Partially ordered sets

Thomas Britz and Peter Cameron

November 2001

These notes have been prepared as background material for the Combinatorics Study Group talks by Professor Rafael Sorkin (Syracuse University) on the topic **Discrete posets and quantum gravity**, which took place in October–November 2001.

1 Binary relations

We begin by taking a closer look at binary relations $R \subseteq X \times X$. Figure 1 shows four of the ways in which to look at a binary relation: as a set R, as a bipartite graph G, as a directed graph D, as an incidence matrix M, where the stars $*$ of the latter are often replaced by numbers or variables. These four ways are completely equivalent but their context is quite varied. For instance, we may be interested in matchings, in which case the bipartite view would be most appropriate; or, we might be interested in paths, and the directed graph view would be more suitable; and so on. Often it is even more useful to translate one context into another. Typical instances of such a translation are when linear algebra is applied to the incidence matrix M in order to describe attributes of the three other structures. In the first two of the following examples, we assume that X contains only finitely many elements.

Example 1.1 Replace each star $*$ of the matrix M with an independent variable. Then a subset $A \subseteq X$ is matched in (or, a partial transversal of) the bipartite graph G if and only if the rows of M corresponding to the elements of A are linearly independent. (The set $A = (a_i)_I \subseteq X$ is matched in G if there is a set $B = (b_i)_I \subseteq X$ such that $(a_i, b_i) \in R$ for all $i \in I$.)

Example 1.2 Replace each star $*$ of the matrix M by the integer 1. Then the (a,b)'th entry of M^k equals the number of paths in D from the vertex a to the vertex b.

Example 1.3 Suppose we have two relations $R, S \subseteq X$, with corresponding incidence matrices M and N. Replace each star $*$ of the matrices M and N by the Boolean 1 (i.e. 1+1=1). Then $M+N$ is the incidence matrix of the relation $R \cup S$.

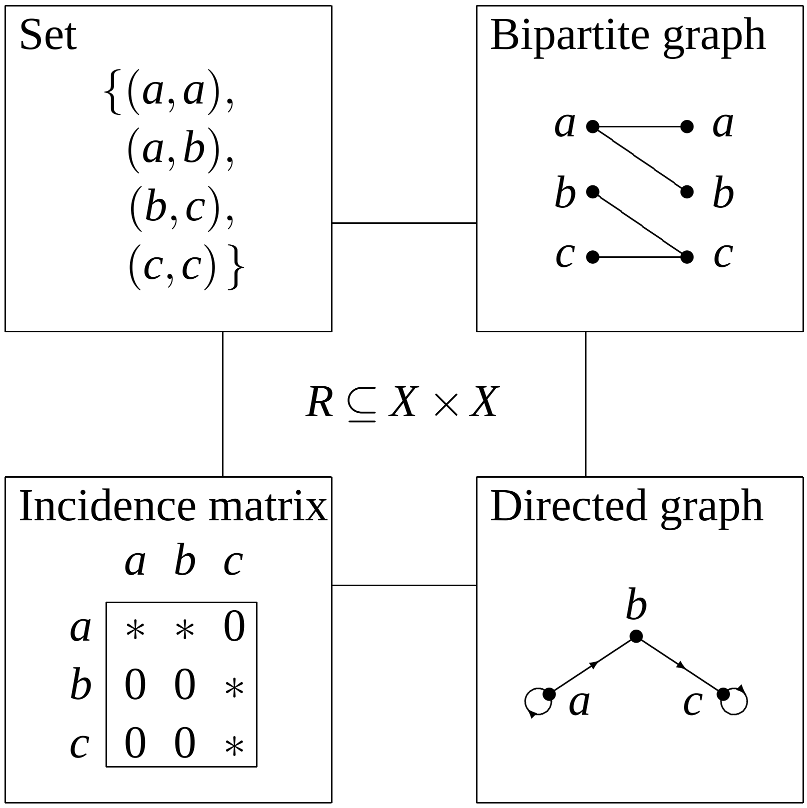

Figure 1: There are always four sides to a relation

The empty set $\emptyset$ is a relation, and as sets, relations may be operated upon by complement R^C, intersection $\cap$, and union $\cup$. Apart from these, we also introduce the identity relation D_X, the inverse relation R^{-1}, and the composition operation $\circ$:

$$\begin{array}{lcl} D_X & = & \{(x,x); x \in X\}; \\ R^{-1} & = & \{(y,x); (x,y) \in R\}; \\ R \circ R' & = & \{(x,z); (x,y) \in R \text{ and } (y,z) \in R' \text{ for some } y \in X\}. \end{array}$$

The inverse R^{-1} is easy to visualise: in terms of sets, the order of each element (x,y) in the relation R is reversed; the two parts of G are interchanged; the direction of each arc of D is reversed; and the matrix M is transposed. The composition $\circ$ is not hard to visualise either. Figure 2 illustrates the composition in terms of bipartite graphs.

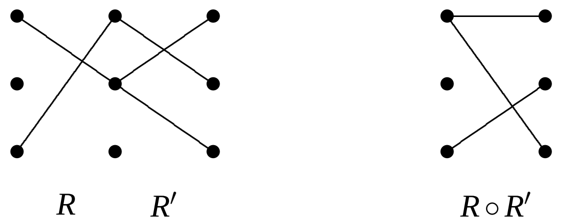

Figure 2: Composition of binary relations on a set

Example 1.4 Let $R, S \subseteq X$ be two relations on X, with corresponding incidence matrices M and N. Replace each star $*$ of the matrices M and N by the Boolean 1. If X contains only finitely many elements, then $M \cdot N$ is the incidence matrix of the relation $R \circ S$.

The *transitive closure* $\overline{R}$ of a relation R is the relation

$$D_X \cup R \cup (R \circ R) \cup (R \circ R \circ R) \cup \ldots .$$

It corresponds precisely to the transitive closure of the directed graph D, that is the graph obtained by adding to D the arc (a,b) whenever D contains a path from a to b. In other words, the transitive closure $\overline{R}$ is the smallest transitive relation on S which contains R. If X contains a finite number n of elements, then by replacing each star $*$ of M by a Boolean 1, we obtain an incidence matrix of $\overline{R}$, namely M^n.

2 What is a poset?

The term "poset" is short for "partially ordered set", that is, a set whose elements are ordered but not all pairs of elements are required to be comparable in the order. Just as an order in the usual sense may be strict (as $<$) or non-strict (as $\leq$), there are two versions of the definition of a partial order:

A *strict partial order* is a binary relation S on a set X satisfying the conditions

(R$-$) for no $x \in X$ does $(x,x) \in S$ hold;

(A$-$) if $(x,y) \in S$, then $(y,x) \notin S$;

(T) if $(x,y) \in S$ and $(y,z) \in S$, then $(x,z) \in S$.

A *non-strict partial order* is a binary relation R on a set X satisfying the conditions

(R$+$) for all $x \in X$ we have $(x,x) \in R$;

(A) if $(x,y) \in R$ and $(y,x) \in R$ then $x = y$;

(T) if $(x,y) \in R$ and $(y,z) \in R$ then $(x,z) \in R$.

Condition (A$-$) appears stronger than (A), but in fact (R$-$) and (A) imply (A$-$). So we can (as is usually done) replace (A$-$) by (A) in the definition of a strict partial order. Conditions (R$-$), (R$+$), (A), (T) are called *irreflexivity*, *reflexivity*, *antisymmetry* and *transitivity* respectively. We can restate these conditions in terms of the identity D_X, the inverse R^{-1}, and the composition operation $\circ$:

(R$-$) $D_X \cap R = \emptyset$;

(R$+$) $D_X \subseteq R$;

(A) $R \cap R^{-1} \subseteq D_X$;

(T) $R \circ R \subseteq R$.

The two definitions of a poset are essentially the same:

Proposition 2.1 *Let X be a set.*

(a) If S is a strict partial order on X, then $S \cup D_X$ is a non-strict partial order on X.

(b) If R is a non-strict partial order on X, then $R \setminus D_X$ is a strict partial order on X.

(c) These two constructions are mutually inverse. ■

Exercise: Prove this proposition.

Thus, a *poset* is a set X carrying a partial order (either strict or non-strict, since we can obtain each from the other in a canonical way). If we have to choose, we use the non-strict partial order. If R and S are corresponding non-strict and strict partial orders, we write $x \leq_R y$ to mean $(x,y) \in R$, and $x <_R y$ to mean $(x,y) \in S$; thus $x \leq_R y$ holds if and only if either $x <_R y$ or $x = y$. (The slightly awkward notation $<_R$ means that we regard R as the name of the partial order.) If there is no ambiguity about R, we simply write $x \leq y$ or $x < y$ respectively.

Exercise: How many different posets are there with 3 elements?

A *total order* is a partial order in which every pair of elements is comparable, that is, the following condition (known as *trichotomy*) holds:

- for all $x, y \in X$, exactly one of $x <_R y$, $x = y$, and $y <_R x$ holds.

In a poset (X, R), we define the *interval* $[x,y]_R$ to be the set

$$[x,y]_R = \{z \in X : x \leq_R z \leq_R y\}.$$

By transitivity, the interval $[x,y]_R$ is empty if $x \not\leq_R y$. We say that the poset is *locally finite* if all intervals are finite.

The set of positive integers ordered by divisibility (that is, $x \leq_R y$ if x divides y) is a locally finite poset.

3 Preorders

Sometimes we need to weaken the definition of a partial order. We say that a *partial preorder* or *pseudo-order* is a relation R on a set X which satisfies conditions (R) (reflexivity) and (T) (transitivity). So it is permitted that distinct elements x and y satisfy $(x,y) \in R$ and $(y,x) \in R$.

Proposition 3.1 *Let R be a partial preorder on X. Define a relation $\sim$ on X by the rule that $x \sim y$ if and only if $(x,y), (y,x) \in R$. Then $\sim$ is an equivalence relation on X. Moreover, if $x \sim x'$ and $y \sim y'$, then $(x,y) \in R$ if and only if $(x',y') \in R$. Thus, R induces in a natural way a relation $\overline{R}$ on the set $\overline{X}$ of equivalence classes of X; and $\overline{R}$ is a non-strict partial order on $\overline{X}$.* ■

The partial order obtained in this way is the *canonical quotient* of the partial preorder R.

Exercise: Prove this proposition.

Exercise: How many different partial preorders are there on a set of 3 elements?

Many of the definitions for posets are also valid for preorders: chains, antichains, upsets, downsets, minimal and maximal elements, local finiteness (see below), and so on. However, the intuition behind these definitions is sometimes different than for posets. For instance, a non-trivial finite chain does not necessarily have a maximal element.

A less well-known characterisation of finite preorders is in terms of the incidence matrix M. Replace the stars $*$ of the incidence matrix M of some reflexive relation R on X by independent variables, or more precisely, elements which are algebraically independent over some field. Then R is a preorder if and only if the inverse matrix M^{-1} is also an incidence matrix of R.

4 Properties of posets

An element x of a poset (X,R) is called *maximal* if there is no element $y \in X$ satisfying $x <_R y$. Dually, x is *minimal* if no element satisfies $y <_R x$.

In a general poset there may be no maximal element, or there may be more than one. But in a finite poset there is always at least one maximal element, which can be found as follows: choose any element x; if it is not maximal, replace it by an element y satisfying $x <_R y$; repeat until a maximal element is found. The process must terminate, since by the irreflexive and transitive laws the chain can never revisit any element. Dually, a finite poset must contain minimal elements.

An element x is an *upper bound* for a subset Y of X if $y \leq_R x$ for all $y \in Y$. *Lower bounds* are defined similarly. We say that x is a *least upper bound* or *l.u.b.* of Y if it is an upper bound and satisfies $x \leq_R x'$ for any upper bound x'. The concept of a *greatest lower bound* or *g.l.b.* is defined similarly.

A *chain* in a poset (X,R) is a subset C of X which is totally ordered by the restriction of R (that is, a totally ordered subset of X). An *antichain* is a set A of pairwise incomparable elements.

Infinite posets (such as $\mathbb{Z}$), as we remarked, need not contain maximal elements. *Zorn's Lemma* gives a sufficient condition for maximal elements to exist:

> Let (X,R) be a poset in which every chain has an upper bound. Then X contains a maximal element.

As well known, there is no “proof” of Zorn’s Lemma, since it is equivalent to the Axiom of Choice (and so there are models of set theory in which it is true, and models in which it is false). Our proof of the existence of maximal elements in finite posets indicates why this should be so: the construction requires (in general infinitely many) choices of upper bounds for the elements previously chosen (which form a chain by construction).

The *height* of a poset is the largest cardinality of a chain, and its *width* is the largest cardinality of an antichain. We denote the height and width of (X,R) by $h(X)$ and $w(X)$ respectively (suppressing as usual the relation R in the notation).

In a finite poset (X,R), a chain C and an antichain A have at most one element in common. Hence the least number of antichains whose union is X is not less than the size $h(X)$ of the largest chain in X. In fact there is a partition of X into $h(X)$ antichains. To see this, let A_1 be the set of maximal elements; by definition this is an antichain, and it meets every maximal chain. Then let A_2 be the set of maximal elements in $X \setminus A_1$, and iterate this procedure to find the other antichains.

There is a kind of dual statement, harder to prove, known as *Dilworth’s Theorem*:

Theorem 4.1 *Let (X,R) be a finite poset. Then there is a partition of X into $w(X)$ chains.* ■

An *up-set* in a poset (X,R) is a subset Y of X such that, if $y \in Y$ and $y \leq_R z$, then $z \in Y$. The set of minimal elements in an up-set is an antichain. Conversely, if A is an antichain, then

$$\uparrow(A) = \{x \in X : a \leq_R x \text{ for some } a \in A\}$$

is an up-set. These two correspondences between up-sets and antichains are mutually inverse; so the numbers of up-sets and antichains in a poset are equal.

Down-sets are, of course, defined dually. The complement of an up-set is a down-set; so there are equally many up-sets and down-sets.

5 Hasse diagrams

Let x and y be distinct elements of a poset (X,R). We say that *y covers x* if $[x,y]_R = \{x,y\}$; that is, $x <_R y$ but no element z satisfies $x <_R z <_R y$. In general, there may be no pairs x and y such that y covers x (this is the case in the rational numbers, for example). However, locally finite posets are determined by their covering pairs:

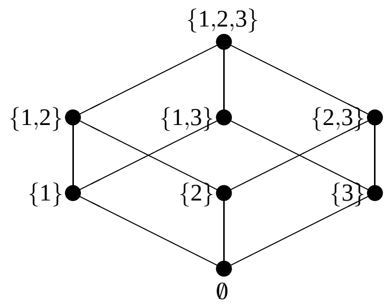

Figure 3: A Hasse diagram

Proposition 5.1 *Let (X,R) be a locally finite poset, and $x,y \in X$. Then $x \leq_R y$ if and only if there exist elements $z_0,\ldots,z_n$ (for some non-negative integer n) such that $z_0 = x$, $z_n = y$, and z_{i+1} covers z_i for $i = 0,\ldots,n-1$.* ■

Exercise: Prove this proposition.

The *Hasse diagram* of a poset (X,R) is the directed graph whose vertex set is X and whose arcs are the covering pairs (x,y) in the poset. We usually draw the Hasse diagram of a finite poset in the plane in such a way that, if y covers x, then the point representing y is higher than the point representing x. Then no arrows are required in the drawing, since the directions of the arrows are implicit.

For example, the Hasse diagram of the poset of subsets of $\{1,2,3\}$ is shown in Figure 3.

Exercise: Find the height and width of the poset in Figure 3. Count its chains and antichains of maximum size, and verify the conclusion of Dilworth's theorem in this case.

Note that the Hasse diagram of a poset corresponds to an operation on posets, or more generally acyclic directed graphs, which is dual to the transitive closure. This operation, R°, which we might call the transitive opening, acts by removing from R an element (a,b) whenever there is a path from a to b, distinct from the arc (a,b), in the corresponding directed graph D. Note that $\overline{R} = \overline{R^\circ}$ and that $R^\circ = \overline{R}^\circ$. If R is a poset, then R° is the unique minimal relation among the relations whose transitive closure is R; indeed, R° is the intersection of these. This explains why the Hasse diagram representation is unique.

Locally finite preorders have unique representations which are similar to Hasse diagrams of posets. Indeed, we may even use labeled Hasse diagrams, where the dots represent subsets of X, rather than just elements as in the case of posets. For instance, consider the Hasse diagram in Figure 4a. The elements f, g, and h are all smaller than i and larger than d and a, and $f \leq g \leq h \leq f$. This sort of representation is well-defined and unique since the property that both $x \leq y$ and $y \leq x$ hold defines an equivalence relation. If we wish to do without the labeling, we may replace each dot by a cluster of dots (see Figure 4b).

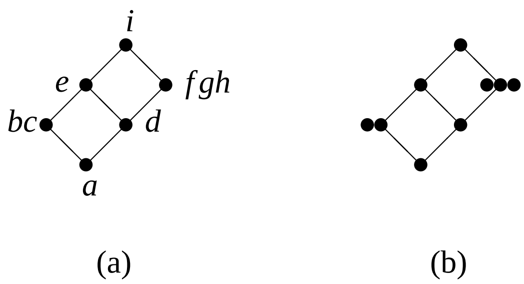

Figure 4: Hasse diagrams of preorders

6 Linear extensions and dimension

One view of a partial order is that it contains partial information about a total order on the underlying set. This view is borne out by the following theorem. We say that one relation *extends* another if the second relation (as a set of ordered pairs) is a subset of the first.

Theorem 6.1 *Any partial order on a finite set X can be extended to a total order on X.*

This theorem follows by a finite number of applications of the next result. ■

Proposition 6.2 *Let R be a partial order on a set X, and let a, b be incomparable elements of X. Then there is a partial order R' extending R such that $(a,b) \in R'$ (that is, $a < b$ in the order R').*

Proof It is just a matter of checking the consequences of putting a below b. Let $A = \{x : x \leq_R a\}$ and $B = \{y : b \leq_R y\}$. Clearly, if the extension of R is possible, then every element of A must lie below every element of B. So we take $R' = R \cup (A \times B)$, in other words, we include these obvious consequences of putting a below b and no others. We have to show that R' is a partial order. This is just a matter of checking a number of cases.

Note first that A and B are disjoint. For, if $x \in A \cap B$, then $x \leq_R a$ and $b \leq_R x$, so $b \leq_R a$, contrary to assumption.

Suppose that $x \leq_{R'} y$ and $y \leq_{R'} x$. Then either $x \leq_R y$ or $(x,y) \in A \times B$, and similarly either $y \leq_R x$ or $(y,x) \in A \times B$. This gives four cases, of which $(x,y),(y,x) \in A \times B$ is impossible by the previous remark. If $x \leq_R y$ and $y \leq_R x$ then $x = y$; if $x \leq_R y$ and $(y,x) \in A \times B$, then $b \leq_R x$, $x \leq_R y$, $y \leq_R a$, so $b \leq_R a$, contrary to assumption, with a similar contradiction in the other case.

For the transitive law, suppose that $x \leq_{R'} y \leq_{R'} z$. Then either $x \leq_R y$ or $(x,y) \in A \times B$, and similarly either $y \leq_R z$ or $(y,z) \in A \times B$. Again there are four cases to consider. This reduces to three, since $y \in A$ and $y \in B$ cannot both occur. If $x \leq_R y$ and $y \leq_R z$, then $x \leq_R z$. If $x \leq_R y$ and $y \in A, z \in B$, then $y \leq a$, so $x \leq a$, and $x \in A$, giving $(x,z) \in A \times B$. The remaining case is similar. ■

A total order extending R in this sense is referred to as a *linear extension* of R. (The term "linear order" is an alternative for "total order".)

Exercise: Find the number of linear extensions of each of the 3-element posets.

This proof does not immediately show that every infinite partial order can be extended to a total order. If we assume Zorn's Lemma, the conclusion follows. For let S be a maximal element (under inclusion) in the set of partial orders extending R. Then S must be a total order, since if a and b were incomparable in S then Proposition 6.2 would give an extension S' of S such that $a <_{S'} b$, contradicting the maximality of S. (We have first to show that the union of a chain of posets is a poset. This is a standard Zorn's Lemma argument.)

It is known that the truth of the infinite analogue of Theorem 6.1 cannot be proved from the Zermelo–Fraenkel axioms alone (assuming their consistency), but is strictly weaker than the Axiom of Choice, that is, the Axiom of Choice (or Zorn's Lemma) cannot be proved from the Zermelo–Fraenkel axioms and this assumption. In other words, assuming the axioms consistent, there is a model in which Theorem 6.1 is false for some infinite poset, and another model in which Theorem 6.1 is true for all posets but Zorn's Lemma is false.

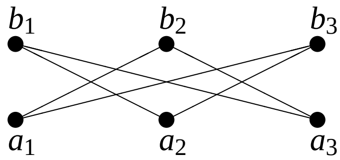

Figure 5: A crown

The theorem gives us another measure of the size of a partially ordered set. To motivate this, we use another model of a partial order. Suppose that a number of products are being compared using several different attributes. We regard object a as below object b if b beats a on every attribute. If each beats the other on some attributes, we regard the objects as being incomparable. This defines a partial order (assuming that each attribute gives a total order). More precisely, given a set S of total orders on X, we define a partial order R on X by $x <_R y$ if and only if $x <_s y$ for every $s \in S$. In other words, R is the intersection of the total orders in S.

Theorem 6.3 *Every partial order on a finite set X is the intersection of some set of total orders on X.*

Proof Let R be a partial order on X, and let S be the set of all total orders which extend R. The intersection of the orders in S certainly contains R. We show it is no bigger. So suppose that a and b are incomparable in R. By Proposition 6.2, there is a total order extending R in which a is less than b, and another in which b is less than a. So in the intersection of these total orders, a and b are still incomparable. ∎

Now we define the *dimension* of a partial order R to be the smallest number of total orders whose intersection is R. In our motivating example, it is the smallest number of attributes which could give rise to the observed total order R.

The *crown* on $2n$ elements $a_1,\ldots,a_n,b_1,\ldots,b_n$ is the partial order defined as follows: for all indices $i \neq j$, the elements a_i and a_j are incomparable, the elements b_i and b_j are incomparable, but $a_i < b_j$; and for each i, the elements a_i and b_i are incomparable. Figure 5 shows the Hasse diagram of the 6-element crown.

Now we have the following result:

Proposition 6.4 *The crown on $2n$ elements has dimension n.* ∎

Exercise: Prove this proposition.

Exercise: Find the dimension of each of the 3-element posets.

Proposition 6.4 shows that the dimension of a finite partial order may be arbitrarily large. For infinite posets (assuming the axiom of choice), the dimension exists but may be infinite. However, a standard application of the Compactness Theorem gives the following:

Theorem 6.5 *Let R be a partial order on a set X, and suppose that the restriction of R to any finite subset of X has dimension at most n (for some integer n). Then R has dimension at most n.*

Proof We take a first order language containing $n+1$ binary relation symbols $R, L_1, \ldots, L_n$, and a constant symbol c_x for each $x \in X$. Let Σ consist of the following set of sentences: a sentence asserting that $R(x,y)$ is equivalent to the conjunction of $L_1(x,y), \ldots, L_n(x,y)$; sentences asserting that each L_i is a linear order; and sentences asserting that the elements c_x for $x \in X$ are all distinct and that $R(c_x, c_y)$ holds if and only if $x \leq y$ in the given poset. Any finite subset of these sentences is satisfiable. So, by the Compactness Theorem, they are all satisfiable, and the result follows. ■

7 Posets and topologies

The number of topologies on an infinite set is greater than the number of relational structures of any fixed type. However, on a finite set, a topology is equivalent to a particular type of relational structure, namely a partial preorder.

A *topology* consists of a set X, and a set $\mathcal{T}$ of subsets of X (called *open sets*), satisfying the following axioms:

- $\emptyset \in \mathcal{T}$ and $X \in \mathcal{T}$;
- the union of any collection of sets in $\mathcal{T}$ is in $\mathcal{T}$;
- the intersection of any two sets in $\mathcal{T}$ is in $\mathcal{T}$.

It follows by induction from the third axiom that the intersection of any finite number of members of $\mathcal{T}$ is a member of $\mathcal{T}$. If X is finite, the second axiom need only deal with finite unions, and so it too can be simplified to the statement that the union of any two sets in $\mathcal{T}$ is in $\mathcal{T}$; then the axioms are 'self-dual'. This is not the case in general!

Theorem 7.1 *Let X be finite. Then there is a one-to-one correspondence between the topologies on X, and the partial preorders (i.e., reflexive and transitive relations) on X.*

Proof The correspondence is simple to describe.

Construction 1 Let $\mathcal{T}$ be a topology on X. Define a relation R by the rule that $(x,y) \in R$ if every open set containing x also contains y. It is trivial that R is reflexive and transitive; that is, R is a partial preorder.

Construction 2 Let R be a partial preorder on X. Call a subset U of X *open* if, whenever $x \in U$, we have $\uparrow x \subseteq U$. (Recall that $\uparrow x = \{y : (x,y) \in R\}$.) Let $\mathcal{T}$ be the set of all open sets. We have to verify that $\mathcal{T}$ is a topology. The first axiom requires no comment. For the second axiom, let $U_1, U_2, \ldots$ be open, and $x \in \bigcup_i U_i$; then $x \in U_j$ for some j, whence

$$\uparrow x \subseteq U_j \subseteq \bigcup_i U_i.$$

For the third axiom, let U and V be open and $x \in U \cap V$. Then $R(x) \subseteq U$ and $R(x) \subseteq V$, and so $R(x) \subseteq (U \cap V)$; thus $U \cap V$ is open.

All this argument is perfectly general. It is the fact that we have a bijection which depends on the finiteness of X. We have to show that applying the two constructions in turn brings us back to our starting point.

Suppose first that R is a partial preorder, and $\mathcal{T}$ the topology derived from it by Construction 2. Suppose that $(x,y) \in R$. Then $y \in \uparrow x$, so every open set containing x also contains y. Conversely, suppose that every open set containing x also contains y. The set $\uparrow x$ is itself open (this uses the transitivity of R: if $z \in \uparrow x$, then $\uparrow z \subseteq \uparrow x$), and so $y \in \uparrow x$; thus $(x,y) \in R$. Hence the partial preorder derived from $\mathcal{T}$ by Construction 1 coincides with R. (We still haven't used finiteness!)

Conversely, let $\mathcal{T}$ be a topology, and R the partial preorder obtained by Construction 1. If $U \in \mathcal{T}$ and $x \in U$, then $\uparrow x \subseteq U$; so U is open in the sense of Construction 2. Conversely, suppose that U is open in this sense, that is, $x \in U$ implies $\uparrow x \subseteq U$. Now each set $\uparrow x$ is the intersection of all members of $\mathcal{T}$ containing x. (This follows from the definition of R in Construction 1.) But there are only finitely many such open sets (here, at last, we use the fact that X is finite!); and the intersection of finitely many open sets is open, as we remarked earlier; so $\uparrow x$ is open in $\mathcal{T}$. But, by hypothesis, U is the union of the sets $\uparrow x$ for all points $x \in U$; and a union of open sets is open, so U is open in $\mathcal{T}$, as required. ∎

In the axiomatic development of topology, the next thing one meets after the definition is usually the so-called 'separation axioms'. A topology is said to satisfy the axiom T_0 if, given any two distinct points x and y, there is an open set containing one but not the other; it satisfies axiom T_1 if, given distinct x and y, there is an open set containing x but not y (and *vice versa*).

These two axioms for finite topologies have a natural interpretation in terms of the partial preorder R. Axiom T_1 asserts that R never holds between distinct points x and y; that is, R is the trivial relation of equality. Construction 2 in the proof of the theorem then shows that every subset is open. (This is called the *discrete topology*.) It follows that any stronger separation axiom (in particular, the so-called 'Hausdorff axiom' T_2) also forces the topology to be discrete.

Axiom T_0 translates into the condition that the relation R is antisymmetric; thus, it is a partial order. So there is a one-to-one correspondence between T_0 topologies on the finite set X and partial orders on X.

We conclude with an extension of this principle closely related to Rafael Sorkin's views: we describe how an arbitrary topological space can be approximated by posets. Let $(X,\mathcal{T})$ be a topological space, and let $\mathcal{U}$ be an open cover of X (that is, a subset of $\mathcal{T}$ whose union is X) which is *locally finite* (that is, every point of X lies in only finitely many members of $\mathcal{U}$). Now we define a relation R on X as follows: $(x,y) \in R$ if and only if every open set in $\mathcal{U}$ containing x also contains y. This relation is always a partial preorder (by the same argument as above), and so it has a canonical quotient which is a partial order. This partial order approximates the original space X: points are identified if they are not distinguished by the open sets in $\mathcal{U}$. Moreover, the partial order is locally finite.

Now if we take a sequence of successively finer open coverings of X by open sets, we obtain a sequence of partial orders, whose "inverse limit" captures the structure of $(X,\mathcal{T})$, at least in nice cases. We do not give further details here.

8 The Möbius function

Let R be a partial order on the finite set X. We take any linear order extending R, and write $X = \{x_1,\ldots,x_n\}$, where $x_1 < \ldots < x_n$ (in the linear order S): this is not essential but is convenient later.

The *incidence algebra* $\mathcal{A}(R)$ of R is the set of all functions $f : X \times X \to \mathbb{R}$ which satisfy $f(x,y) = 0$ unless $x \leq_R y$ holds. We could regard it as a function on R, regarded as a set of ordered pairs. Addition and scalar multiplication are defined pointwise; multiplication is given by the rule

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

If we represent f by the $n \times n$ matrix A_f with (i,j) entry $f(x_i,x_j)$, then this is precisely the rule for matrix multiplication. Also, if $x \not\leq_R y$, then there is no point z such that $x \leq_R z$ and $z \leq_R y$, and so $(fg)(x,y) = 0$. Thus, $\mathcal{A}(R)$ is closed under multiplication and does indeed form an algebra, a subset of the matrix algebra $M_n(\mathbb{R})$. Also, since f and g vanish on pairs not in R, the sum can be restricted to the interval $[x,y]_R = \{z : x \leq_R z \leq_R y\}$:

$$(fg)(x,y) = \sum_{z \in [x,y]_R} f(x,z)g(z,y).$$

Incidentally, we see that the (i,j) entry of A_f is zero if $i > j$, and so $\mathcal{A}(R)$ consists of upper triangular matrices. Thus, an element $f \in \mathcal{A}(R)$ is invertible if and only if $f(x,x) \neq 0$ for all $x \in X$.

The *zeta-function* ζ_R is the element of $\mathcal{A}(R)$ defined by

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

Its inverse (which also lies in $\mathcal{A}(R)$) is the *Möbius function* μ_R of R. Thus, we have, for all $(x,y) \in R$,

$$\sum_{z \in [x,y]_R} \mu(x,z) = \begin{cases} 1 & \text{if } x = y, \\ 0 & \text{otherwise.} \end{cases}$$

This relation allows the Möbius function of a poset to be calculated recursively. We begin with $\mu_R(x,x) = 1$ for all $x \in X$. Now, if $x <_R y$ and we know the values of $\mu(x,z)$ for all $z \in [x,y]_R \setminus \{y\}$, then we have

$$\mu_R(x,y) = -\sum_{z \in [x,y]_R \setminus \{y\}} \mu_R(x,z).$$

In particular, $\mu_R(x,y) = -1$ if y covers x.

The definition of the incidence algebra and the Möbius function extend immediately to locally finite posets, since the sums involved are over intervals $[x,y]_R$.

Exercise: Find the Möbius function of the poset shown in Figure 3.

Exercise: Let X be the set of positive integers and R the relation of divisibility. Prove that $\mu_R(1,n)$ is equal to the classical Möbius function $\mu(n)$, which is defined to be $(-1)^k$ if n is the product of k distinct primes, and 0 if n is not squarefree. Prove also that $\mu_R(m,n) = \mu(n/m)$ if m divides n (and is zero otherwise).

Exercise: The *disjoint union* of posets (X_1,R_1) and (X_2,R_2) is defined to be $(X_1 \cup X_2, R_1 \cup R_2)$. In other words, if $x \in X_1$ and $y \in X_2$, then x and y are incomparable; but comparability within each part remains unchanged. Describe a Hasse diagram, the number of linear extensions, the incidence algebra, and the Möbius function, of the disjoint union in terms of those of the two parts.

The following are examples of Möbius functions.

- The subsets of a set:
  $\mu(A,B) = (-1)^{|B \setminus A|}$ for $A \subseteq B$;
- The subspaces of a vector space $V \subseteq GF(q)^n$:
  $\mu(U,W) = (-1)^k q^{\binom{k}{2}} \quad$ for $U \subseteq W$, where $k = \dim U - \dim W$.

Much of the work done on incidence algebras of posets has been concentrated on finding tools and methods with which to determine the Möbius function of various classes of posets. It may seem odd that it generally is quite hard to determine the Möbius function of a poset; after all, it just amounts to inverting an upper triangular $(0,1)$-matrix. The problem lies in that we wish to express the Möbius function in terms of general properties of the poset, and not in terms of the particular zeta function. However, the efforts in finding the Möbius function are well-rewarded, as the remaining part of this section indicate.

The following trivial result is the Möbius inversion for locally finite posets.

Theorem 8.1 *$f = g\zeta \Leftrightarrow g = f\mu$. Similarly, $f = \zeta g \Leftrightarrow g = \mu f$.*

Example 8.1 Suppose that f and g are functions on the natural numbers which are related by the identity $f(n) = \sum_{d|n} g(d)$. We may express this identity as $f = g\zeta$ where we consider f and g as vectors and where ζ is the zeta function for the lattice of positive integer divisors of n. Theorem 8.1 implies that $g = f\mu$, or

$$g(n) = \sum_{d|n} \mu(d,n) f(d) = \sum_{d|n} \mu\left(\frac{d}{n}\right) f(d),$$

which is precisely the classical Möbius inversion.

Example 8.2 Suppose that f and g are functions on the subsets of some fixed (countable) set X which are related by the identity $f(A) = \sum_{B \supseteq A} g(B)$. We may express this identity as $f = \zeta g$ where ζ is the zeta function for the lattice of subsets of X. Theorem 8.1 implies that $g = \mu f$, or

$$g(A) = \sum_{B \supseteq A} \mu(A,B) f(B) = \sum_{B \supseteq A} (-1)^{|B \setminus A|} f(B)$$

which is a rather general form of the inclusion/exclusion principle.

9 Lattices

A *lattice* is a poset (X,R) with the properties

- X has an upper bound 1 and a lower bound 0;
- for any two elements $x,y \in X$, there is a least upper bound and a greatest lower bound of the set $\{x,y\}$.

A simple example of a poset which is not a lattice is the poset 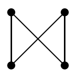.

In a lattice, we denote the l.u.b. of $\{x,y\}$ by $x \vee y$, and the g.l.b. by $x \wedge y$. We commonly regard a lattice as being a set with two distinguished elements and two binary operations, instead of as a special kind of poset.

Lattices can be axiomatised in terms of the two constants 0 and 1 and the two operations $\vee$ and $\wedge$. The result is as follows, though the details are not so important for us. The axioms given below are not all independent. In particular, for finite lattices we don't need to specify 0 and 1 separately, since 0 is just the meet of all elements in the lattice and 1 is their join.

Proposition 9.1 *Let X be a set, $\wedge$ and $\vee$ two binary operations defined on X, and 0 and 1 two elements of X. Then $(X, \vee, \wedge, 0, 1)$ is a lattice if and only if the following axioms are satisfied:*

- *Associative laws: $x \wedge (y \wedge z) = (x \wedge y) \wedge z$ and $x \vee (y \vee z) = (x \vee y) \vee z$;*
- *Commutative laws: $x \wedge y = y \wedge x$ and $x \vee y = y \vee x$;*
- *Idempotent laws: $x \wedge x = x \vee x = x$;*

- $x \wedge (x \vee y) = x = x \vee (x \wedge y)$;
- $x \wedge 0 = 0$, $x \vee 1 = 1$.

Proof The proof will not be given here (you may regard it as an exercise). We remark merely that the order must be defined by $x \leq y$ if $x \vee y = y$ (this is equivalent to $x \wedge y = x$, by the second-last axiom). ■

A *sublattice* of a lattice is a subset of the elements containing 0 and 1 and closed under the operations $\vee$ and $\wedge$. It is a lattice in its own right.

The following are a few examples of lattices.

- The subsets of a (fixed) set:
$$\begin{aligned} A \wedge B &= A \cap B \\ A \vee B &= A \cup B \end{aligned}$$
- The subspaces of a vector space:
$$\begin{aligned} U \wedge V &= U \cap V \\ U \vee V &= \mathrm{span}(U \cup V) \end{aligned}$$
- The partial pseudo-orders on a set:
$$\begin{aligned} R \wedge T &= R \cap T \\ R \vee T &= \overline{R \cup T} \end{aligned}$$

The last example has as a sublattice the equivalence relations on a set, with the same l.u.b. and g.l.b. The posets on a set do not form a lattice. See C. H. Yan, *Discrete Math.* **183** (1998), 285–292, for this example.

These examples illustrate the typical situation when the order is defined by set-inclusion: the g.l.b. is the intersection, and the l.u.b. is an appropriate closure of the union. A closely related example of a lattice is the subsets of a multiset (see Figure 6a).

If the multiset is finite, the lattice may be viewed from an alternative viewpoint. In particular, it is the lattice of positive divisors of a fixed integer, which has as g.l.b. $\wedge$ and l.u.b. $\vee$ the greatest common denominator and the least common factor. For instance, the lattice of positive integer divisors of 18 (Figure 6b) corresponds to the lattice of submultisets of the multiset $\{a, b, b\}$ (Figure 6a).

10 Distributive and modular lattices

A lattice is *distributive* if it satisfies the *distributive laws*

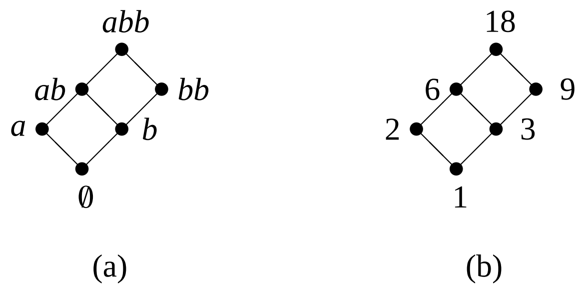

Figure 6: A class of lattices

(D) $x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$ and $x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z)$ for all x, y, z.

A lattice is *modular* if it satisfies the *modular law*

(M) $x \vee (y \wedge z) = (x \vee y) \wedge z$ for all x, y, z such that $x \leq z$.

Exercise: Prove that a distributive lattice is modular.

Exercise: Prove that the two conditions in (D) are equivalent for any lattice.

Exercise: State the modular law as a law valid for all choices of the variables.

Figure 7 presents a lattice, N_5, which is not modular, as well as a modular lattice, M_3, which is not distributive.

Not only are N_5 and M_3 the smallest lattices with these properties, they are, in a certain sense, the only lattices with these properties. The following theorem states this more precisely.

Theorem 10.1 *A lattice is modular if and only if it does not contain the lattice N_5 as a sublattice. A lattice is distributive if and only if it contains neither the lattice N_5 nor the lattice M_3 as a sublattice.*

In the same way as in Proposition 9.1, we are able to describe distributive lattices axiomatically.

Proposition 10.2 *Let X be a set, $\wedge$ and $\vee$ two binary operations defined on X, and 0 and 1 two elements of X. Then $(X, \vee, \wedge, 0, 1)$ is a distributive lattice if and only if the condition (D) and the following axioms are satisfied:*

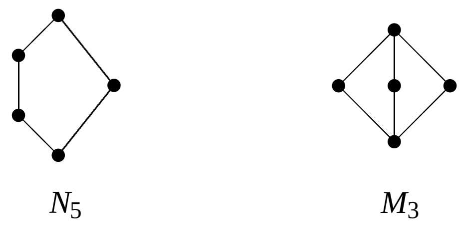

Figure 7: Two lattices

- *Idempotent law:* $x \wedge x = x$*;*
- $x \vee 1 = 1 \vee x = 1$*;*
- $x \vee 0 = 0 \vee x = 0$.

Exercise: Prove this.

Proposition 10.2 seems (erroneously) to suggest that fewer conditions are needed for a distributive lattice than for lattices in general. This is due to the fact that the distributive conditions (D) are strong enough to imply, together with the three conditions stated in Proposition 10.2, the second idempotent law, as well as the associative and commutative laws.

The poset of all subsets of a set S (ordered by inclusion) is a distributive lattice: we have $0 = \emptyset$, $1 = S$, and l.u.b. and g.l.b. are union and intersection respectively. Hence every sublattice of this lattice is a distributive lattice.

Conversely, every finite distributive lattice is a sublattice of the lattice of subsets of a set. We describe how this representation works. This is important in that it gives us another way to look at posets.

Let (X,R) be a poset. Recall that an *down-set* in X is a subset Y with the property that, if $y \in Y$ and $z \leq_R y$, then $z \in Y$.

Let L be a lattice. A non-zero element $x \in L$ is called *join-irreducible* if, whenever $x = y \vee z$, we have $x = y$ or $x = z$.

Theorem 10.3 *(a) Let (X,R) be a finite poset. Then the set of down-sets in X, with the operations of union and intersection and the distinguished elements $0 = \emptyset$ and $1 = X$, is a distributive lattice.*

(b) Let L be a finite distributive lattice. Then the set X of non-zero join-irreducible elements of L is a sub-poset of L.

(c) These two operations are mutually inverse. ■

Meet-irreducible elements are defined dually, and there is of course a dual form of Theorem 10.3.

Figure 8 illustrates a simple algorithm with which to extract the poset from the corresponding distributive lattice.

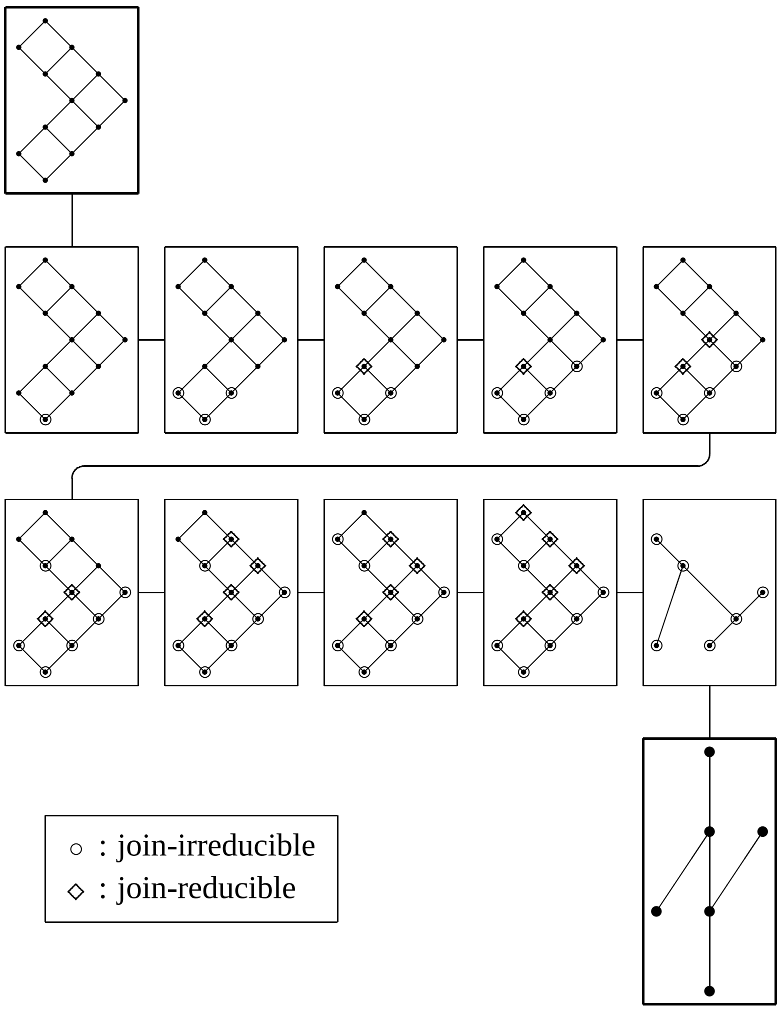

Figure 8: Posets and distributive lattices