LTCC Enumerative Combinatorics

Notes 6 | Alex Fink | Fall 2015

6 Deletion-contraction

The common thread that runs through this section is one of my favourite recurrence relations, which solves a surprising number of counting problems about graphs and related objects.

6.1 Hyperplane arrangements

Let $\mathbf{k}$ be a field. A *hyperplane* H in a $\mathbf{k}$-vector space V is a subset of V of the form $ax = b$, where $a \in \mathbf{k}^*$ is a nonzero linear form on V, and $b \in \mathbf{k}$. It is a *central* hyperplane if it is a linear subspace, i.e. $b = 0$. A *hyperplane arrangement* in V is a multiset of hyperplanes, which is *central* if all its hyperplanes are. The space V is called the *ambient space* of the arrangement. More generally, the ambient space need not be a vector space; it is sufficient that it be an *affine* space, like a vector space but "without a distinguished origin". The intersection of any number of hyperplanes in a vector space is such an affine space. (Formally we could take the ambient space to be a *torsor* for some vector space W, that is a space on which W acts freely and transitively by addition.)

If $\mathbf{k} = \mathbb{R}$, then each hyperplane H partitions V into the two connected components of its complement. Similarly, given a hyperplane arrangement $\mathscr{H}$, the complement $V \setminus \bigcup_{H\in\mathscr{H}} H$ falls into connected components; these are called the *regions* of $\mathscr{H}$.

Given a hyperplane arrangement $\mathscr{H}$ and one of its hyperplanes H_i, there are two fundamental ways to produce a smaller arrangement based on eliminating H_i. One is simply to remove it from the multiset: the result is the *deletion* $\mathscr{H} \setminus H_i$, whose ambient space is also V. The other is the *contraction*

$$\mathscr{H}/H_i = \{H \cap H_i : H \in \mathscr{H},\ H \text{ not parallel to } H_i\},$$

in the ambient space H_i. Note that, as defined here, while any deletion of $\mathscr{H}$ has one fewer hyperplane than $\mathscr{H}$, a contraction may have any number fewer. Deletions and contractions of a central arrangement are central.

A *face* of a hyperplane arrangement $\mathscr{H}$ over $\mathscr{R}$ is a region of some repeated contraction of $\mathscr{H}$. To say this another way, if the hyperplanes of $\mathscr{H}$ have equations $a_i x = b_i$, then a face of $\mathscr{H}$ is a nonempty set of form

$$\{x \in V : a_i x \gtreqless_i b_i\}$$

where each $\gtreqless_i$ is either $<$, $=$, or $>$. The faces of $\mathscr{H}$ partition its ambient space. We get the inequalities describing the regions with the same procedure but disallowing the relation $=$.

(1) One of the main examples of a central hyperplane arrangement is the *braid arrangement* $\mathscr{A}_{n-1}$ in $\mathbb{R}^n$,

$$\mathscr{A}_{n-1} = \{\{x_i = x_j\} : i \neq j \in [n]\}.$$

To specify a region of A_{n-1}, we need to indicate for each distinct i and j in $[n]$ whether x_i or x_j is greater. So the regions are in bijection with total orders on $[n]$. To specify a face, we allow x_i and x_j to be declared equal as well; so the faces of A_{n-1} are in bijection with ordered set partitions.

(2) Successive deletions of $\mathscr{A}_{n-1}$ yield *graphical arrangements*. For any loopless graph G on vertex set $[n]$, the corresponding graphical arrangement is

$$\{\{x_i = x_j\} : \{i,j\} \text{ an edge of } G\}.$$

If G has parallel edges, we also get a graphical arrangement by repeating hyperplanes appropriately.

Now, to specify a region we must indicate whether x_i or x_j is greater for each edge $\{i,j\}$ of G. This can be encoded as an orientation of G, say by orienting each edge towards the greater vertex. To ensure that the order relation described by the orientation is transitive is equivalent to insisting that there are no directed cycles. These orientations of G are called *acyclic*. Thus acyclic orientations of G are in bijection with regions of the graphical arrangement of G.

Let us write $r(\mathscr{H})$ for the number of regions of $\mathscr{H}$, and $b(\mathscr{H})$ for the number of these which are bounded. (The above examples being central, they had no bounded regions.) Both of these quantities are computed by a *deletion-contraction recurrence*.

Lemma 6.1 *For any $H \in \mathscr{H}$,*

$$r(\mathscr{H}) = r(\mathscr{H} \setminus H) + r(\mathscr{H}/H)$$

and

$$b(\mathscr{H}) = \begin{cases} 0 & \text{if } \bigcap_{J \in \mathscr{H} \setminus H} J \text{ contains a line that } H \text{ doesn't} \\ b(\mathscr{H} \setminus H) + b(\mathscr{H}/H) & \text{otherwise.} \end{cases}$$

Proof Inserting H into $\mathscr{H} \setminus H$ cuts some of the regions into two, and so

$$r(\mathscr{H}) = r(\mathscr{H} \setminus H) + \#\{\text{regions cut in two}\}.$$

But each region that was cut in two intersects H in one of the regions of $\mathscr{H} \setminus H$, so the latter summand is $r(\mathscr{H}/H)$. For b the argument is similar, except that when H is the only hyperplane of $\mathscr{H}$ not containing some line, then all regions of $\mathscr{H} \setminus H$ are unions of translates of that line, so none are bounded, nor do any become bounded when H is reinserted.

For example, an arrangement of n hyperplanes is in *general position* in $\mathbb{R}^r$ if no $(r-k)$-dimensional affine subspace is contained in more than k of them. Deletions and contractions of arrangements in general position are also in general position. So if G_n^r is an arrangement of n hyperplanes in $\mathbb{R}^r$, the above lemma yields recursions

$$\begin{aligned} r(G_n^r) &= r(G_{n-1}^r) + r(G_{n-1}^{r-1}) \\ b(G_n^r) &= b(G_{n-1}^r) + b(G_{n-1}^{r-1}), \end{aligned}$$

where the genericity assures that the lower subscript in the contraction is $n-1$ in every case. The differences lies in the base cases: $r(G_0^r) = 1$ while $b(G_0^r) = 0$, and $r(G_n^0) = b(G_n^0) = 1$ for n positive. These recurrences easily solve to

$$\begin{aligned} r(G_n^r) &= \binom{n}{r} + \binom{n}{r-1} + \cdots + \binom{n}{0} \\ b(G_n^r) &= \binom{n-1}{r}. \end{aligned}$$

6.1.1 The characteristic polynomial

The *intersection poset* $L(\mathscr{H})$ of $\mathscr{H}$ is the poset whose elements are the intersections of subsets of $\mathscr{H}$, known as *flats*, and whose partial order is reverse containment, i.e. $F \leq G$ iff $F \supseteq G$.

A *meet semilattice* is a poset in which every pair of elements has a meet (it is "semi-" because half of the two symmetric conditions for a poset to be a lattice are satisfied[1]). The intersection poset is a meet semilattice, with $F \wedge G = F \cap G$. Any finite meet semilattice with a maximum element $\hat{1}$ is a lattice, for the least upper bound of any two elements is the greatest lower bound of the set of all their common upper bounds, and this set is nonempty, containing as it does the element $\hat{1}$. As such, the intersection poset of a central hyperplane arrangement is a lattice.

[1] If only the term "semigroup" were so sensible.

Let X be a locally finite poset. Then X is *graded* if it can be endowed with a rank function $\operatorname{rank} : X \to \mathbb{Z}$ such that $\operatorname{rank}(y) = \operatorname{rank}(x) + 1$ for every covering relation $x \lessdot y$. When X is finite, we will demand that the minimum value attained by the rank function is zero; this is merely a normalisation and does not affect which posets are graded. Every intersection poset is graded, with the rank function $\operatorname{rank} F = \operatorname{codim} F$.

The *characteristic polynomial* of a hyperplane arrangement $\mathscr{H}$, denoted $\chi(\mathscr{H};q)$, is the generating function for the Möbius function of $L(\mathscr{H})$, by opposite rank:

$$\chi(\mathscr{H};q) = \sum_{F \in \mathscr{H}} \mu(\hat{0},F) q^{\operatorname{rank}(\mathscr{H}) - \operatorname{rank}(F)},$$

where $\operatorname{rank}(\mathscr{H})$ denotes the maximal rank attained on $\mathscr{H}$. For example, the general-position arrangement G^n_n has the Boolean lattice $\mathscr{B}_n$ as its intersection lattice, in which the Möbius function is $\mu(\hat{0},F) = (-1)^{\operatorname{rank}(F)}$, so $\chi(G^n_n;q) = (q-1)^n$. Indeed, if G^r_n is a generic arrangement for any $n \geq r$, all intervals in $L(G^r_n)$ are Boolean lattices, and similar argumentation gives

$$\chi(G^r_n;q) = q^r - \binom{n}{1} q^{r-1} + \cdots + (-1)^r \binom{n}{r} q^0.$$

The Crosscut Theorem, Theorem 5.13, applied to the multiset of individual hyperplanes, gives another interpretation of the characteristic polynomial:

Proposition 6.2 *Let $\mathscr{H}$ be a hyperplane arrangement. Then*

$$\chi(\mathscr{H};q) = \sum_{S \subseteq \mathscr{H}} (-1)^{|S|} q^{\dim(\bigcap S) - \dim(\ell)},$$

where $\bigcap S$ is short for $\bigcap_{s \in S} s$, and ℓ is the largest linear subspace a translate of which is contained in every $H \in \mathscr{H}$.

Note that the exponent of q equals $\operatorname{rank}(\mathscr{H}) - \operatorname{rank}(\bigvee S)$. The above formula yields a deletion-contraction recurrence.

Theorem 6.3 *Let $\mathscr{H}$ be a hyperplane arrangement and $H \in \mathscr{H}$. Then*

$$\chi(\mathscr{H};q) = \varepsilon \chi(\mathscr{H} \setminus H;q) - \chi(\mathscr{H}/H;q),$$

where ε equals 1 unless $\bigcap_{J \in \mathscr{H} \setminus H} J$ contains a line that H doesn't, in which case it equals q.

In the latter case, $\chi(\mathscr{H} \setminus H;q) = \chi(\mathscr{H}/H;q)$ so $\chi(\mathscr{H};q) = (q-1)\varepsilon\chi(\mathscr{H} \setminus H;q)$.

Proof Using our crosscut expansion,

$$\begin{aligned}\chi(\mathscr{H};q) &= \sum_{S\subseteq\mathscr{H}:H\notin S}(-1)^{|S|}q^{\dim(\bigcap S)-\dim(\ell)} + \sum_{S\subseteq\mathscr{H}:H\in S}(-1)^{|S|}q^{\dim(\bigcap S)-\dim(\ell)} \\ &= \sum_{S\subseteq\mathscr{H}\setminus H}(-1)^{|S|}q^{\dim(\bigcap S)-\dim(\ell)} + \sum_{T\subseteq\mathscr{H}\setminus H}(-1)^{|T|+1}q^{\dim(\bigcap T\cap H)-\dim(\ell)}.\end{aligned}$$

The first term is $\chi(\mathscr{H}\setminus H;q)$, unless the hyperplanes of $\mathscr{H}\setminus H$ share a larger linear space than H, in which case it contains an extra factor of q; this accounts for the ε. The second is seen to be $-\chi(\mathscr{H}/H;q)$, by rewriting $(\bigcap_{t\in T}t)\cap H$ as $\bigcap_{t\in T}(t\cap H)$, and noting that ℓ plays the same role in $\mathscr{H}/H$ as in $\mathscr{H}$.

By comparing the base cases and noting that the recursions satisfied are the same, we conclude the following result:

Proposition 6.4 (Zaslavsky) *Let $\mathscr{H}$ be a hyperplane arrangement over $\mathbb{R}$. Then $r(\mathscr{H}) = (-1)^{\operatorname{rank}\mathscr{H}}\chi(\mathscr{H};-1)$ and $b(\mathscr{H}) = (-1)^{\operatorname{rank}\mathscr{H}}\chi(\mathscr{H};1)$.*

It is also easy to see the following by induction on the deletion-contraction recurrence, noting that deleting a hyperplane preserves the rank except for the $\varepsilon = q$ case, and contracting always decreases it by one.

Corollary 6.5 *Let $\mathscr{H}$ be a hyperplane arrangement. Then $\chi(\mathscr{H};q)$ is a polynomial in q with integer coefficients. This polynomial is of degree* $\operatorname{rank}(\mathscr{H})$*, and its coefficients alternate in sign, with the leading coefficient equalling 1.*

Finally, we give another way to compute characteristic polynomials of arrangements.

Proposition 6.6 (Athanasiadis) *Let $\mathscr{H}$ be a hyperplane arrangement over the field of order q in the ambient space V, and let ℓ be the largest linear subspace a translate of which is contained in every hyperplane of $\mathscr{H}$. Then $q^{\dim(\ell)}\chi(\mathscr{H};q)$ is the number of points of V not on any hyperplane of $\mathscr{H}$.*

Proof Apply the inclusion-exclusion principle to count the points not on any hyperplane, with the sets A_i being the sets of points on each hyperplane. The result exactly matches the right hand side of Proposition 6.2, up to the factor of $q^{\dim(\ell)}$.

As an aside, Zaslavsky's result for $r(\mathscr{H})$ can be proved in the same way with compactly supported Euler characteristic in place of cardinality.

It may seem that Proposition 6.6 only gives us one evaluation of the characteristic polynomial and not the whole polynomial, but its utility is greater than it appears. By extending scalars from the field of order q to the field of order q^n

for positive integers n, we can use the proposition to evaluate the characteristic polynomial at countably many values, which is enough to recover it. Moreover, if $\mathscr{H}$ is a hyperplane arrangement over $\mathbb{Q}$, then for all but finitely many primes p, reducing the defining equations $a_i x = b_i$ of its hyperplanes modulo p gives a well-defined hyperplane arrangement over the field of order p in which none of the dimensions in Proposition 6.2 have changed. Indeed, this reduction will only fail if p divides a denominator of one of the a_{ij} or b_i or divides the determinant of some square submatrix whose rows are chosen from the a_i.

For example, let us compute the characteristic polynomial of the braid arrangement $\mathscr{A}_{n-1}$. If p is a prime, there are no problems reducing the defining equations $x_i = x_j$ of the hyperplanes mod p. So what we wish to count are the points in $\mathbb{F}_p^n$ all of whose coordinates are distinct: there are clearly $(p)_n$ of these. The linear space ℓ is one-dimensional, consisting of all the points with all coordinates equal. We conclude

$$\chi(\mathscr{A}_{n-1};q) = \frac{(q)_n}{q^1} = (q-1)_{n-1}.$$

6.2 Graphs

In this subsection we will allow our graphs to have *loops*, edges both of whose endpoints are at the same vertex, and multiple *parallel edges* which share the same endpoints. Formally, we could say a graph G is the data of a set $V(G)$ of vertices and a multiset $E(G)$ of edges, where an edge is a multiset on $V(G)$ of size 2.

6.2.1 Colourings

No more need be said to motivate the problem of graph colouring than "the four-colour theorem". A colouring of a graph by a set S of colours is a set function from its vertices to S such that no edge has the same colour assigned to both its endpoints. As usual, we may as well take the set of colours to be a standard set $[q]$ for some natural number q; a q-colouring will be a colourings by the set $[q]$.

Let us count graph colourings. Define $\chi(G;q)$ to be the number of q-colourings of G. This quantity is known as the *chromatic polynomial* of G, for as we will see shortly it turns out to be a polynomial in q.

If G is a graph and $e \in E(G)$ one of its edges, the *deletion* of e in G is the new graph $G \setminus e$ obtained in the obvious way by subtracting the edge e from the edge set. We also define the *contraction* of e in G. If v and w are the vertices of e, then the contraction is the graph G/e obtained by discarding e and "merging" v and w into a new vertex $\bar{v}$: that is, the vertex set of G/e is $(V(G) \setminus \{v,w\}) \cup \{\bar{v}\}$, and its edge set is obtained from $E(G)\setminus e$ by changing each edge endpoint that was v or w to $\bar{v}$. Note that contracting a loop has the same effect as simply deleting it.

The chromatic polynomial has a deletion-contraction recurrence:

Proposition 6.7 *Let G be a graph, and e an edge of G. Then*

$$\chi(G;q) = \begin{cases} 0 & \textit{if e is a loop} \\ \chi(G\setminus e;q) - \chi(G/e;q) & \textit{otherwise.} \end{cases}$$

Proof Clearly a graph with a loop has no colourings. Otherwise, let v and w be the endpoints of e. The colourings of $G\setminus e$ are of two sorts: those assigning different colours to v and w, which are in bijection with the colourings of G; and those assigning the same colour to v and w, which are in bijection with the colourings of G/e.

The base cases needed to use this recurrence are the graphs with no edges. Clearly, if G is the graph with n vertices and no edges, then $\chi(G;q) = q^n$. Thus, inductively, $\chi(G;q)$ is indeed a polynomial.

In fact, it's a polynomial we've seen before:

Proposition 6.8 *Let G be a graph without loops. Then the chromatic polynomial of G is the characteristic polynomial of the graphical arrangement of G.*

Proof If $\mathbb{F}_q$ is a finite field of order q, then the colourings of G by $\mathbb{F}_q$ are exactly the points lying off the graphical arrangement counted by Proposition 6.6. So the two polynomials agree at each prime power q, and must therefore be equal.

Zaslavsky's result and our above interpretation of the regions of the graphical arrangement yield another nice example of combinatorial reciprocity, between colourings of a graph, counted by $\chi(G,q)$, and acyclic orientations of that graph, counted by $(-1)^{|V(G)|-b_0(G)}\chi(G,-1)$, where $b_0(G)$ is the number of connected components of G. We note without further discussion that Stanley has extended this to a full reciprocity, giving a meaning to $(-1)^{|V(G)|-b_0(G)}\chi(G,-q)$ for all naturals q: these count certain "compatible" pairs of acyclic orientations and q-colourings.

6.2.2 Spanning trees

A *spanning tree* of a graph is a subset of its edges which make up a tree. Clearly a spanning tree can contain no loops, and at most one of any set of parallel edges. Let us write $b(G)$ for the number of spanning trees of G. This also has a deletion-contraction recurrence.

Proposition 6.9 *Let G be a graph, and e an edge of G. Then*

$$b(G) = \begin{cases} b(G \setminus e) & \text{if } e \text{ is a loop} \\ b(G \setminus e) + b(G/e) & \text{otherwise.} \end{cases}$$

Proof The spanning trees of G which don't use e are in bijection with the spanning trees of $G \setminus e$. There are no spanning trees which do use e if e is a loop; if it isn't, these are in bijection with the spanning trees of G/e.

The base cases for this recurrence can again be taken to be the graphs G with no edges; these have one spanning tree if $|V(G)| = 1$ and none if $|V(G)| > 1$. However, a disconnected graph has no spanning trees, so the moment the graph becomes disconnected by removal of an edge (we call such an edge an *isthmus* or *bridge* or *coloop*) we may as well shortcut the recurrence and take $b(G) = 0$.

There is however a higher-tech and much faster way to count spanning trees. The *Laplacian matrix* of a graph G is the matrix $L(G)$, its rows and columns indexed by $V(G)$, with

$$L(G)_{vw} = \begin{cases} -\#\{\text{edges connecting } v \text{ and } w\} & v \neq w \\ \#\{\text{non-loop edges incident to } v\} & v = w. \end{cases}$$

Note that $L(G)$ is symmetric, and each of its rows sum to zero, so it is singular. However:

Theorem 6.10 (Kirchhoff's matrix-tree theorem) *Let G be a connected graph. Then $b(G)$ equals the determinant of any principal cofactor of $L(G)$.*

Recall that a principal cofactor of a matrix is obtained by deleting its vth row and vth column, for some v. Said otherwise, if the eigenvalues of $L(G)$ with multiplicity are $\lambda_1, \ldots, \lambda_{n-1}, \lambda_n = 0$, then $b(G) = \lambda_1 \cdots \lambda_{n-1}/n$.

The proof of Kirchhoff's matrix-tree theorem uses the *signed incidence matrix* $M(G)$ of G. (For readability we will suppress the argument "(G)" of our notations from here on.) In fact "the" is not the correct article, since M depends on some choices: namely, we must give each edge of G an orientation, distinguishing its two endpoints as a head and a tail. Then M is the V-by-E matrix given by

$$M_{ve} = \begin{cases} 1 & v \text{ is the head but not the tail of } e \\ -1 & v \text{ is the tail but not the head of } e \\ 0 & \text{otherwise.} \end{cases}$$

Observe that $L = MM^{\mathrm{T}}$.

Given a matrix A and subsets S, T of its sets of row and column indices, let us write A_{ST} for the submatrix consisting of rows S and columns T. Write $W = V \setminus \{v\}$. Then the vth principal cofactor of L is $L_{WW} = M_{WE} \cdot M_{WE}^{\mathrm{T}}$.

Lemma 6.11 *Let $S \subseteq E$ be a set of edges of size $|V|-1$. Then*

$$\det M_{WS} = \begin{cases} \pm 1 & \text{if } S \text{ is a spanning tree of } G \\ 0 & \text{otherwise.} \end{cases}$$

Proof If S is a spanning tree, let e be one of its edges incident to v. There is only one nonzero entry in the eth column of M_{WS}, which is ± 1, and cofactor expansion along this column reduces the computation to the determinant of a block matrix, each of whose blocks has the same form as M_{WS} for one of the connected components of $S \setminus \{e\}$. This proves the first case by induction.

If S is not a spanning tree, then it contains a cycle C. We may assume v is not a vertex of C. If it is, then let w be a vertex not in C, and use row operations to replace the wth row of M_{WS} with the vth row of M_{VS}, which is possible since the rows of M sum to the zero vector. This amounts to changing W to $V \setminus \{v\}$, which as we see preserves the determinant (up to sign). Now M_{WS} is a block-diagonal matrix with one of the blocks being the signed incidence matrix of C. This is singular (its rows sum to zero), and so M_{WS} is singular as well.

Now the proof of Kirchhoff's matrix-tree theorem is complete using the Cauchy–Binet formula:

$$\begin{aligned} \det(L_{WW}) &= \det(M_{WE} \cdot M_{WE}^{\mathrm{T}}) \\ &= \sum_{S \subseteq E \,:\, |S| = |V|-1} \det(M_{WS}) \det(M_{WS}^{\mathrm{T}}) \\ &= \sum_{S \subseteq E \,:\, |S| = |V|-1} \det(M_{WS})^2 \\ &= \sum_{S \text{ a spanning tree}} 1. \end{aligned}$$

For example, let us revisit the problem of counting labelled trees on n vertices from Section 4.1.1. Such a labelled tree is exactly a spanning tree of the complete graph K_n, whose Laplacian is

$$L(K_n) = nI_n - J_n.$$

Here I_n is, as usual, the $n \times n$ identity matrix, and J_n is the $n \times n$ matrix with all entries 1. After deleting one row and a matching column, we can easily find the determinant by row operations on the resulting $(n-1) \times (n-1)$ matrices:

$$b(K_n) = \det \begin{bmatrix} n-1 & -1 & \cdots & -1 \\ -1 & n-1 & \cdots & -1 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & \cdots & n-1 \end{bmatrix}$$

$$= = \det \begin{bmatrix} 1 & 1 & \cdots & 1 \\ -1 & n-1 & \cdots & -1 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & \cdots & n-1 \end{bmatrix}$$

$$= = \det \begin{bmatrix} 1 & 1 & \cdots & 1 \\ 0 & n & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & n \end{bmatrix} = n^{n-2}.$$

Exercise Show that the complete bipartite graph $K_{m,n}$ has $m^{n-1}n^{m-1}$ spanning trees. Can you find a bijective proof?

6.3 Matroids

The fullest setting in which the deletion-contraction recurrence is at home is that of *matroids*. I will not explore them in depth here, but couldn't bear not to mention them.

Matroids are a combinatorial object generalising hyperplane arrangements, and therefore graphs. They were introduced in the mid-1930s by Hassler Whitney (as is well known) and Takeo Nakasawa (as is not). Since then they have spent time in and out of fashion; at present they are a hot topic, with the success of the programme of Geelen, Gerards and Whittle to prove Rota's excluded minors conjecture through intense structural study, as well as the exploitation of connections to algebro-geometric combinatorics.

A curiosity of matroid theory is the large number, easily dozens, of different-looking but equivalent definitions of a matroid. Rota gave this phenomenon the name "cryptomorphism". We give two definitions, the first closer to our development above and the second perhaps more usual.

Definition 6.12 A matroid M on the finite ground set E is a graded lattice X, whose *atoms* i.e. elements of rank 1 are labelled by a set partition with nonempty parts of a subset of E, satisfying:

- *atomicity*: every element of X is the join of some set of atoms (regarding $\hat{0}$ as the join of the empty set);
- *submodularity*: for every $x, y \in X$,

$$\mathrm{rank}(x) + \mathrm{rank}(y) \geq \mathrm{rank}(x \wedge y) + \mathrm{rank}(x \vee y).$$

So the intersection lattice of any central hyperplane arrangement $\mathscr{H}$ is a matroid $M(\mathscr{H})$ on its multiset of hyperplanes (or, more carefully, on a set indexing its multiset of hyperplanes). The labelling by a set partition lets us annotate the intersection lattice to say when several of the hyperplanes are equal. By allowing the set partition to be of only a subset of the ground set, we also allow for degenerate "hyperplanes" of rank 0, i.e. containing the whole ambient space, which is something we had ruled out in the definition of a hyperplane arrangement.

To non-central hyperplane arrangements there is a way to associate a generalisation called a *semimatroid*, but we haven't space to go into these. One can also associate a matroid by replacing the arrangement with a central one, namely the *cone* over it: replace the ambient space V by $V \oplus \mathbf{k}$, and each hyperplane H by $\{(\lambda x, \lambda) : x \in H\}$.

Definition 6.13 A matroid M on the finite ground set E is a nonempty set $\mathscr{B}$ of subsets of E, called its *bases*, satisfying the *exchange axiom*: for every two bases $A, B \in \mathscr{B}$, and every element $a \in A \setminus B$, there exists an element $b \in B \setminus A$ such that $(A \setminus \{a\}) \cup \{b\} \in \mathscr{B}$.

In fact, all bases have the same cardinality, which is called the *rank* of the matroid and denoted $\operatorname{rank}(M)$.

Given a matroid realised as a graded lattice X, its rank is the rank r of the lattice, and its bases are all subsets of E of cardinality r, all of which are used to label some atom, and so that the join of the corresponding atoms is $\hat{1}$. Conversely, given a matroid realised as a set $\mathscr{B}$ of bases on ground set E, the corresponding lattice is the lattice, under inclusion, of all subsets $S \subseteq E$ such that no element may be added to S without increasing $\max_{B \in \mathscr{B}} |S \cap B|$; the value of this maximum is the rank of S in the lattice. The other piece of data, the set partition, is given by associating to each atom S the set $S \setminus \hat{0}$ (bearing in mind that S and $\hat{0}$ are both subsets of E). We leave as an exercise to the reader to check that these operations are mutually inverse and yield objects satisfying the requisite properties.

In the matroid $M(G)$ associated to a connected graph G (by passing through its graphical arrangement), the ground set is the edge set $E(G)$, and the bases are exactly the spanning trees of G.

We will state the prerequisites for our main theorem on matroids in the basis language. Let M be a matroid with bases $\mathscr{B}$ and e an element of its ground set, Unless e appears in every basis, the *deletion* $M \setminus e$ is the matroid with bases

$$\{B \in \mathscr{B} : e \notin B\}.$$

And unless e appears in no basis, the *contraction* M/e is the matroid with bases

$$\{B \setminus \{e\} : B \in \mathscr{B}, e \in B\}.$$

An element appearing in no basis is called a *loop*, and one appearing in every basis is called a *coloop*. The *rank* $\mathrm{rank}(S)$ of a subset $S \subseteq E$ is defined to be $\max_{B\in\mathscr{B}} |S \cap B|$.

The *Tutte polynomial* of a matroid M is the polynomial in $\mathbb{Z}[x,y]$ defined as

$$T(M;x,y) = \sum_{S\subseteq E} (x-1)^{\mathrm{rank}(M)-\mathrm{rank}(S)}(y-1)^{|S|-\mathrm{rank}(S)}.$$

The Tutte polynomial satisfies a deletion-contraction recurrence,

$$T(M;x,y) = T(M\setminus e;x,y) + T(M/e;x,y)$$

unless e is a loop or a coloop. The recurrence bottoms out at matroids with one basis, where every element is a loop or a coloop. The Tutte polynomial of such a matroid is a monomial $x^{|B|}y^{|E|-|B|}$; we conclude that the Tutte polynomial has nonnegative coefficients.

What's more important is that the Tutte polynomial is "universal" for the deletion-contraction recurrence: *every* invariant that satisfies it can be written as an evaluation of the Tutte polynomial.

Theorem 6.14 *Let* $\mathbf{k}$ *be a field of characteristic 0, and* f *a function associating a value in* $\mathbf{k}$ *to each matroid (variation: each graph) such that*

- $f(M) = 1$ *when the ground set of* M *is empty;*
- *if* e *is a loop of* M*, then* $f(M) = Af(M\setminus e)$*, where* $f(\text{a single loop}) = A$*;*
- *if* e *is a coloop of* M*, then* $f(M) = Bf(M/e)$*, where* $f(\text{a single coloop}) = B$*;*
- *if* e *is neither a loop nor a coloop of* M*, then* $f(M) = \alpha f(M\setminus e) + \beta f(M/e)$ *where* α *and* β *are nonzero constants in* $\mathbf{k}$*.*

Then

$$f(M) = \alpha^{|E|-\mathrm{rank}(M)}\beta^{\mathrm{rank}(M)}T(M;\frac{A}{\beta},\frac{B}{\alpha}).$$

We close with a small list of evaluations of the Tutte polynomial of a matroid M, and their interpretations when $M = M(G)$ for a graph G, without proof.

- $T(M;1,1)$ is the number of bases of M, that is the number of spanning trees of G if G is connected; more generally, it is the number of choices of a spanning tree in each component of G.
- $T(M;2,1)$ is the number of subsets of bases of M, that is the number of subforests of G.

- $T(M;1,2)$ is the number of supersets of bases of M, that is the number of connected subgraphs of G if G is connected, where a subgraph must contain all the vertices of G; more generally, it is the number of subgraphs of G with no more connected components than G.
- $T(M;2,2) = 2^{|E|}$ is vacuous, as is $T(M;0,0) = 0$.
- $T(M(G);1-q,0) = (-1)^{\mathrm{rank}(M(G))} q^{-b_0(G)} \chi(G;q)$ is the chromatic polynomial.
- In particular, $T(M(G);2,0) = (-1)^{|V(G)|} \chi(G;-1)$ is the number of acyclic orientations of G.
- $T(M(G);0,2)$ is the number of orientations of G that result in a *strongly connected* directed graph, i.e. one in which there is a directed path from any vertex to any other.
- For any abelian group $(A,+)$, an *flow* on G valued in A is an element of the right kernel of the signed adjacency matrix $M(G)$, interpreted as a matrix over A: that is, it is an assignment of an element of A to each edge of G such that, at each vertex, the sum on incoming edges equals the sum on outgoing edges. Then, if $|A| = q$ is finite, $T(M(G);0,1-q)$ is the number of flows on G valued in A not assigning zero to any edge.