remains nonsingular modulo p, then every subset of columns in U' has the same rank in $U'_{\mathscr{S}}$, and that implies $\mathscr{L}(\mathscr{A}_p) \cong \mathscr{L}(\mathscr{A})$. A sufficient condition for preserving nonsingularity is that p does not divide the determinant of any nonsingular square submatrix of U'. It follows that almost all primes, and all sufficiently large primes, give the desired lattice isomorphism. That proves:

Theorem 3 (Finite Field Method)**.** *Given an integral arrangement $\mathscr{A}$ in $\mathbb{A}^n(\mathbb{R})$, for every sufficiently large prime p the modular arrangement $\mathscr{A}_p$ has the same characteristic polynomial as does $\mathscr{A}$.*

Thus, the Critical Theorem enables us to obtain $p_{\mathscr{A}}(\lambda)$ by computing the number of points of $\mathbb{A}^n(\mathbb{F}_p) \setminus \bigcup \mathscr{A}_p$ for all large primes p. This is how the finite field method works. Note that we do not need finite fields, only prime fields. In other words, we may work modulo prime numbers. In fact (but this is not part of the finite field method), we could work modulo any positive integer m that is relatively prime to all the nonzero subdeterminants of U' (no one does this).

0.1. **Affinographic arrangements.** For affinographic arrangements, where every equation has the form $x_j - x_i = c$, the finite-field method is simpler because the matrix U, the top part of U', is totally unimodular (every subdeterminant is 0 or ± 1) so all primes are good as far as concerns determinants in U. The other subdeterminants of U' are those that use c. The only nonzero ones we need to worry about are those associated with a circle C. The circle has l vertices and edges. The $l \times l$ submatrix U_C of U corresponding to those vertices and edges has determinant 0, so it is not of concern, but what is of concern is the $l \times l$ submatrix of U' obtained by substituting for one row (any one row) of U_C the row c_C of c. Call this matrix U'_C; then $\det U'_C = \pm\varphi(C)$. Therefore, in the finite field method we can use any prime that does not divide the gain of an unbalanced circle.

Exercise 4. Prove the preceding paragraph. In particular, prove the determinant formula.

This works. But it is simpler to use modular coloring. For one thing, we need not be restricted to primes. For a second, we know exactly which moduli are valid: every positive integer m that is not a divisor of any nonzero circle gain. Modular coloring does not appear to count points an affine space, unlike the critical Theorem, but in fact it is not so different. Suppose we have an integral gain graph Φ and consider a coloration $\gamma : V \to \mathbb{Z}_m$. We can view γ as the vector $(\gamma(v_1), \ldots, \gamma(v_n))$ in $\mathbb{Z}_m^V = \mathbb{Z}_m^n$. The rule for γ to be a proper coloration is that it avoids all the hyperplanes of $\mathscr{A}_m[\Phi]$. In other words, the difference is not that great. However, it is not that little, since the Critical Theorem is false in $\mathbb{Z}_m^n$ if m is composite. It is only the special form of affinographic hyperplanes that lets us use vectors (which we call colorations) in $\mathbb{Z}_m^n$ to get the characteristic polynomial.

Example: The Shi arrangement. This is the arrangement $\mathscr{S}_n = \mathscr{A}[\{0,1\}\vec{K}_n]$ associated with the Shi gain graph, $\{0,1\}\vec{K}_n$. The computation via modular coloring is simple. The 0-edges ensure that no two vertices have the same color, so as with the Catalan arrangement we can put the n vertices into spaces between $m-n$ markers to make a sequence of m places labelled by the colors $0, 1, \ldots, m-1 \in \mathbb{Z}_m$. The rule for the Shi arrangement is that two vertices may have adjacent colors but if they do, say $\gamma(v_i) = \gamma(v_j) \pm 1$ where $i < j$, then $\gamma(v_j) \neq \gamma(v_i) + 1$ due to the 1-edges. That means that if we have a (cyclically) consecutive sequence of colors applied to a bunch of vertices, those vertices must be in decreasing order by subscript. And that means that the order of vertices in a bunch that have consecutive colors is determined. So, all we need to do is place v_1 in the last place of our sequence (position $m-1$) and distribute the other $n-1$ vertices into the $m-n$ spaces arbitrarily (there are $(m-n)^{n-1}$ ways to do that. Then we rotate the sequence so v_1 is in any position (m ways). That gives each proper $\mathbb{Z}_m$ coloration exactly once, so we have the characteristic polynomial.

Proposition 5. *For the Shi arrangement $\mathscr{S}_n$ with $n \geq 1$:*

(1) *The characteristic polynomial is*

$$p_{\mathscr{S}_n}(\lambda) = \lambda(\lambda - n)^{n-1}.$$

(2) *The number of regions is*

$$(n+1)^{n-1}.$$

(3) *The number of bounded regions is*

$$(n-1)^{n-1}.$$

The number of regions equals the number of labelled trees of order $n+1$; this suggests finding an explicit bijection, which has been done (cf. Stanley).