LTCC Enumerative Combinatorics

Notes 5 Alex Fink Fall 2015

5 Posets and Möbius inversion

Möbius inversion can be viewed as a generalisation of the inclusion-exclusion principle with an apparatus to keep track of how the conditions intersect, as an apparatus to reduce the number of terms. The apparatus, that of partial orders, turns out to be of great combinatorial utility in its own right.

5.1 The inclusion-exclusion principle

Often we are in the situation where we have a number of conditions on a set of combinatorial objects, and we have information about the number of objects which satisfy various combinations of these conditions (inclusion), while we want to count the objects satisfying none of the conditions (exclusion), or perhaps satisfying some but not others. What is known as the *sieve method* is of general use in this situation: overcount the objects satisfying the conditions, and then make corrections and subtract off elements that have been multiply counted, and so forth. The sieve of Eratosthenes gave its name to the class (although, alone, it's not especially helpful for the enumeration of primes): the primes are the integers which satify none of the conditions of having the forms $2n, 3n, 5n, 7n, \dots$ for $n \geq 2$.

Let $A_1, \ldots, A_n$ be subsets of a finite set X. For any non-empty subset J of the index set $[n]$, we put

$$A_J = \bigcap_{j \in J} A_j,$$

and take $A_\emptyset = X$ by convention.

Theorem 5.1 (Inclusion-Exclusion Principle) *The number of elements of X lying in none of the sets A_i is equal to*

$$\sum_{J \subseteq [n]} (-1)^{|J|} |A_J|.$$

Proof The expression in the theorem is a linear combination of the cardinalities of the sets A_J, and so we can calculate it by working out, for each $x \in X$, the contribution of x to the sum. If K is the set of all indices j for which $x \in A_j$, then x contributes to the terms involving sets $J \subseteq K$, and the contribution is

$$\sum_{J \subseteq K} (-1)^{|J|}.$$

Just as in Proposition 2.3, this is the sum of the terms encountered when expanding the product

$$\prod_{k\in K}(1-1)=\begin{cases}1 & |K|=0\\ 0 & \text{otherwise.}\end{cases}$$

So the points with $K=\emptyset$ (those lying in no set A_i) each contribute 1 to the sum, and the remaining points contribute nothing. So the theorem is proved.

Here are some examples.

(1) In section 2.4 we counted the surjective functions from $[n]$ to $[k]$, obtaining the number $k!\,S(n,k)$. We can also apply inclusion-exclusion. Let X be the set of all functions $f:[n]\to[k]$, and A_i the set of functions whose range does not include the point i. Then A_J is the set of functions whose range includes none of the points of J, that is, functions from $[n]$ to $[k]\setminus J$; so $|A_J|=(k-j)^m$ when $|J|=j$. A function is a surjection if and only if it lies in none of the sets A_i. Collapsing together all the terms with $|J|=j$ for each j, the inclusion-exclusion formula for the number of surjections is

$$\sum_j(-1)^j\binom{k}{j}(k-j)^n.$$

We thus get an alternating sum for the Stirling number of the second kind,

$$S(n,k)=\sum_{j=0}^{n}\frac{(-1)^j(k-j)^n}{j!\,(k-j)!}.$$

(2) We can count derangements similarly. Let X be the set of all permutations of $[n]$, and A_i the set of permutations fixing i. Then A_j is the set of permutations fixing every point in J; so $|A_J|=(n-j)!$ when $|J|=j$. The permutations lying in none of the sets A_i are the derangements, and so their number is

$$\sum_j(-1)^j\binom{n}{j}(n-j)!=n!\sum_{j=0}^{n}\frac{(-1)^j}{j!},$$

agreeing with example (6) of Section 4.

(3) Number theory was one of the first subjects in which Möbius inversion, to be discussed below, found application. The formula for Euler's *totient function*,

$$\phi(n)=\#\{k\in\mathbb{Z}:0\le k<n,\gcd(k,n)=1\},$$

can be seen in this light. Let $n=p_1^{a_1}\cdots p_e^{a_e}$ be the prime factorisation of n, so that $\gcd(k,n)=1$ iff no p_i divides n.

Hence take A_i to be the set of nonmultiples of p in $X = \{0, \dots, n-1\}$. The elements of A_J for $J \subseteq [e]$ are the multiples of $\prod_{j \in J} p_J$, of which there are $n / \prod_{j \in J} p_J$. Hence

$$\begin{aligned} \phi(n) &= \sum_{J \subseteq [e]} (-1)^{|J|} \frac{n}{\prod_{j \in J} p_J} \\ &= n \left(1 - \frac{1}{p_1}\right) \cdots \left(1 - \frac{1}{p_e}\right). \end{aligned}$$

To count the elements of X in exactly a given collection of the sets A_k, those with $k \in K$, we apply the inclusion-exclusion principle taking A_K for X.

Corollary 5.2 *The number of elements of X lying in exactly the sets A_k for $k \in K$ but no others is*

$$\sum_{K \subseteq J \subseteq [n]} (-1)^{|J|-|K|} |A_J|.$$

Since this can be done for each K, the corollary can be interpreted as giving a change of basis between the set of indicator functions of the A_J and the set of indicator functions of the sets of elements in exactly the sets A_k but none of the remaining ones. The components of the vectors being transformed need not be natural numbers for the linear algebra to go through:

Proposition 5.3 *Let elements a_J and b_J of an abelian group $(G, +)$ be given for each subset J of $[n]$. Then the following are equivalent:*

(a) $a_J = \sum_{J \subseteq I \subseteq [n]} b_I$ *for all* $J \subseteq [n]$;

(b) $b_J = \sum_{J \subseteq I \subseteq [n]} (-1)^{|I|-|J|} a_I$ *for all* $J \subseteq [n]$.

By taking a_J to depend only on $|J|$, we recover the change-of-basis result for binomial coefficients at the end of Section 2.3.2.

5.2 Posets

A *partial order* on a set X is a binary relation $\leq$ on X which satisfies the following conditions:

- *reflexivity*: $x \leq x$;
- *antisymmetry*: if $x \leq y$ and $y \leq x$ then $x = y$;

- *transitivity*: if $x \leq y$ and $y \leq z$ then $x \leq z$.

A set bearing a partial order is very frequently called a *poset*, for "partially ordered set".

We use other inequality symbols for relations based on $\leq$ in the obvious way: so $x \geq y$ means $y \leq x$, and $x < y$ means $x \leq y$ and $x \neq y$. Note that there may be *incomparable* pairs of elements in a poset, that is pairs x and y with neither $x \leq y$ nor $x \geq y$; as such, $x < y$ does not mean $x \ngeq y$. A *total order* is a partial order in which every pair of elements is comparable.

We say x *covers* y, and write $x \lessdot y$, if $x < y$ and there exists no $z \in X$ such that $x < z < y$. A poset X is conventionally drawn as a *Hasse diagram*, which is a graph-type picture with a vertex for each element of X and an edge going *up* from x to y for each covering relation $x \lessdot y$. The figure below is a Hasse diagram for a poset on $[7]$, under which for example $3 < 2$ but 3 is incomparable to 4, and 7 is not comparable to any other element.

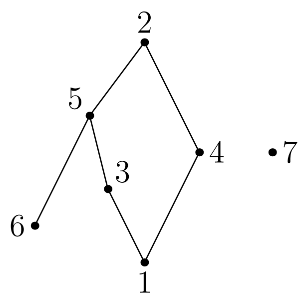

An *interval* in a poset X is a subset of form

$$[x,y] := \{z \in X : x \leq z \leq y\}$$

for $x, y \in X$. This interval is itself a poset, and one that contains a unique minimum x and a unique maximum y. We often use the name $\hat{0}$ for the unique minimum element of any poset that has one, and $\hat{1}$ for the unique maximum element likewise.

Some standard examples of posets are:

(1) For $n \in \mathbb{N}$, the poset $\mathbf{n}$ is the set $[n]$ with the usual partial order on integers. This is in fact a total order. So are $\mathbb{N}$ and $\mathbb{Z}$ with their usual order.

(2) The *Boolean lattice* $\mathscr{B}_n$ is the poset of subsets of $[n]$ ordered by containment, $S \leq T$ iff $S \subseteq T$.

(3) The *partition lattice* Π_n is the poset of set partitions of $[n]$ ordered by refinement, i.e. $\pi \leq \rho$ iff every part of the partition π is a subset of some part of the partition ρ.

An isomorphism $f : P \xrightarrow{\sim} Q$ of posets is a bijection of the underlying sets preserving the order, i.e. $x \leq y$ if and only if $f(x) \leq f(y)$.

Given two posets P and Q, their product $P \times Q$ is the poset whose underlying set is the Cartesian product $P \times Q$, with $(p,q) \leq (p',q')$ iff $p \leq p'$ and $p \leq q'$. For example, $\mathscr{B}_n$ is isomorphic to the n-fold product of $\mathbf{2}$.

(4) For n a positive integer, the poset D_n of positive divisors of n bears the order by divisibility, $a \leq b$ iff $a \mid b$. If $n = p_1^{a_1} \cdots p_e^{a_e}$ is the prime factorisation of n, then D_n is isomorphic to $\mathbf{b_1} \times \cdots \times \mathbf{b_e}$, where $b_i = a_i + 1$ (due to typographical awkwardness).

We can also define the poset of all positive integers under divisibility. We'd like to say this is isomorphic to a countable direct product of $\mathbb{N}$. The requisite definition cannot quite be made naturally at the level of posets, though, as we must demand that all but finitely many elements of the factors $\mathbb{N}$ are 0, and the element $0 \in \mathbb{N}$ has no distinguished role.

We define a few more items of terminology. The *opposite* of a poset $(X,\leq)$ is the poset $(X,\geq)$, that is the poset with the same underlying set but order relations reversed. A *chain* in a poset is a sequence of elements $x_0,\ldots,x_k$ such that

$$x_0 < \cdots < x_k;$$

we also call the set $\{x_0,\ldots,x_k\}$ a chain. A *multichain* is a sequence of elements $x_0,\ldots,x_k$ such that

$$x_0 \leq \cdots \leq x_k.$$

So a multichain can repeat elements, where a chain cannot; the name is by analogy to "set" vs. "multiset". The chain and the multichain above each have *length* k — that is, the length is the number of relations, not the number of elements.

We will deal mostly with finite posets here. The class of infinite posets to which some of our results will extend are those that are *locally finite*. A partially ordered set X is locally finite if, for any $x,y \in X$, the interval $[x,y]$ is finite. Most of our above examples of finite posets have locally finite analogues with an infinite ground set.

5.3 The incidence algebra and Möbius inversion

Let R be a ring; for the purposes of enumeration we will want $\mathbb{Z} \subseteq R$. The *incidence algebra* $I(X)$ of a finite, or more generally locally finite, partially ordered set X, over R, is defined to be the set of all functions $f : X \times X \to R$ which have the property that $f(x,y) = 0$ unless $x \leq y$. The algebra structure is the one given by regarding these functions as matrices whose (x,y)th entry is $f(x,y)$. That is,

$$(f+g)(x,y) = f(x,y) + g(x,y),$$
$$(fg)(x,y) = \sum_z f(x,z)g(z,y).$$

We will nearly always omit the ring R from the notation, as above, but will write the incidence algebra as $I_R(X)$ when we wish to foreground it.

The identity element of the incidence algebra is given by

$$1(x,y) = \begin{cases} 1 & x = y \\ 0 & \text{otherwise.} \end{cases}$$

Proposition 5.4 *An element $f \in I(X)$ is invertible if and only if each entry $I(x,x)$ is invertible (on the same side) in R.*

In particular, if invertible elements of R have two-sided inverses, then the same is true of $I(X)$. This proof is much the same as that of Proposition 3.1: the equation $fg = 1$ gives a recurrence for entries of g that can be solved on larger and larger intervals sequentially, and only diagonal entries of f ever need to be inverted.

The *zeta function* of X is the element $\zeta \in I(X)$ given by

$$\zeta(x,y) = 1$$

for all $x \leq y$ in X. So the zeta function has the largest possible support of any element of the incidence algebra. We compute

$$\zeta^2(x,y) = \sum_z \zeta(x,z)\zeta(z,y) = \sum_{x \leq z \leq y} 1 = |[x,y]|$$

and more generally

$$\zeta^n(x,y) = \sum_{x \leq z_1 \leq \cdots z_{n-1} \leq y} 1,$$

the number of multichains from x to y of length n. Similarly, since

$$(\zeta - 1)(x,y) = \begin{cases} 1 & x < y \\ 0 & \text{otherwise,} \end{cases}$$

we get that $(\zeta-1)^n(x,y)$ is the number of chains from x to y of length n, so the generating function for chains by their length is the geometric series

$$1 + (\zeta-1)t + (\zeta-1)^2t^2 + \cdots = \frac{1}{1 - t(\zeta - 1)}, \tag{1}$$

written in the indeterminate t. For instance, setting $t = 1$ shows that $1/(2-\zeta)$ is the unweighted enumerator for all chains.

The *Möbius function* of X is the inverse of the zeta function, which exists by Proposition 5.4. Due to its importance, we spell the defining recurrence out:

$$\mu(x,y) = \begin{cases} 1 & x = y \\ -\sum\limits_{x \le z < y} \mu(x,z) & x < y. \end{cases} \tag{2}$$

By taking the inverse on the other side, one could also take the sum over $x < z \le y$ of $\mu(z,y)$. Note that all entries of μ are integers. Indeed, by equation (1) with $t = -1$, we have an observation of Philip Hall:

$$\mu(x,y) = c_0 - c_1 + c_2 - c_3 + \cdots$$

where c_i is the number of chains of length i from x to y.

Proposition 5.5 (Möbius inversion) *Let X be a poset, and $f, g : X \to R$ functions. The following are equivalent:*

- $g(y) = \sum\limits_{x \le y} f(x)$ *for all* $y \in X$*;*
- $f(y) = \sum\limits_{x \le y} g(x)\mu(x,y)$ *for all* $y \in X$*;*

The proof is a direct consequence of the fact that ζ and μ are inverses. To formalise it the machinery we have developed so far, one can give the set R^X of functions $X \to R$ the structure of an R-module by the usual matrix-vector multiplication.

5.3.1 Linear extensions

A relation σ is an *extension* of a relation ρ if $x\,\rho\,y$ implies $x\,\sigma\,y$; that is, if ρ is a subset of σ, regarding a relation in the usual way as a set of ordered pairs. A *linear extension* is an extension which is a total order.

Proposition 5.6 *Every partial order has a linear extension.*

This theorem is easily proved for finite sets: take any element x which is maximal in the poset, and declare it maximum in the total order. Then delete x from the poset and recurse onto the remainder, picking from it the second-largest element in the total order; repeat in this vein until all elements have been chosen. The proof for infinite sets is harder. The Zermelo-Fraenkel axioms of set theory do not suffice; an additional principle such as the axiom of choice is required.

Note that, if f is an element of the incidence algebra $I(X)$, then its matrix is lower triangular if the ordering of the rows and columns is given by a linear extension of X. So the incidence algebra of any finite poset is isomorphic to a subalgebra of the algebra of upper-triangular matrices. (However, we do not need to give a special role to any linear extension of X to develop the theory, just as the matrix algebra can be defined without a distinguished total order on the row and column positions).

We may wish to count the linear extensions of a finite poset X. Let $e(X)$ be the number thereof. For example, $e(\mathbf{2}\times\mathbf{n})$ is the Catalan number $C_n = \binom{2n}{n}/(n+1)$. If $\mathbf{m}+\mathbf{n}$ is the disjoint union of the posets $\mathbf{m}$ and $\mathbf{n}$ with no comparability between elements of the two, then $e(\mathbf{m}+\mathbf{n}) = \binom{m+n}{m}$.

A recurrence for $e(X)$ is easy to extract from the proof of the last proposition.

Proposition 5.7 *Let $x_1,\dots,x_k$ be the maximal elements of a finite poset X. Then*

$$e(X) = \sum_{i=1}^{k} e(X \setminus \{x_i\}).$$

As a base case, the empty poset $\emptyset$ has $e(\emptyset) = 1$.

There is another technique necessitating a bit more setup. An *order ideal*, or *downset*, of a poset X is a subset of X "closed under going down", i.e. such that if it contains some $x \in X$ then it also contains any element less than or equal to x. The opposite notion, a subset "closed under going up", is called an *order filter* or *upset* of X.

If X is a poset, the set of all downsets of X is itself a poset under containment, denoted $J(X)$. In fact, $J(X)$ is a *lattice*. A poset is a lattice if, given any two of its elements x and y, there is a unique least upper bound for both, and a unique greatest lower bound. The least upper bound is denoted $x \vee y$ and called the *join* of x and y; the greatest lower bound is denoted $x \wedge y$ and called their *meet*. Note that every finite lattice has a minimum element $\hat{0}$, namely the meet of all its elements, and a maximum element $\hat{1}$, the join of all its elements. As other examples, any total order is trivially a lattice, with $x \vee y = \max\{x,y\}$ and $x \wedge y = \min\{x,y\}$. The Boolean lattice and the partition lattice are also lattices, as their names suggest.

Proposition 5.8 *Let X be a finite poset. Linear extensions of X are in bijection with chains of maximal length in $J(X)$.*

The bijection sends a linear extension $\preceq$ of X to the chain

$$\{\emptyset\} \cup \{\{y \in X : y \preceq x\} : x \in X\}$$

in $J(X)$.

5.4 Some Möbius functions

Proposition 5.9 *Let X and Y be posets. The Möbius function of the direct product $X \times Y$ is given by*

$$\mu((x,y),(x',y')) = \mu(x,x')\mu(y,y').$$

Proof It is enough to show that this product satisfies the recurrence it should, i.e.

$$\sum_{x \le x'' \le x', y \le y'' \le y'} \mu(x,x'')\mu(y,y'') = 0.$$

Now the left-hand side of this expression factorises as

$$\left(\sum_{x \le x'' \le x'} \mu(x,x'') \right) \left(\sum_{y \le y'' \le y'} \mu(y,y'') \right)$$

and the inner sum is zero by the recurrence for the Möbius function on X and Y.

Note that there is a natural isomorphism of incidence algebras $I_R(X \times Y) \cong I_{I_R(X)}(Y)$.

(1) The Möbius function of a total order, including the posets **n** and $\mathbb{N}$ and $\mathbb{Z}$ with their usual orders, is

$$\mu(x,y) = \begin{cases} 1 & x = y \\ -1 & x \lessdot y \\ 0 & \text{otherwise.} \end{cases}$$

(2) The Möbius function of the Boolean lattice $\mathscr{B}_n$ is

$$\mu(S,T) = (-1)^{|T|-|S|}$$

for all $S \subseteq T$. For, by Proposition 5.9, $\mu(S,T)$ is equal to the product over all $i \in [n]$ of a Möbius function in $\mathscr{B}_1 \cong \mathbf{2}$, which is -1 if T contains i but S does not, and 1 otherwise.

Thus we recognise the Inclusion-Exclusion Principle as Möbius inversion on the Boolean lattice.

(3) In similar fashion, Proposition 5.9 tells us the Möbius function for the positive integers under divisibility. It is

$$\mu(m,n) = \begin{cases} (-1)^s & \text{if } n/m \text{ is the product of } s \text{ distinct primes} \\ 0 & \text{otherwise.} \end{cases}$$

Example (3), recast as a function in a single argument n/m, is the classical *Möbius function* of number theory, first of that name.

$$\mu(n) = \begin{cases} (-1)^s & \text{if } n \text{ is the product of } s \text{ distinct primes} \\ 0 & \text{otherwise.} \end{cases}$$

This features in classical Möbius inversion:

Corollary 5.10 *Let f and g be functions on the positive integers. Then the following are equivalent:*

- $g(n) = \sum_{m|n} f(x)$;
- $f(n) = \sum_{m|n} g(m)\mu(n/m)$.

For instance, we have $\sum_{m|n} \phi(m) = n$, since each element of $\{0,\dots,n-1\}$ has some greatest common divisor with n, which is a divisor m of n; dividing through by n/m puts these in bijection with the set counted by $\phi(m)$. Möbius inversion then yields

$$\phi(n) = \sum_{m|n} m\mu(n/m)$$

which can be unpacked to the formula in example (3) of Section 5.1.

As a longer example of an application, let us count the monic irreducible polynomials over the finite field of order q. To start with, every monic polynomial is uniquely a product of monic irreducible polynomials. We express that in ordinary generating function machinery: if the number of monic irreducibles of degree k is m_k, then

$$\frac{1}{1-qx} = \prod_{k\geq 1}(1-x^k)^{-m_k}.$$

The right side is the EGF of monic polynomials, while when expanding the left side, we pick a term from one of m_k copies of the geometric series

$$1 + x^n + x^{2k} + \cdots$$

for each k, encoding how many copies of each of the m_k monic irreducibles of degree k we use in the factorisation. Taking logarithms of both sides, we obtain

$$\begin{aligned} \sum_{n\geq 1} \frac{q^n x^n}{n} &= \\ -\log(1-qx) &= \sum_{k\geq 1} -m_k \log(1-x^k) \\ &= \sum_{k\geq 1} m_k \sum_{i\geq 1} \frac{x^{ki}}{i}. \end{aligned}$$

The coefficient of x^n in the last expression is the sum, over all divisors k of n, of $m_k/i = km_k/n$. This must be equal to the coefficient on the left, which is q^n/n. We conclude that

$$q^n = \sum_{k|n} km_k.$$

We cannot omit mentioning that there is also an algebraic proof of this last equality, through which it looks more perspicuous: each of the q^n elements α of the field of order q^n has a unique minimal polynomial, which is monic of some degree k with $k \mid n$, because the field of order q^n is an extension of the field generated by α. A monic polynomial of degree k has k roots, and the equality follows.

At any rate, Möbius inversion can be used on this last equality, giving us the formula

$$m_n = \frac{1}{n} \sum_{d|n} q^d \mu(n/d).$$

(4) We find the Möbius function of the partition lattice Π_n. Note that any interval in Π_n is isomorphically a product of smaller partition lattices: to specify a partition in the interval $[\pi, \rho]$, for each block R of ρ we must specify how those blocks of π contained in R are glued together, and to do this is to choose an element of $\Pi_{\#\{P \in \pi : P \subseteq R\}}$. So it is enough to find $\mu_n := \mu(\hat{0}, \hat{1})$ in Π_n, in terms of n.

The result is that

$$\mu_n = (-1)^{n-1}(n-1)!.$$

I know no very direct proof of this, but one approach uses Möbius inversion in a way cooked so that the single value of the Möbius function we want is easy to isolate. If π is a set partition of $[n]$, how many q-colourings are there of the blocks of π? Clearly, q^{π}. On the other hand, we can count them according to the partition greater than or equal to π induced by merging all parts of the same colour together, resulting in a partition coloured with *distinct* colours. This gives the right side of the formula

$$q^{|\pi|} = \sum_{\rho \geq \pi} (q)_{|\rho|}.$$

Möbius inversion, on the oppoite of the interval $[\pi, \hat{1}]$, transforms this to

$$(q)_{|\pi|} = \sum_{\pi} \mu(\pi, \rho) q^{|\rho|}.$$

Now compare coefficients of the linear term in q: on the right this is $\mu(\pi, \hat{1}) = \mu_{|\pi|}$ since $\hat{1} = \{[n]\}$ is the only set partition with one part, while on the left it is $(-1)^{n-1}(n-1)!$ since we must select the constant term in every factor but the first from the expansion

$$q \cdot (q-1) \cdots (q-(n-1)).$$

Finally, we take a peek at the rich field of topological combinatorics by giving a topological interpretation of the Möbius function. An *abstract simplicial complex* Δ on a set X of vertices is a nonempty collection of subsets (called *faces*) of X closed under taking subsets. These simplicial complexes can also be regarded as topological spaces: the space we associate to Δ, called its *realisation* $|\Delta|$, is a subset of the Euclidean space $\mathbb{R}^X$ with basis $\{e_x : x \in X\}$, given as

$$|\Delta| = \bigcup_{F \in \Delta} \mathrm{conv}\{e_x : x \in F\}.$$

Given such a simplicial complex Δ, let $f_i(\Delta)$ be the number of faces of X with $i+1$ elements; the indexing reflects that the topological counterparts of these faces have dimension i. The *reduced Euler characteristic* of Δ is

$$\tilde{\chi}(\Delta) = -f_{-1}(\Delta) + f_0(\Delta) - f_1(\Delta) + \cdots,$$

agreeing with the reduced Euler characteristic of $|\Delta|$. Note that $f_{-1}(\Delta) = 1$, because $\emptyset \in \Delta$.

The *order complex* $\Delta(X)$ of a poset X is the set of all its chains, which is an abstract simplicial complex. Philip Hall's identity implies

Proposition 5.11 *Let X be a poset, and $\hat{X}$ the poset obtained by adjoining a new minimum element $\hat{0}$ and a new maximum element $\hat{1}$. Then the Möbius function $\mu(\hat{0}, \hat{1})$ in $\hat{X}$ equals $\tilde{\chi}(\Delta(X))$.*

It is also possible to pass from abstract simplicial complex to poset: indeed, the faces of an abstract simplicial complex are themselves a poset, under inclusion.

Proposition 5.12 *Let Δ be an abstract simplicial complex. Then $\tilde{\chi}(\Delta)$ equals the Mobius function $\mu(\hat{0}, \hat{1})$ in the poset obtained by adding to Δ, with the relation of inclusion, a new maximal element $\hat{1}$ (but using the extant $\emptyset \in \Delta$ as $\hat{0}$).*

The passages from complex to poset in Proposition 5.12 and from poset to complex in Proposition 5.11 are not exact inverses of one another. However, beginning with a simplicial complex Δ and carrying out both translations yields the new simplicial complex known by topologists as the *barycentric subdivision* of Δ, which is homeomorphic to Δ.

Proposition 5.12 is easily proved using Rota's Crosscut Theorem, taking the Y in the theorem statement to be the set of singletons.

Theorem 5.13 (Crosscut Theorem) *Let X be a finite lattice, and Y a subset of $X \setminus \hat{0}$ such that every element of $X \setminus \hat{0}$ is greater than or equal to some element of Y. Then*

$$\mu(\hat{0}, x) = \sum_{Z \subseteq Y : \bigvee Z = x} (-1)^{|Z|},$$

where $\bigvee Z$ is short for $\bigvee_{z \in Z} z$.

The proof easiest to state is to observe that the summation at the right hand side satisfies the recurrence (2) defining $\mu(\hat{0}, x)$, because

$$\begin{aligned} \sum_{z \leq x} \sum_{Z \subseteq Y : \bigvee Z = z} (-1)^{|Z|} &= \sum_{Z \subseteq Y : \bigvee Z \leq x} (-1)^{|Z|} \\ &= \sum_{Z \subseteq \{y \in Y : y \leq x\}} (-1)^{|Z|} \\ &= (1 + (-1))^{|\{y \in Y : y \leq x\}|} \end{aligned}$$

by the binomial theorem, and this is 1 if $x = \hat{0}$ and 0 otherwise.