Chapter 12

Partially Ordered Sets

In a totally ordered set, such as the natural numbers under "$\leq$", any two elements are related. Other orderings are "partial"; examples include the divisibility relation on positive integers and the inclusion relation on subsets of a set. Partially ordered sets can model precedence constraints in scheduling, preferences among objects, partial information about numbers being sorted, etc.

12.1. Structure of Posets

DEFINITIONS AND EXAMPLES

12.1.1. Example. The subsets of a finite set, ordered by inclusion, form a natural poset. We spent most of Chapter 11 studying aspects of it. Elementary understanding of containment yields three natural properties: (1) $A \subseteq A$, (2) $A \subseteq B$ and $B \subseteq A$ together imply $A = B$, and (3) If $A \subseteq B$ and $B \subseteq C$, then $A \subseteq C$.

Similarly, the divisibility relation on the divisors of an integer N satisfies: (1) $x \mid x$, (2) $x \mid y$ and $y \mid x$ imply $x = y$, and (3) If $x \mid y$ and $y \mid z$, then $x \mid z$. ■

12.1.2. Definition. A **relation** R on a set X is a subset of the cartesian product $X \times X$. We write $(x, y) \in R$ or xRy or say (x, y) *satisfies* R. An **order relation** (or **partial order**) on X is a relation that is

reflexive (xRx for all x),
antisymmetric (xRy and yRx imply $x = y$), and
transitive (xRy and yRz imply xRz).

A **partially ordered set** (or **poset**) is a set P plus an order relation on P. When (x, y) satisfies the relation, we write $x \leq_P y$ or simply $x \leq y$. If $x \leq y$ or $y \leq x$, then x and y are **comparable** (or *related*) in P; otherwise they are **incomparable** (or *unrelated*). All of $\leq$, $<$, $\geq$, $>$ describe ways that elements can be comparable. We write $x \| y$ to mean that x and y are incomparable.

The word "poset" illustrates the evolution of terminology. Originally it was written as "PO-set", emphasizing grammatically that a poset is a set equipped with a partial order.

Like any binary relation, an order relation is specified by a $0,1$-matrix, with 1 in position (x, y) if $x \le y$. To facilitate study, we seek visual representations. Since an order relation on P is a set of ordered pairs from P, we can view it as a directed graph with vertex set P. Ignoring which way pairs are ordered treats the edges as unordered pairs.

12.1.3. Definition. The **comparability digraph** of a poset P is the digraph whose vertices are the elements of P and whose edges are the ordered pairs xy such that $x \le_P y$. The **comparability graph** is the graph whose vertices are the elements of P and whose edges are the unordered pairs of distinct vertices that are comparable in P.

12.1.4. Remark. *Comparability graphs and digraphs.* The comparability digraph specifies a poset completely. Since order relations are reflexive, a comparability digraph has a loop at each vertex. Since this is understood, we omit the loops when drawing a comparability digraph and just draw the digraph of the **strict order relation**: the pairs xy with $x < y$.

The comparability graph discards information and does not specify the poset. Below we show distinct posets with the same comparability graph. Reversing all comparable pairs also yields another poset with the same comparability graph.

A comparability graph becomes a comparability digraph by transitively orienting the edges. A **transitive orientation** of a graph G is an orientation such that if xy and yz are (directed) edges, then xz is an edge; that is, x and z are adjacent in G and the orientation directs the edge from x to z. The transitive orientations of G correspond to the posets for which G is the comparability graph.

Exercise 6 gives a necessary condition for comparability graphs that turns out also to be sufficient. Another characterization in Exercise 7 leads to a fast algorithm for testing whether a graph is a comparability graph. Nevertheless, our focus in this chapter is on the order-theoretic properties of posets, so we will not study comparability graphs further. ■

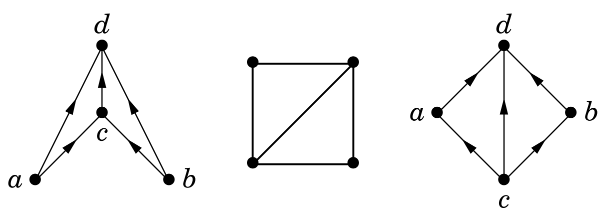

For a useful visual presentation, transitivity of the order relation allows us to describe a poset completely while drawing only some of the comparable pairs.

12.1.5. Definition. If $x < y$ in P and there is no z with $x < z < y$, then y **covers** x (in P), written as $x \prec y$ or $y \succ x$. The **cover relation** is the set of pairs (x, y) such that $x \prec y$. The **cover digraph** is the digraph on the elements of P whose edge set is $\{xy \colon x \prec y\}$.

A **cover diagram** (or **Hasse diagram** or **diagram**) of P is obtained by erasing the directions on edges after drawing the cover digraph in the plane such that each (straight-line) edge points upward. The **cover graph** is the graph on the elements of P whose edge set is $\{xy \colon x \prec y \text{ or } y \prec x\}$.

12.1.6. Example. *The diagram.* The comparability digraphs in Remark 12.1.4 direct all edges upward. Erasing the edges implied by transitivity yields the cover digraphs. Since the edges all point upward, erasing the arrowheads then yields the diagrams, which appear on the left below. One of these is the poset of subsets of $\{1, 2\}$, ordered by inclusion.

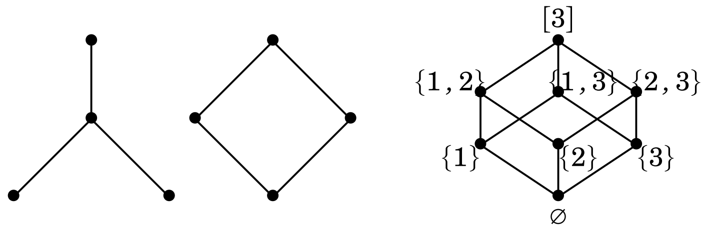

On the right is the poset of subsets of $[3]$, ordered by inclusion. The diagram specifies the poset completely; $x < y$ in P if and only if we can reach y from x in the diagram by moving only upward along edges. Omitting the edges implied by transitivity makes it easier to see the structure. By the convention that the edges for the order relation point upward, the diagram provides the cover digraph and, by transitive closure, the comparability digraph. ■

The cover graph (but not the diagram) discards the order information and does not specify the poset. For example, C_4 is the cover graph of two posets (one shown above), and C_6 is the cover graph of seven posets (Exercise 3). In contrast to comparability graphs, testing whether a graph is a cover graph is NP-complete (Nešetřil–Rödl [1987, 1993]; Brightwell [1993] gave a simpler proof).

The simplest posets are totally ordered or totally unordered. Indeed, we say "partial order" as a generalization of total order.

12.1.7. Definition. A **chain** is a poset whose elements are linearly ordered so that $x < y$ if and only if x comes before y in that order. The chain with k elements is denoted $\mathbf{k}$ and called a k-**chain**. An **antichain** is a poset in which no two elements are comparable.

An element of a poset is **maximal** if no other element is greater than it, **minimal** if no other element is less than it.

The first poset in Remark 12.1.4 has one maximal element (d) and two minimal elements (a and b). A k-chain has one maximal and one minimal element. In an antichain, every element is maximal and minimal. The cover graph of $\mathbf{k}$ is the path P_k; its comparability graph is K_k.

12.1.8. Definition. Posets P and Q are **isomorphic** if some bijection from the set P to the set Q preserves the order relation. A **subposet** of a poset P is a poset R on a subset of P defined by restricting the comparability relation to R. If Q is isomorphic to a subposet of P, then P **contains** Q or Q **embeds** in P.

The **dual** of P, written P^*, is the poset on the elements of P defined by $y \le_{P^*} x$ if and only if $x \le_P y$. A poset isomorphic to its dual is **self-dual**. A poset is **finite** if it has finitely many elements.

The notion of isomorphism allows us to view a diagram of a poset as the poset itself, just as a drawing is a graph. We use $\mathbf{k}$ for a k-element chain, as an isomorphism class, just as C_k denotes a k-vertex cycle.

12.1.9. Example. *Subposets.* If $n \le m$, then $\mathbf{n}$ embeds in $\mathbf{m}$, and all n-element subposets of $\mathbf{m}$ are isomorphic to $\mathbf{n}$.

We write $\mathbf{2}^n$ for the inclusion poset on the subsets of $[n]$, rather than $\mathbf{2}^{[n]}$ as in Section 11.3; we will soon explain why. In $\mathbf{2}^3$ (Example 12.1.6) we find six chains of size 4, two antichains of size 3, and 15 subposets isomorphic to $\mathbf{2}^2$ (Exercise 2). Furthermore, $\mathbf{2}^3$ is self-dual, via complementation of subsets of $[3]$.

A chain in a poset is a set of pairwise comparable elements; an antichain is a set of pairwise incomparable elements. They become cliques and independent sets in the comparability graph. When $x_1, \ldots, x_k$ is a chain in P, with $x_1 < \cdots < x_k$, *it need not hold* that x_{i+1} covers x_i. For example, $\mathbf{2}^3$ contains 19 chains of size 2. ■

12.1.10. Remark. *Subposet* in posets corresponds to *induced subgraph* in graph theory. When G is the comparability graph of P, and Q is a subposet of P, the comparability graph of Q is the subgraph of G induced by the set Q. However, the cover graph of Q need not be a subgraph of the cover graph of P.

A graph consists of a vertex set and an adjacency relation on it. Similarly, a poset consists of a set of elements with an order relation on it. Nevertheless, we generally use the same notation (P) for a partially ordered set and for its set of elements. This abuse of notation works because we treat P as a partially ordered *set* and study as "subposets" only the structures that inherit all comparabilities.

For graphs we manipulate both vertices and edges, so we write $v \in V(G)$ and $e \in E(G)$ instead of $v \in G$ and $e \in G$ to avoid ambiguity. For posets, we just write $x \in P$ when x is in the set of elements, and we use $|P|$ for the **size** of that set. We rarely consider analogues of edge deletion or contraction for posets. ■

12.1.11. Definition. Let P be a poset. Its **width** $w(P)$ is the size of a largest antichain in P. Its **height** $h(P)$ is the size of a largest chain in P. Its **length** is one less than its height.

12.1.12. Example. The poset $\mathbf{2}^3$ has width 3, height 4, and length 3. An antichain of size $w(P)$ in P is a **maximum antichain**, but a **maximal antichain** (one not contained in a larger antichain) may be smaller. For example, the antichain $\{\{1\}, \{2, 3\}\}$ in $\mathbf{2}^3$ is a maximal antichain but not a maximum antichain.

A chain in a poset with finite height is a **maximal chain** if and only if extends from a minimal element to a maximal element and its successive pairs satisfy the cover relation. For elements, "maximal" and "minimal" refer to the order relation in P, not to containment on sets of elements. ■

12.1.13. Definition. In a poset P, a **down-set** is a subset D such that $x \in D$ and $y < x$ imply $y \in D$. The complement of a down-set is an **up-set**. The down-set $D[x]$ **generated by** x is $\{y \in P\colon y \le x\}$. The down-set $D[A]$ generated by a family A is $\{y\colon y \le x$ for some $x \in A\}$. We also write $D(x) = \{y \in P\colon y < x\}$. Similarly define $U[x]$, $U[A]$, and $U(x)$.

Down-sets have also been called **ideals** or **order ideals**, and up-sets have been called **dual ideals** or **filters**. The terms "down-set" and "up-set" are less sophisticated, but they are explicit and do not conflict with other uses of "ideal" and "filter" in mathematics. In Chapter 11, we used "ideal of sets" for a down-set in $\mathbf{2}^n$ because this is common and is consistent with algebraic notions.

In $\mathbf{2}^3$, the down-set generated by $\{1, 23\}$ is $\{1, 23, 2, 3, \varnothing\}$. There are exactly 20 antichains and 20 down-sets in $\mathbf{2}^3$; one of each is empty.

12.1.14. Remark. *There is a one-to-one correspondence between the antichains and the down-sets in any poset.* The natural bijection maps an antichain A to the down-set $D[A]$. The inverse assigns to each down-set the antichain consisting of its maximal elements. The empty antichain corresponds to the empty down-set. ∎

The product operation is a fundamental way to combine posets.

12.1.15. Definition. The **product** $P \times Q$ of two posets P, Q is the poset on $\{(x, y)\colon x \in P, y \in Q\}$ defined by $(x, y) \le (x', y')$ if and only if $x \le x'$ and $y \le y'$. The **sum** $P + Q$ consists of disjoint copies of P and Q with no comparabilities between them.

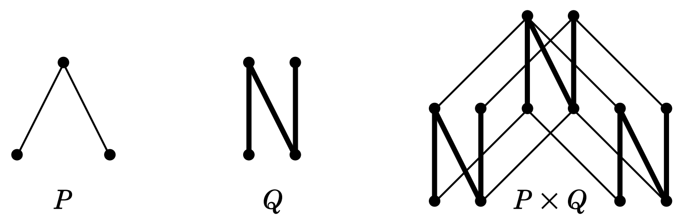

12.1.16. Remark. *Products of posets.* The product $P \times Q$ has disjoint copies of Q for each element of P and disjoint copies of P for each element of Q. If the cover graphs of P and Q are G and H, then the cover graph of $P \times Q$ is the cartesian product $G \,\square\, H$. However, the comparability graph of $P \times Q$ is not the cartesian product of the comparability graphs of P and Q.

In the sense of isomorphism, the operation is commutative and associative, and we write the product of n copies of P as P^n. ∎

12.1.17. Example. *The subset poset.* The poset $\mathbf{2}$ is a 2-element chain; call its elements 0 and 1, with $0 < 1$. With each element of $[n]$, associate a copy of $\mathbf{2}$. Let P be the poset of subsets of $[n]$, ordered by inclusion. Mapping each member of P to its incidence vector expresses P as $\mathbf{2}^n$. Thus we denote P as $\mathbf{2}^n$. (We suggest pronouncing $\mathbf{2}^n$ as "$\mathbf{2}$ sup n".) For the inclusion poset on the subsets of a set X, we write $\mathbf{2}^X$ and call X the **ground set**. The poset $\mathbf{2}^n$ is also called the **Boolean algebra** on n elements, sometimes denoted B_n.

Since $\mathbf{2}^n$ is the product of n copies of $\mathbf{2}$, its cover graph is the n-dimensional cube. Note that if x is a k-set and y is an l-set in $[n]$ with $x < y$, then the subposet of $\mathbf{2}^n$ consisting of $\{z\colon x \le z \le y\}$ is isomorphic to $\mathbf{2}^{l-k}$. Also, $\mathbf{2}^n$ is self-dual (via complementation of subsets of $[n]$). Finally, when writing elements of $\mathbf{2}^n$ as sets or as their incidence vectors, we typically drop set brackets, commas, and parentheses and just write strings of elements. ∎

DILWORTH'S THEOREM AND BEYOND

The archetypal extremal problem for posets is finding a largest antichain. Explicit answers exist only for special classes; we discuss these later. First, we obtain a min-max relation. Since an antichain has no two elements on any chain, the width is bounded by the number of chains needed to cover the elements. Dilworth proved that equality holds.

12.1.18. Theorem. (**Dilworth's Theorem**; Dilworth [1950]) If P is a finite poset, then the maximum size of an antichain in P equals the minimum number of chains needed to cover the elements of P.

Proof: (Perles [1963]) We use induction on $|P|$, with trivial basis. For $|P| > 1$, suppose first that some largest antichain A omits both a maximal element and a minimal element of P. Thus neither the down-set $D[A]$ nor the up-set $U[A]$ generated by A contains all of P, and we can apply the induction hypothesis to each subposet. Also, they share only A.

Let $k = w(P)$. Since both $D[A]$ and $U[A]$ are subposets of P, they have width at most k. Since they both contain A, equality holds. From the induction hypothesis, we obtain k chains covering $D[A]$ and k chains covering $U[A]$, the former with elements of A at the top and the latter with elements of A at the bottom. These combine to form k chains. These chains cover all of P, because the maximality of A implies that $D[A] \cup U[A] = P$. This case is shown on the left below.

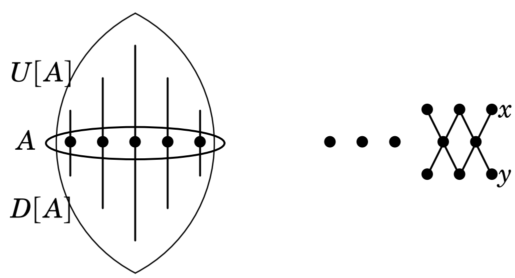

In the remaining case, every maximum antichain of P consists of all maximal elements or all minimal elements. Thus $w(P - \{x, y\}) \le k - 1$ if x is a minimal element and y is a maximal element. Choose a maximal element x and a minimal element y below it (they may be equal). Since $P - \{x, y\}$ is smaller (and has width $k-1$), we can apply the induction hypothesis to cover $P - \{x, y\}$ with $k-1$ chains and add the chain consisting of x and y to complete the desired covering. ∎

Since subposets of chains are chains, we conclude that P is covered by $w(P)$ *pairwise disjoint* chains. A partition of P into chains is a **chain decomposition** of P. In honor of Theorem 12.1.18, a smallest chain decomposition (size $w(P)$) is called a **Dilworth decomposition** of P.

Perles' proof is like Pym's proof of Menger's Theorem. Each uses splicing of paths, with a special argument for a degenerate case. Here we seek a maximum cut (the antichain) instead of a minimum cut. In the language of network flows, the arguments are essentially equivalent.

Meanwhile, we relate Dilworth's Theorem to the min-max relations of graph theory. A poset is **bipartite** if every element is minimal or maximal (or both); equivalently, the comparability graph is bipartite.

12.1.19. Theorem. (Fulkerson [1956]) Dilworth's Theorem is equivalent to the König–Egerváry Theorem on matchings in bipartite graphs.

Proof: To apply Dilworth's Theorem to bipartite matching, view an n-vertex bipartite graph G as a bipartite poset. The vertices of one part become maximal elements; the others are minimal. Chains have one or two elements. Every covering of the poset by $n-k$ chains uses k disjoint 2-chains; these yield a matching of size k in G. Each antichain corresponds to an independent set in G; if it is maximal, then the remaining vertices form a vertex cover. Hence Dilworth's equality between the sizes of a maximum antichain and minimum chain-covering yields a matching and a vertex cover of the same size in G. Since every vertex cover is at least as large as every matching, this proves the König–Egerváry Theorem.

For the converse implication, consider any poset P of size n. We define a bipartite graph $S(P)$, the **split** of P, in which to study matchings. For each element $x \in P$, create two vertices x^- and x^+ (see figure below). The parts of the bipartition of $S(P)$ are $\{x^-: x \in P\}$ and $\{x^+: x \in P\}$. The edge set is $\{x^-y^+: x <_P y\}$.

Every matching in $S(P)$ yields a chain-covering in P as follows: if x^-y^+ is in the matching, then y is immediately above x on a chain in the cover. If x^- or x^+ is unmatched, then x is the top or bottom of its chain, respectively. Since each vertex of $S(P)$ appears in at most one edge of the matching, this defines disjoint chains covering P. If the matching has k edges, then the cover has $n-k$ chains, since each additional edge links the top of one chain with the bottom of another to form a single chain.

Given a vertex cover R of $S(P)$, let $A = \{x \in P: x^-, x^+ \notin R\}$. A relation between elements of A would yield an edge of $S(P)$ uncovered by R, so A is an antichain. No *minimal* vertex cover of $S(P)$ uses both of $\{x^+, x^-\}$, because by transitivity the sets $\{z^-: z \in D(x)\}$ and $\{y^+: y \in U(x)\}$ induce a complete bipartite subgraph in $S(P)$. A vertex cover of a complete bipartite graph must use all of one part. Since x^+ and x^- have no other neighbors, we can drop from R the one of $\{x^+, x^-\}$ that is not in that part.

Thus a minimum vertex cover of size k yields an antichain of size $n-k$. Now the existence of a matching and a vertex cover in $S(P)$ of equal size yields an antichain and a chain-covering of equal size in P. ∎

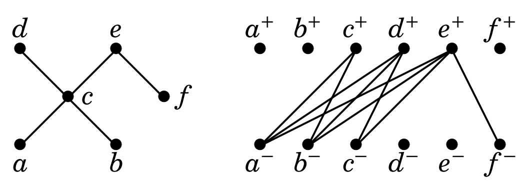

12.1.20. Remark. A simpler min-max relation holds for maximum chains and minimum antichain coverings (Mirsky [1971]). The antichain of maximal elements intersects each maximal chain, so by induction on $h(P)$ we can cover P with $h(P)$ antichains. Via the Perfect Graph Theorem (Corollary 8.3.34), this easy observation *implies* Dilworth's Theorem. ∎

Dilworth's original proof was a complicated induction, cutting and pasting chains. An appropriate generalization is simpler. Here $\alpha(D)$ denotes the independence number of the graph obtained by viewing the edges as unordered pairs.

12.1.21. Theorem. (Gallai–Milgram [1960]) The vertices of an n-vertex digraph D can be covered using at most $\alpha(D)$ disjoint paths.

Proof: Since $V(D)$ is covered by n disjoint paths of length 0, it suffices to prove a stronger claim: If $\mathbf{C}$ is a set of k disjoint paths covering $V(D)$, with $k > \alpha(G)$, and S is the set of sources (initial vertices) of these paths, then $V(D)$ can be covered using fewer than k disjoint paths with sources in S. The statement holds vacuously when $n = 1$; we proceed by induction. The added requirement about sources simplifies the induction step.

Suppose $n > 1$. Since $k > \alpha(D)$, there is an edge xy with $x, y \in S$. Let A and B be the paths in $\mathbf{C}$ starting with x and y. If A has no edge, then moving x to the beginning of B saves one path. Hence A has an edge xz. Deleting x from A yields a cover $\mathbf{C}'$ of $V(D - x)$ by k paths with sources in S', where $S' = S - x + z$. Since $\alpha(D - x) \le \alpha(D)$, the induction hypothesis yields a cover $\mathbf{C}''$ of $V(D - x)$ by fewer than k paths with sources in S'.

All of S' is in S except z. If z is a source in $\mathbf{C}''$, then prepend x to the path starting at z to cover $V(D)$ using fewer than k paths. If z is not a source but y is, then prepend x to the path starting at y. If neither y nor z is a source, then $|S'| = k$ implies that $\mathbf{C}''$ has at most $k - 2$ paths, and letting x be a path by itself covers $V(D)$ using fewer than k paths. In all cases, the resulting paths are disjoint and have sources in S. ■

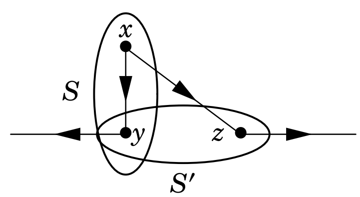

In the special case where D is the comparability digraph of a poset P, the initial covering is the set of $|P|$ trivial paths. Since $\alpha(D) = w(P)$, the paths given by Theorem 12.1.21 form the desired chain-covering, a Dilworth decomposition.

Dilworth's Theorem also generalizes beyond antichains.

12.1.22. Definition. A **k-family** in a poset is a subposet containing no $(k+1)$-chain. The **k-norm** of a set partition $C_1, \dots, C_m$ is $\sum_{i=1}^{m} \min\{k, |C_i|\}$. A chain decomposition of a poset P is **k-saturated** if its k-norm equals the maximum size of a k-family in P.

Dilworth decompositions are 1-saturated partitions. By Mirsky's observation, k-families are just the unions of (at most) k antichains. Each chain C in a chain partition contributes at most $\min\{k, |C|\}$ to a k-family. Hence the size of a k-family is bounded by the k-norm of a chain partition.

12.1.23. Theorem. Greene–Kleitman Theorem; Greene–Kleitman [1976a]) For every poset P and natural number k, there is a chain decomposition of P that is both k-saturated and $(k+1)$-saturated. ■

The original proof of the Greene–Kleitman Theorem was quite long. There are now shorter combinatorial proofs (Saks [1979]) and proofs applying other min-max relations (such as Frank [1980b] via minimum-cost network flow). We leave these to a more advanced book.

For acyclic digraphs, there are extensions of the Greene–Kleitman Theorem due to Linial [1981], Saks [1980], Cameron [1982], and Hoffman [1983]. Berge conjectured an extension for arbitrary digraphs. Let **C** be a partition of the vertices of a digraph D using paths. A partition with smallest k-norm is **k-optimal**. On the other hand, a **partial k-coloring** is a union of k independent sets of vertices; view each set as a color class. The number of colors on a path C is bounded by $\min\{k, |V(C)|\}$. We seek a partial k-coloring where equality holds.

12.1.24. Conjecture. (Berge [1982]) For every digraph D, integer k, and k-optimal path partition **C** of D, there is a partial k-coloring such that each path C in **C** intersects $\min\{k, |V(C)|\}$ color classes. ■

Berge [1985] noted that the Greene–Kleitman Theorem is the special case for transitive digraphs; he proved it also for $k = 1$ in general. Cameron [1986] proved it for acyclic digraphs, Berger–Hartman [2008] proved it for $k = 2$, and Berger–Hartman [2012] gave a unified proof for various cases using network flows. The general conjecture remains open; Hartman [2006] surveyed partial results.

EXERCISES 12.1

12.1.1. (−) For $n = 3$ and $n = 4$, list the isomorphism classes of posets with n elements. Determine how many are self-dual.

12.1.2. (−) Show that $\mathbf{2}^3$ contains 15 copies of $\mathbf{2}^2$.

12.1.3. (−) Draw the diagrams of the seven posets (isomorphism classes) whose cover graph is a 6-cycle. Determine their widths.

12.1.4. (−) By Remark 12.1.20, every k-family decomposes into k antichains. Use the poset below to show that it may not be possible to obtain a maximum 2-family by adding an antichain to a maximum antichain.

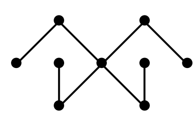

12.1.5. Prove that the graphs below are comparability graphs by exhibiting cover diagrams of posets for which they are the comparability graphs. Use the necessary condition of Exercise 6 to prove that their complements are not comparability graphs.

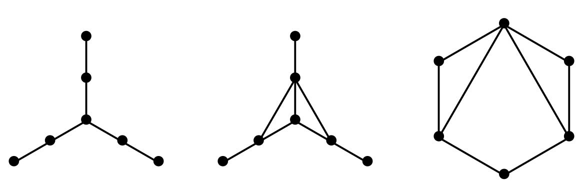

12.1.6. ($\diamond$) For a closed walk $[x_1, \ldots, x_k]$ in a graph, a **triangular chord** is an edge of the form $x_i x_{i+2}$ (with indices modulo k). Prove that in a comparability graph, every closed walk of odd length has a triangular chord. Conclude that the complement of the cycle C_n is not a comparability graph when $n \geq 5$. (Comment: Gilmore–Hoffman [1964] showed that this necessary condition is also sufficient.)

12.1.7. ($\diamond$) For a graph G, the **Ghouilà-Houri graph** G' is the graph defined on the ordered pairs of adjacent vertices of G by putting $(x, y) \leftrightarrow (y, z)$ in G' if and only if $xz \notin E(G)$. Prove that G' is bipartite if and only if every closed odd walk of G has a triangular chord. (Comment: With Exercise 12.1.6, this yields a polynomial-time recognition algorithm for comparability graphs.) (Ghouilà-Houri [1962])

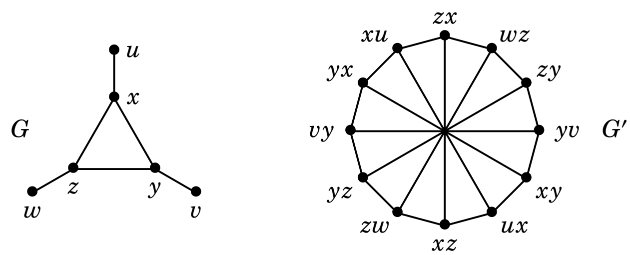

12.1.8. ($\diamond$) Prove that a graph is the cover graph of some poset if and only if it has an acyclic orientation without dependent edges, where a **dependent edge** in a digraph is an edge whose reversal completes a cycle. Conclude that if the chromatic number of a graph is less than its girth, then it is a cover graph.

12.1.9. Use Exercise 12.1.8 to prove that the **Grötzsch graph** below is not a cover graph. (Fisher–Fraughnaugh–Langley–West [1997])

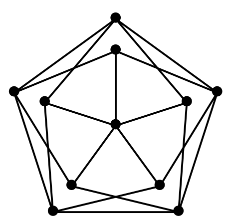

12.1.10. A digraph D is a **cover digraph** if $E(D)$ is the cover relation of some poset. Derive a polynomial-time checkable characterization of cover digraphs. Assume that one can find in polynomial time all vertices reachable from a specified vertex. (Comment: No efficient algorithm is known to check whether an undirected graph is a cover graph.)

12.1.11. Obtain simple formulas (constant number of terms) for

(a) the number of chains of size 2 and the number of chains of size 3 in $\mathbf{2}^n$.

(b) the numbers of antichains of size 2 and size 3 in $\mathbf{2}^n$. (Popadić [1970])

12.1.12. A **cutset** in a partial order is a set of elements intersecting every maximal chain. For each n, prove or disprove: the maximum size of a minimal cutset in $\mathbf{2}^n$ is $\binom{n}{\lceil n/2 \rceil}$. (See Füredi–Griggs–Kleitman [1989])

12.1.13. ($\diamond$) Prove that a poset of size greater than mn has a chain of size greater than m or an antichain of size greater than n. Use this to prove the Erdős–Szekeres Theorem: every list of $mn + 1$ distinct integers has an increasing sublist with more than m elements or a decreasing sublist with more than n elements.

12.1.14. Suppose that the red subgraph in a red/blue-coloring of $E(K_n)$ has a transitive orientation. Prove that the coloring has a monochromatic complete subgraph with at least $\sqrt{n}$ vertices. Show that the bound is sharp.

12.1.15. (◇) A family of sets is **union-free** if it has no two distinct members whose union is a third member. Moser asked for $f(n)$, the maximum size of a union-free subfamily that can be guarantee to exist in any family of n sets.

(a) Use Dilworth's Theorem to prove $f(n) \geq \sqrt{n}$. (Riddell, Erdős–Komlós [1970])

(b) Prove $f(n) > \sqrt{2n} - 1$. (Erdős–Shelah [1972], Kleitman [1973])

(c) Let $A_{i,j}$ consist of the integers from $-i$ to $+j$, and let $F = \{A_{i,j}\colon (i,j) \in [t]^2\}$. Prove that the maximum size of a union-free family in F is $2t-1$. Thus $f(n) \leq 2\sqrt{n} - 1$ when n is a perfect square. (Erdős–Shelah [1972])

(Comment: Fox–Lee–Sudakov [2012] proved $f(n) = \lfloor \sqrt{4n+1} \rfloor - 1$ for all n.)

12.1.16. For $w = (w_1, \ldots, w_d)$, let P denote the poset whose elements are the d-tuples x with $0 \leq x_i < w_i$ for all i, ordered by $x < y$ if and only if $x_i < y_i$ for all i. Prove that P has the property that every maximal antichain is a maximum antichain, and determine its size. (Tsai [2017])

12.1.17. Use the Matroid Intersection Theorem (Theorem 11.3.44) to prove the Gallai–Milgram Theorem (Theorem 12.1.21) when D is an acyclic digraph: the vertices of D can be covered with at most $\alpha(D)$ pairwise-disjoint paths, where $\alpha(D)$ is the maximum number of pairwise nonadjacent vertices. (Chappell)

12.1.18. *Another proof of Dilworth's Theorem.* The poset on the right below arises from that on the left by deleting the central cover relation. All other comparabilities remain.

(a) Suppose that x covers y and z in P. Let Q and R be the posets obtained by deleting (y,x) and (z,x), respectively, from the set of relations in P. Prove that $\min\{w(Q), w(R)\} = w(P)$. (Hint: Take maximum antichains in Q and R, and consider the maximal elements and the minimal elements of the union of these antichains as a subposet of P.)

(b) Use part (a) to prove Dilworth's Theorem. (Harzheim [1983])

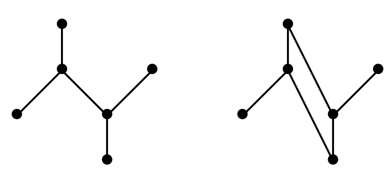

12.1.19. For each poset below and each k, find a chain partition that is both k-saturated and $(k+1)$-saturated. Is some chain partition k-saturated for all k?

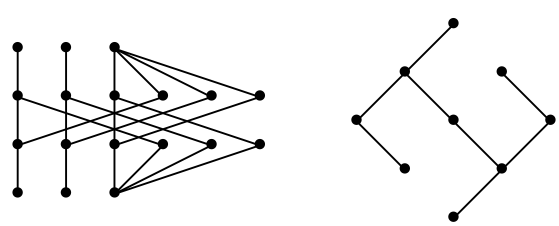

12.1.20. (◇) Let d_k be the maximum size of a k-family in a poset P, and let $\Delta_k = d_k - d_{k-1}$.

(a) In a k-saturated chain partition $\mathbf{C}$ of P, let α be the number of chains of size at least k. Prove $\Delta_k \geq \alpha \geq \Delta_{k+1}$, (Comment: The Greene–Kleitman Theorem thus implies $\Delta_1 \geq \cdots \geq \Delta_h$, where P has height h. No direct proof is known.) (Greene–Kleitman [1976a])

(b) Prove $\alpha = \Delta_k$ when $\mathbf{C}$ is both k-saturated and $(k-1)$-saturated.

12.1.21. A poset P is **polyunsaturated** (West [1986]) if for $1 \le k < l - 1 < h(P)$ no chain partition of P is both k-saturated and l-saturated.

(a) Use the Greene–Kleitman Theorem and Exercise 12.1.20 to prove that $h(P) \le w(P) + 2$ when P is polyunsaturated.

(b) Construct a poset P_k iteratively as follows. For $k = 1$, let $P_k = \mathbf{3}$. For $k > 1$, obtain P_k from P_{k-1} by adding a chain C of $k + 1$ new elements and making the element just below the maximal element on C cover the element just below the maximal element of the (unique) longest chain in P_{k-1}. The diagrams of P_2 and P_3 appear below. Prove that P_k is polyunsaturated. (Comment: Thus the maximum height of a polyunsaturated poset of width k is $k + 2$.) (Chappell [2002])

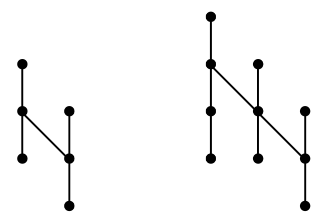

12.2. Symmetric Chains and LYM

Although Dilworth's Theorem applies to all (finite) posets, much of the study of partially ordered sets concerns special families of posets. When suitable constraints are placed on posets, much more can be said about their maximum antichains and other structural aspects.

GRADED POSETS

In this section we consider only finite posets. This property makes the notion of the "rank" of an element well defined.

12.2.1. Definition. In a poset, the **rank** of an element x, written $r(x)$, is the maximum length of a chain having x as its top element. A poset P is **graded** if all its maximal chains have the same length, and its **rank** $r(P)$ is that length.

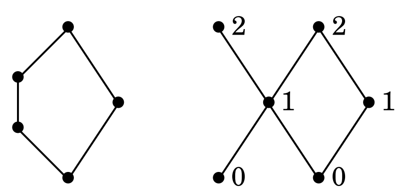

not graded graded

In a graded poset, $r(x) = r(y) + 1$ when x covers y. The rank of the poset is the rank of its maximal elements. The notion of rank function can be extended to more general posets, but we will study only graded posets in this section.

12.2.2. Definition. If P is graded, then the elements with rank k are the kth **rank** or kth **level** P_k, and we write $N_k(P)$ for the **rank size** $|P_k|$ (also called the kth **Whitney number** of P). A graded poset is **rank-symmetric** if $N_k = N_{r(P)-k}$ for all k. It is **rank-unimodal** if there is a rank k such that $N_i \le N_j$ whenever $i \le j \le k$ or $i \ge j \ge k$. The **rank generating function** is the formal power series $\sum_{k\ge 0} N_k x^k$.

12.2.3. Example. *Subsets.* The poset $\mathbf{2}^n$ is graded, with $r(x) = |x|$ and $r(\mathbf{2}^n) = n$. The kth rank is $\binom{[n]}{k}$. The poset is rank-symmetric and rank-unimodal. Since $N_k(\mathbf{2}^n) = \binom{n}{k}$, the rank generating function is $(1+x)^n$.

Every maximal chain in $\mathbf{2}^n$ has length n (size $n+1$). There are maximal antichains of size 1, but for maximum antichains we will reprove Sperner's Theorem (Theorem 11.2.14) that $w(\mathbf{2}^n) = \binom{n}{\lfloor n/2\rfloor}$. ■

12.2.4. Example. *Divisors of an integer, or multisets.* The divisors of a positive integer N form a poset $D(N)$ under divisibility. It is graded; the rank of an element is the number of primes (with multiplicity) in its factorization. It is rank-symmetric; mapping x to N/x also shows that $D(N)$ is self-dual.

When N is a product of n distinct primes, $D(N) \cong \mathbf{2}^n$; the subsets select the prime factors. When $N = \prod_{i=1}^n p_i^{e_i-1}$, the divisors of N correspond to integer lists $(a_1,\ldots,a_n)$ with $0 \le a_i < e_i$, ordered by $a \le b$ if and only if $a_i \le b_i$ for all i.

Thus $D(N)$ can be viewed as the containment order on multisets from $[n]$ with at most $e_i - 1$ copies of the ith element, where $A \subseteq B$ if each element appears at least as many times in B as in A. Since $e_1,\ldots,e_n$ determine the multiset poset, we denote it by M^e, where $e = (e_1,\ldots,e_n)$.

The multiset description expresses the poset as a product of chains: $M^e \cong \mathbf{e_1}\times\cdots\times\mathbf{e_n}$. Study of M^e arose from divisibility questions, and some early proofs were phrased using divisibility, but the structure is completely captured by the chain-product description, and we usually view the elements as n-tuples.

For $a \in M^e$, we have $r(a) = \sum a_i$. The rank generating function enumerates multisets, indexed by size: $\prod_{i=1}^n (1 + x + \cdots x^{e_i-1})$. When the ith multiplicity is unbounded, we let $e_i = \infty$ and use $1/(1-x)$ as the ith factor. ■

SYMMETRIC CHAIN DECOMPOSITIONS

In light of Dilworth's Theorem, we can determine the width of a poset by exhibiting an antichain and a chain decomposition of the same size. Given a rank-symmetric rank-unimodal poset, we can hope for a special decomposition.

12.2.5. Definition. A chain in a graded poset P is **symmetric** if it has an element of rank $r(P) - k$ whenever it has an element of rank k. A chain is **consecutive** or **skipless** if its elements lie in consecutive ranks. A **symmetric chain decomposition** is a partition into symmetric skipless chains. A poset with a symmetric chain decomposition is a **symmetric chain order**.

A graded poset has the **Sperner property** if a largest-sized rank is a largest antichain. It has the **strong Sperner property** if for all k the union of any k largest ranks form a largest k-family.

12.2.6. Example. Since every chain in a symmetric chain decomposition intersects the middle rank (or both middle ranks), it is a Dilworth decomposition. A symmetric chain order must be rank-symmetric and rank-unimodal. The poset on the left fails, although it is graded and has the Sperner property. The poset on the right is graded but does not have the Sperner property. It has an antichain of size 4, but the maximum rank-size is 3. ∎

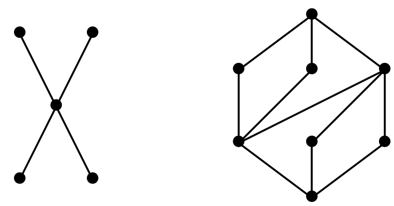

12.2.7. Proposition. Every symmetric chain order satisfies the strong Sperner property.

Proof: In a graded poset P, any k largest ranks form a k-family. Every chain C in a symmetric chain decomposition contributes exactly $\min\{k, |C|\}$ elements to the union of the k largest ranks. No chain C can contribute more than $\min\{k, |C|\}$ elements to a k-family, so these ranks form a maximum k-family. ∎

12.2.8. Theorem. (de Bruijn–Tengbergen–Kruyswijk [1951]) $\mathbf{2}^n$ is a symmetric chain order.

Proof: We use induction on n. When $n = 0$, there is only one element. For $n > 0$, take two copies of a chain decomposition of $\mathbf{2}^{n-1}$. Add $\{n\}$ to each member of each chain in the second copy. This decomposes $\mathbf{2}^n$ into skipless chains, but they are not symmetric and are too many when n is odd. Alter the two copies of a chain C by transferring the top element of the second copy to the top of the first copy. Since the element moved is the union of $\{n\}$ with the previous top element, the result is still skipless. All chains are now symmetric within the full poset. ∎

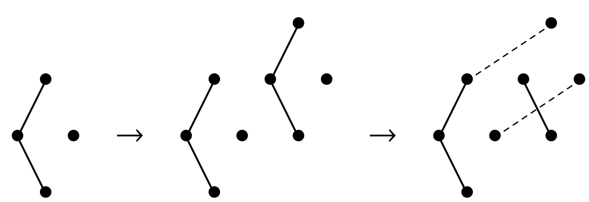

The result was proved more generally for the divisor order, solving a problem posed by the Dutch Wiskundig Genootschap in 1949. Katona observed that the method applies to any product of symmetric chain orders.

12.2.9. Theorem. (Katona [1972a]) Products of symmetric chain orders are symmetric chain orders.

Proof: We need only consider a product of two such orders. If P and Q have symmetric chain decompositions $B_1, \ldots, B_k$ and $D_1, \ldots, D_l$, then the products $B_i \times D_j$

partition $P \times Q$ into "symmetric rectangles". The product $B_i \times D_j$ has elements from ranks k through $r(P \times Q) - k$ for some k, and it has as many elements from rank j as from rank $r(P \times Q) - j$, for each j. Partitioning each such rectangle into symmetric chains provides a symmetric chain decomposition of $P \times Q$.

This reduces the problem to P and Q being chains of sizes $s + 1$ and $t + 1$, with $s \geq t$. The product consists of $\{(i, j)\colon 0 \leq i \leq s \text{ and } 0 \leq j \leq t\}$. Use chains $C_0, \ldots, C_t$, where C_k consists of $(0, k), (1, k), \ldots, (s-k, k), (s-k, k+1), \ldots, (s-k, t)$. This chain is skipless; it is symmetric because the top and bottom ranks sum to $s + t$. If $i + j \leq s$, then (i, j) belongs only to chain C_j. If $i + j > s$, then (i, j) belongs only to chain C_{s-i}. Thus the chains are disjoint, as shown below. Theorem 12.2.8 is the case $t = 1$. ∎

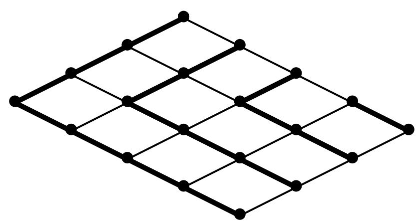

Theorem 12.2.8 is easy but does not explicitly describe the chains. How can we tell whether two given elements lie in the same chain?

12.2.10. Example. *Bracketing decomposition of* $\mathbf{2}^n$. Greene–Kleitman [1976b] and Leeb [unpublished] gave an explicit locally described symmetric chain decomposition of $\mathbf{2}^n$ (see Exercise 12 for multisets); we will see that it is the same as the decomposition in Theorem 12.2.8.

View $S \in \mathbf{2}^n$ as a binary vector x, with $x_i = 1$ if and only if $i \in S$. Encode each 0 as a left bracket and each 1 as a right bracket. Iteratively pair some positions as follows. As long as some unmatched left bracket precedes an unmatched right bracket, there is some closest such pair, say positions l and r with $l < r$. Add (l, r) to the list of paired positions. For any closest unmatched pair, all positions between them are already matched to other positions between them.

The resulting set of position pairs is the **bracketing structure** of x. Having the same bracketing structure is more restrictive than having the same set of matched positions. Below we group the sets in $\mathbf{2}^4$ by their bracketing structure; those in the first column have no paired positions, then three columns have one pair, and finally the two rightmost sets show two ways to pair all four positions.

1234 =))))					
123 =)))(	234 = ()))	134 =)())	124 =))()		
12 =))((	23 = ())(	13 =)()(	14 =)(()	24 = ()()	34 = (())
1 =)(((	2 = ()((	3 = (()(	4 = ((()		
∅ = ((((					

When the process ends, all unmatched rights occur before all unmatched lefts. Let C be the family of subsets of $[n]$ with a fixed bracketing structure. The positions with matched rights correspond to elements of $[n]$ belonging to each set in C. Order C by the number of unmatched rights, starting with the set whose unmatched parentheses are all lefts. This expresses C as a chain in $\mathbf{2}^n$. To move up the chain, change the leftmost unmatched left parenthesis to a right parenthesis, changing a 0 to a 1 in the incidence vector. At the bottom of C, all unmatched positions are left parentheses; at the top they are rights.

When the bracketing structure has j matched pairs, the chain extends from rank j to rank $n-j$ (with $n-2j+1$ members), since every member contains the elements for the j matched right parentheses and omits those for the j matched left parentheses. Hence the chains are skipless and symmetric. Every member of $\mathbf{2}^n$ occurs on exactly one chain, so this is a symmetric chain decomposition. ■

12.2.11. Theorem. The inductive (Theorem 12.2.8) and bracketing (Example 12.2.10) decompositions of $\mathbf{2}^n$ are the same.

Proof: We use induction on n. For $n \le 1$, the poset is a single chain. For $n > 1$, the induction hypothesis states that each chain in the inductive decomposition of $\mathbf{2}^{n-1}$ has a fixed bracketing. It suffices to show that applying the inductive construction to any chain C in the bracketing decomposition of $\mathbf{2}^{n-1}$ yields chains in the bracketing decomposition of $\mathbf{2}^n$.

From C define a chain C' by adding $\{n\}$ to each element of C. In position n, each vector for C gains a left parenthesis, while the corresponding vector for C' gains a right parenthesis. The bracketing structure for C does not change by the added left parenthesis. The top member of C' also has this bracketing structure, since its unmatched positions all have right parentheses, so adding a right at the end does not create another match. Thus this element moves from the top of C' to the top of C in the bracketing decomposition of $\mathbf{2}^n$.

The remaining members of C' all had a rightmost unmatched left parenthesis in the same position. Thus the new right parenthesis matches to the same position to complete the bracketing structure for each member of C' below the top. Hence they lie on the same chain in the bracketing decomposition of $\mathbf{2}^n$.

Furthermore, when the bracketing structure of C consists of j pairs, C starts with $n-2j$ members. It gains one to reach $n+1-2j$ members (still with j pairs in its bracketing structure), while C' ends with $j+1$ pairs and $n-1-2j$ members. Hence these are full chains in the bracketing decomposition of $\mathbf{2}^n$.

We have proved that every chain of size at least 2 in the bracketing decomposition of $\mathbf{2}^{n-1}$ becomes two chains in the decomposition of $\mathbf{2}^n$ when the induction step is applied. By induction on n, the two decompositions are thus the same. ■

The Dilworth decomposition we have described for $\mathbf{2}^n$ arises in many ways. We present a third description, and Exercise 19 has yet a fourth.

12.2.12.* Example. (Aigner [1973]) We use a greedy lexicographic rule. List the k-sets in lexicographic order, from $[k]$ to $[n]-[n-k]$. Match each k-set A in turn to the $(k+1)$-set that is earliest in lexicographic order on level $k+1$ among the $(k+1)$-sets that remain unmatched and contain A. For example, in $\mathbf{2}^4$ we match $\varnothing$ to 1, then $1 \to 12$, $2 \to 23$, $3 \to 13$, and $4 \to 14$, then $12 \to 123$, $13 \to 134$, $14 \to 124$, $23 \to 234$, with none available for 24 and 34, and finally $123 \to 1234$.

The resulting chains (shown below for $n = 4$) are the same as those in Example 12.2.10 (see Exercise 20). It is not even obvious that this produces symmetric chains, and yet we obtain the familiar bracketing decomposition. The construction also extends to multisets. ■

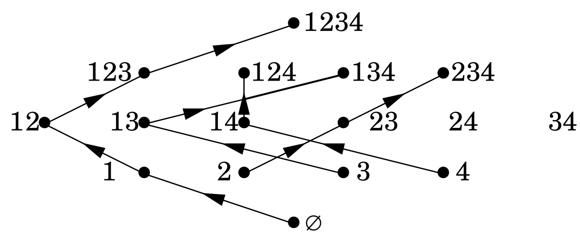

For many classical rank-symmetric rank-unimodal posets, the existence of symmetric chain decompositions remains unknown.

12.2.13. Example. $L(m, n)$*: Bounded integer partitions.* The elements of $L(m, n)$ are the integer lists $(a_1, \ldots, a_m)$ such that $0 \le a_1 \le \cdots < a_m \le n$. These correspond to Ferrers diagrams contained in an m by n rectangle. The order relation puts $a \le b$ if and only if $a_i \le b_i$ for all i. Hence $L(m, n)$ is a subposet of $(\mathbf{n+1})^m$.

The complement of a Ferrers diagram in a rectangle fits in that rectangle, so $L(m, n)$ is rank-symmetric (actually, self-dual). The rank generating function (with formal variable q for algebraic reasons), is $\prod_{j=0}^{m-1} \frac{q^{n+m-j}-1}{q^{m-j}-1}$ (Exercise 30). Rank-unimodality is difficult; algebraic proofs began with Sylvester (see Proctor [1982]). O'Hara [1990] found an intricate combinatorial proof, presented also in Zeilberger [1989]. Stanley [1982] observed that $L(m, n)$ has the Sperner property, using results of Griggs [1977].

Stanley conjectured that $L(m, n)$ is a symmetric chain order. This is easy for $m \le 2$ (see Exercises 5–8) and is known for $m \le 4$ (Riess [1978], Lindström [1980], West [1980]). Solutions for $L(m, n)$ with $m = 5$ and n odd have been rumored. ■

Finally, we present an application of the bracketing decomposition. We want to count antichains in $\mathbf{2}^n$; this is known as **Dedekind's Problem**. By Remark 12.1.14, antichains correspond to down-sets. There are $2^{\binom{n}{\lceil n/2\rceil}}$ down-sets whose maximal elements all have size $\lceil n/2 \rceil$; this yields a lower bound. For the upper bound, we encode down-sets as **monotone Boolean functions**, which order-preserving functions from $\mathbf{2}^n$ to $\{0, 1\}$.

12.2.14. Theorem. (Hansel [1966]) The number of monotone Boolean functions is at most $3^{\binom{n}{\lceil n/2\rceil}}$.

Proof: We construct a monotone Boolean function f by specifying its values in order. Consider chains in the bracketing decomposition in increasing order of size. For each new chain, the values of f on some of it are already forced. If x contains a set already assigned 1, then also $f(x) = 1$. Similarly, a subset of a set assigned 0 must be assigned 0. Since there are $\binom{n}{\lceil n/2\rceil}$ chains, it suffices to show that always there are at most three ways to extend f to the next chain C. To do this, we show that C has at most 2 elements with undetermined labels. When these are x and y with $x > y$, the options $(f(y), f(x))$ will be only $(0, 0)$, $(0, 1)$, and $(1, 1)$.

Consider sets x and y in C, with $x > y$ and $|x - y| > 1$. To show that f is known for x or y, it suffices to construct z between them in $\mathbf{2}^n$ such that $f(z)$ is known. Then $f(z) = 1$ implies $f(x) = 1$, and $f(z) = 0$ implies $f(y) = 0$.

If C has size k, then its bracketing has $n - k + 1$ matched positions. Shorter chains have more matched positions. The unmatched positions for x and y have right parentheses followed by left parentheses, with at least two more rights in x than in y. Construct z from x by changing the next-to-last unmatched right parenthesis in x to a left parenthesis.

Now $y < z < x$. In the bracketing of z, the positions corresponding to the last two unmatched right parentheses in x match, since the positions between them are all matched. The other matches are unchanged, so z has more matched positions and appears on a shorter chain, as desired. ∎

$$x = ---)----)-----$$
$$z = ---(----)-----$$
$$y = ---(----(-----$$

By a more detailed argument along the same lines, showing in effect that almost always the new chain contains at most one undetermined element, Kleitman–Markowsky [1975] improved this upper bound to $2^{\binom{n}{\lceil n/2\rceil}(1+O(\log n/n))}$. Korshunov [1981] obtained a more detailed asymptotic formula, which for even n is

$$2^{\binom{n}{n/2}} e^{\binom{n}{n/2-1}(2^{-n/2}+n^2 2^{-n-5}-n2^{-n-4})}.$$

Exact values for $n \le 7$ are 3, 6, 20, 168, 7581, 7828354, and 2,414,682,040,998, with Korshunov's formula giving 7996118 for $n = 6$.

LYM AND SPERNER PROPERTIES

A short direct proof of Sperner's Theorem counts the maximal chains in $\mathbf{2}^n$ by where they hit an antichain. The proof identifies a property for graded posets that implies the Sperner property and much more.

12.2.15. Theorem. (Sperner [1928]) In $\mathbf{2}^n$, the elements of rank $\lfloor n/2 \rfloor$ form a maximum antichain.

Proof: (Lubell [1966]) A maximal chain has at most one element in an antichain F. We count the chains according to which element of F (if any) they contain. There are $n!$ maximal chains, since they are specified by acquiring the elements of $[n]$ in some order. To pass through x, they must acquire the elements of x first, so exactly $|x|!(n-|x|)!$ maximal chains pass through x. Thus $\sum_{x\in F} |x|!(n-|x|)! \le n!$.

Dividing both sides by $n!$ yields $\sum_{x\in F} 1/N_{r(x)} \le 1$. Replacing $N_{r(x)}$ with the largest binomial coefficient preserves the inequality, and then clearing the fraction yields $|F| \le \binom{n}{\lceil n/2\rceil}$. ∎

The inequality $\sum_{x\in F} N_{r(x)}^{-1} \le 1$ was discovered independently for antichains in $\mathbf{2}^n$ by Lubell [1966], Yamamoto [1954], and Meshalkin [1963].

12.2.16. Definition. For a family F in a graded poset, the inequality $\sum_{x\in F} N_{r(x)}^{-1} \le 1$ is the **LYM inequality**. A graded poset satisfies the **LYM property** if its antichains all satisfy the LYM inequality. Such a poset is an **LYM order**.

By the argument of Theorem 12.2.15, the LYM property implies the Sperner property. An analogue of this argument using circular arrangements yields the strong Sperner property for $\mathbf{2}^n$ (Exercise 27). We will see later that the LYM property implies the strong Sperner property.

We will also see that the argument in Theorem 12.2.15 generalizes; it needs only a list of maximal chains such that, in each rank, each element appears equally often. We introduce a name for such a list of maximal chains in order to use the argument in more generality.

12.2.17. Definition. A nonempty list of maximal chains in a graded poset P is a **regular covering** of P if, for each rank P_k, each element of P_k lies in exactly the fraction $1/|P_k|$ of these chains.

12.2.18. Example. A regular covering of P is a list; chains may be used repeatedly. The list may omit some maximal chains and use some more than others. In $\mathbf{3}\times\mathbf{4}$, shown below, the number of maximal chains containing a specified element of a middle rank is 1, 6, or 3. The full set of maximal chains is not a regular covering. On the other hand, there is a regular covering with 6 chains consisting of two copies of the "outer" indicated chains plus one copy of each "inner" chain. We will prove later that every (finite) product of chains has a regular covering. ■

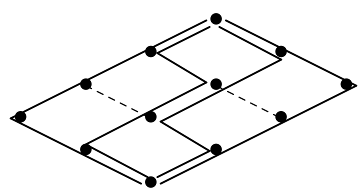

LYM orders are characterized by the existence of a regular covering. To prove this, we use a third and equally important property.

12.2.19. Definition. (Graham–Harper [1969]) A graded poset P has the **normalized matching property** if $|A^*|/N_{k+1} \ge |A|/N_k$ for all k and all $A \subseteq P_k$, where A^* denotes $U[A] \cap P_{k+1}$. The set A^* is called the **shade** of A or its "shadow at the rank above".

12.2.20. Example. $\mathbf{2}^n$ *has the normalized matching property.* For $A \subset \binom{[n]}{k}$, each element of A extends to $n-k$ elements of A^*. Since each element in A^* can lose an element in $k+1$ ways, $(k+1)|A^*| \ge (n-k)|A|$. Dividing both sides by $(n-k)\binom{n}{k}$ yields $|A|/N_k \le |A^*|/N_{k+1}$. This was the essence of the argument we used to prove Sperner's Theorem (Theorem 11.2.14). ■

The subgraph of the diagram induced by $P_k \cup P_{k+1}$ is a P_k, P_{k+1}-bigraph with $N(A) = A^*$. When $N_{k+1} \ge N_k$, normalized matching is stronger than Hall's condition for a matching that covers P_k (Section 6.1). We use Hall's Theorem to prove equivalence for these properties.

12.2.21. Theorem. (Kleitman [1974]) For a graded poset P, the following statements are equivalent:
(A) P has a regular covering.
(B) P has the LYM property.
(C) P has the normalized matching property.

Proof: A $\Rightarrow$ B. As in the proof of Theorem 12.2.15, counting the chains in a regular covering $\mathbf{C}$ that are hit by the elements of an antichain F yields $\sum_{x \in F} |\mathbf{C}|/N_{r(x)} \le |\mathbf{C}|$ and thus the LYM inequality.

B $\Rightarrow$ C. If $A \subseteq P_k$, then $A \cup (P_{k+1} - A^*)$ is an antichain in P. The LYM inequality yields $|A|/N_k + |P_{k+1} - A^*|/N_{k+1} \le 1$. Since $|P_{k+1}| = N_{k+1}$, this becomes $|A|/N_k \le |A^*|/N_{k+1}$.

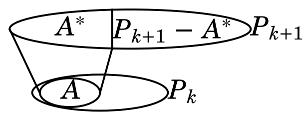

C $\Rightarrow$ A. Letting $M = \prod N_k(P)$, define the **blowup poset** P' with $M/N_{r(x)}$ copies of x for each $x \in P$. A copy of x is less than a copy of y in P' if and only if $x < y$ in P. Use the normalized matching property in P to find perfect matchings joining adjacent ranks in P'; such matchings combine to partition P' into disjoint maximal chains, and these chains collapse to form a regular covering of P with x appearing $M/N_{r(x)}$ times.

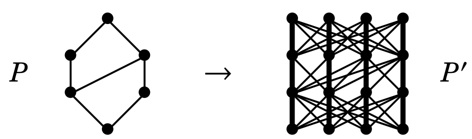

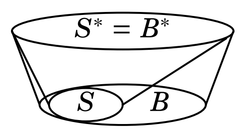

By Hall's Theorem, it suffices to show for $S \subset P'_k$ that $|S^*| \ge |S|$, where S^* is the shade of S (in P'_{k+1}). Let B be the subset of P'_k consisting of all copies of each element of P_k that has at least one copy in S. Now $S^* = B^*$ and $S \subseteq B$, so it suffices to show that $|B^*| \ge |B|$.

Let A be the sets of elements in P that have copies in B. Note that in P' the set B^* consists of all copies of all elements of P_k covering elements of A; that is, all copies of elements of A^*. Note that $|B| = |A|\, M/N_k$ and $|B^*| = |A^*|\, M/N_{k+1}$, since all copies of an element appear when any copies appear. The normalized matching property in P now yields $|B^*| \ge |B|$. ■

12.2.22.* Remark. (1) The LYM inequality and regular coverings are self-dual, so normalized matching is equivalent to requiring $|D[A] \cap P_{k-1}| \ge |A|\, N_{k-1}/N_k$ when $A \subseteq P_k$ (Exercise 28 requests a direct proof without duality).

(2) Theorem 12.2.21 guarantees short proofs. Having the LYM property can be proved by giving a regular covering. Failing the LYM property can be proved by giving a violation of the normalized matching property (Exercises 25–26).

(3) Since the number of chains is divisible by each rank-size, every regular covering has at least lcm $\{N_k\}$ chains, achieved by blowing up x to lcm $\{N_k\}/N_{r(x)}$ copies (Exercise 4). However, every LYM poset has a regular covering using at most $|P| - r(P)$ *distinct* chains, and this is best possible (Exercise 29).

(4) Graham–Harper [1969] found an efficient method to find a regular covering. We seek nonnegative integer weights on the edges from P_k to P_{k+1} such that

the total weight on edges leaving each element of P_k is N_{k+1} and the total weight on edges entering each element of P_{k+1} is N_k. A solution forms a regular covering for the subposet $P_k \cup P_{k+1}$. The weights give the relative usage of the cover relations in the full regular covering. Feasibility can be tested by adding a source and sink, modeling the weights with multiedges, and using Menger's Theorem; thus we can test the LYM property in polynomial time.

(5) When the number of elements an element x covers depends only on the rank of x, the set of all maximal chains is a regular covering (Baker [1969]). ■

We have seen that the LYM property implies the Sperner property. The LYM property (with its equivalence to regular covering) also implies a more general statement that implies the strong Sperner property.

12.2.23. Theorem. (Greene–Kleitman [1978]) Let $\lambda\colon P \to \mathbb{R}$ be a weight function defined on the elements of P. If P has the LYM property, then for every subset $G \subset P$ and every regular covering $\mathbf{C}$ of P,

$$\sum_{x\in G} \frac{\lambda_x}{N_{r(x)}} \le \max_{C\in\mathbf{C}} \sum_{y\in C\cap G} \lambda_y \quad .$$

Proof: We interpret the inequality probabilistically. Choose a chain from $\mathbf{C}$ uniformly at random, and define a random variable $X = \sum_{y\in C\cap G} \lambda_y$. The expectation of X is $\sum_{x\in G} \lambda_x \mathbb{P}(x \in C)$. Since $\mathbf{C}$ is a regular covering, $\mathbb{P}(x \in C) = 1/N_{r(x)}$. Thus the left side of the desired inequality is the expected value of X, while the right side is its maximum value. ■

Suitable choices of λ yield various applications. First let $\lambda_x = N_{r(x)}$.

12.2.24. Corollary. If P is an LYM order with regular covering $\mathbf{C}$ and G is any subset of P, then $|G| \le \max_{C\in\mathbf{C}} \sum_{y\in C\cap G} N_{r(y)}$. ■

Bounds on $|G|$ follow from Corollary 12.2.24 for various chain conditions on $|G|$. Erdős [1945b] proved the next corollary for $\mathbf{2}^n$ by generalizing the argument of Sperner [1928] for the Sperner property.

12.2.25. Corollary. (Erdős [1945b]) The LYM property implies the strong Sperner property.

Proof: If G is a k-family, then Corollary 12.2.24 limits its size to the sum of the k largest rank sizes. ■

Now it is natural to ask whether the LYM property is so strong that it also implies the structural "holy grail" of symmetric chain decomposition. The answer is "yes" when the obvious necessary conditions hold.

12.2.26. Theorem. (Anderson [1976], Griggs [1977]) Every rank-unimodal rank-symmetric LYM poset is a symmetric chain order.

Proof: We use induction on the height of the poset P. For height 1, there is nothing to prove. For even height, the two middle levels have equal size, and normalized matching between them reduces to Hall's Condition. Collapse the two levels to one along the resulting matching. Since normalized matching considers only consecutive levels, the smaller poset P' is still an LYM order, and it also has symmetric and unimodal rank sizes. The guaranteed chain decomposition for P' expands into a symmetric chain decomposition of P using the central matching.

For odd height, let the middle rank be P_{k+1}. If we can match P_k to P_{k+2} through P_{k+1}, then the unused elements of P_{k+1} will become singleton chains. After discarding them, the three levels collapse to one level along the matched edges to obtain a smaller LYM order P'. After obtaining a symmetric chain decomposition of P', expand each element of the middle level into a 3-element chain as the middle of a skipless symmetric chain in P. Together with the discarded singletons, these chains form a symmetric chain decomposition of P.

The cover relations down from P_{k+2} and up from P_k yield two families $\mathbf{A}$ and $\mathbf{B}$ of subsets of P_{k+1}, each with N_k sets. The two desired matchings exist when there is a common system of distinct representatives for $\mathbf{A}$ and $\mathbf{B}$. For subsets $I, J \subseteq [N_k]$, let $A' = \cup_{i\in I} A_i$ and $B' = \cup_{j\in J} B_j$. By the Ford–Fulkerson Condition (Theorem 7.2.15), it suffices to show that $|A' \cap B'| \geq |I| + |J| - N_k$. Because the dual of an LYM order is an LYM order, we have both $|A'| \geq |I|\,(N_{k+1}/N_k)$ and $|B'| \geq |J|\,(N_{k+1}/N_k)$. Thus

$$\begin{aligned}|A' \cap B'| &= |A'| + |B'| - |A' \cup B'| \geq (|I| + |J|)(N_{k+1}/N_k) - N_{k+1} \\ &= (|I| + |J| - N_k)(N_{k+1}/N_k) \geq |I| + |J| - N_k.\end{aligned}$$ ∎

The LYM property has many consequences. Unfortunately, many graded posets of interest are not LYM orders (see Exercises 24–26). For these, one can still ask whether consequences such as symmetric chain decomposition hold. For the poset $L(m, n)$ (Example 12.2.13), this remains a fascinating open problem.

PRODUCTS OF LYM ORDERS (optional)

So far, our only criteria for the LYM property are the three equivalent defining properties. We have not yet proved that chain-products are symmetric chain orders. We consider the behavior of the LYM property under products. First note the behavior of the rank sizes.

12.2.27. Remark. *Rank of poset products.* If P and Q are graded, then their product is also graded and has rank function given by $r_{P\times Q}(x, y) = r_P(x) + r_Q(y)$. Thus $N_k(P \times Q) = \sum_i N_i(P) N_{k-i}(Q)$. ∎

12.2.28. Example. *The product of LYM posets need not be an LYM poset.* In the example below, the set $\{a1, b1\}$ occupies $\frac{2}{3}$ of its rank, but its shadow at the rank above is only $\frac{1}{2}$ of that rank. ∎

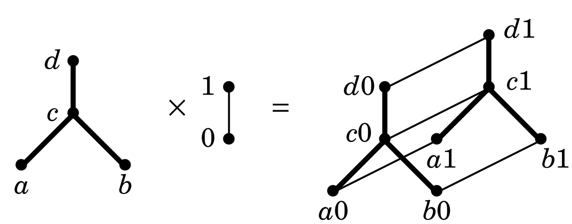

12.2.29. Remark. *Necessary condition for LYM property in a special product.* Let $Q = P \times \mathbf{2}$, where P is an LYM poset. Let $A = \{(x, 1)\colon x \in P_{k-1}\}$. The set A lies in rank k in Q. Note that $A^* \subseteq \{(y, 1)\colon y \in P_k\}$.

Write N_j for $N_j(P)$. Since $\mathbf{2}$ has one element at each rank, Q has $N_{j-1} + N_j$ elements at rank j. Normalized matching thus requires

$$\frac{|A^*|}{N_k + N_{k+1}} \ge \frac{|A|}{N_{k-1} + N_k}.$$

Since $|A| = N_{k-1}$ and $|A^*| = N_k$, we need $N_k^2 \ge N_{k-1}N_{k+1}$. This condition on the sequence of rank sizes is necessary for the product with $\mathbf{2}$ to be an LYM order. We will show that it is sufficient for products in general. ■

12.2.30. Definition. A sequence $\langle N \rangle$ is **log-concave** if $N_k^2 \ge N_{k-1}N_{k+1}$ for all k.

The condition for log-concavity is that the logarithms of the terms form a concave sequence. Every concave sequence is also log-concave.

12.2.31. Theorem. (Harper [1974], Hsieh–Kleitman [1973]) The family of LYM orders with log-concave rank sizes is closed under taking products.

Proof: We have $r(P_1 \times P_2) = r(P_1) + r(P_2)$ and $N_k(P_1 \times P_2) = \sum_{i=0}^{k} N_i(P_1)N_{k-i}(P_2)$. Log-concavity of the convolution of two log-concave sequences is an exercise in algebraic manipulation (Exercise 33).

To prove the LYM inequality for $P_1 \times P_2$, we generalize Theorem 12.2.23. There we picked a random chain from a regular covering. Here, we pick a random pair C_1, C_2 from regular coverings $\mathbf{C}_1$ and $\mathbf{C}_2$ of P_1 and P_2. This produces a random rectangle in the product.

Given any subset G of $P_1 \times P_2$ and weights λ_x for $x \in P$, let $X = \sum_{x \in G \cap (C_1 \times C_2)} \lambda_x$. The expectation is $\sum_{x \in G} \lambda_x \mathbb{P}(x \in C_1 \times C_2)$. Since $\mathbf{C}_1$ and $\mathbf{C}_2$ are regular coverings, $\mathbb{P}(x \in C_1 \times C_2) = N_{r(x_1)}(P_1)^{-1} N_{r(x_2)}(P_2)^{-1}$, where $x = (x_1, x_2)$. Thus the left side of the inequality below is the expectation of X, while the right side is its maximum.

$$\sum_{x \in G} \frac{\lambda_x}{N_{r(x_1)}(P_1)N_{r(x_2)}(P_2)} \le \max_{C_i \in \mathbf{C}_i} \sum_{y \in G \cap C_1 \times C_2} \lambda_y$$

Setting $\lambda_x = N_{r(x_1)}(P_1)N_{r(x_2)}(P_2)/N_x$ yields

$$\sum_{x \in G} \frac{1}{N_{r(x)}} \le \max_{C_i \in \mathbf{C}_i} \sum_{y \in G \cap C_1 \times C_2} \frac{N_{r(y_1)}(P_1)N_{r(y_2)}(P_2)}{N_{r(y)}}.$$

To obtain the LYM inequality, it suffices to show that the right side of this inequality is at most 1 when G is an antichain.

If G lies entirely in the kth rank of $P_1 \times P_2$, then we compute a bound over rank k for arbitrary $C_1 \times C_2$. The elements of rank k in $P_1 \times P_2$ that lie in $C_1 \times C_2$ have the form (y_1, y_2), where y_1 is the element of C_1 with rank i in P_1 and y_2 is the element of C_2 with rank $k - i$ in P_2. There is at most one such element for each $i \in \{0, \ldots, k\}$. Thus

$$\sum_{y \in G \cap C_1 \times C_2} \frac{N_{r(y_1)}(P_1)N_{r(y_2)}(P_2)}{N_{r(y)}} \le \sum_{i=0}^{k} \frac{N_i(P_1)N_{k-i}(P_2)}{N_k(P_1 \times P_2)} = 1.$$

Suppose that G has elements from more than one rank in $P_1 \times P_2$. We show that pushing those in the lowest occupied rank upward toward the highest occupied rank cannot decrease the specified sum. This suffices, since we have shown that the sum is bounded by 1 when G lies in a single rank.

Let T be a maximal set of "consecutive" elements of G at rank k in $C_1 \times C_2$; "consecutive" means $T = \{(x_i, y_{k-i})\colon r \le i \le s\}$ for some $0 \le r \le s$, where x_i is the element of rank i in C_1 and y_j is the element of rank j in C_2. Let T^* be the set of elements covering members of T in $C_1 \times C_2$; this is another consecutive set. Because T is a maximal consecutive set, replacing all of T with T^* does not violate the antichain condition.

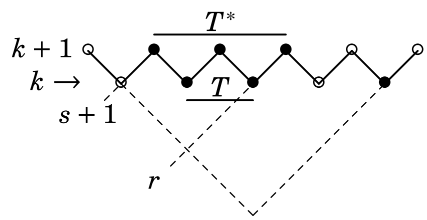

Let $g_k(r,s) = \sum_{i=r}^{s} N_i(P_1)N_{k-i}(P_2)$. The contribution of T to the original sum is $\frac{g_k(r,s)}{g_k(0,\infty)}$. The contribution of T^* to the new sum is $\frac{g_{k+1}(r,s+1)}{g_{k+1}(0,\infty)}$. We need only show that $\frac{g_{k+1}(r,s+1)}{g_{k+1}(0,\infty)} \ge \frac{g_k(r,s)}{g_k(0,\infty)}$. Since this is equivalent to $\frac{g_{k+1}(r,s+1)}{g_k(r,s)} \ge \frac{g_{k+1}(0,\infty)}{g_k(0,\infty)}$, it suffices to show that $\frac{g_{k+1}(r,s+1)}{g_k(r,s)}$ is increasing in r and decreasing in s. See Exercise 34. ∎

12.2.32. Corollary. M^e is an LYM order.

Proof: M^e is a product of chains, each of which is an LYM order with log-concave rank sizes. ∎

Algebraic techniques can yield log-concavity of sequences. When the rank generating function is known, the following result may apply (see Stanley [1989]).

12.2.33. Theorem. If all roots of a polynomial with real coefficients are real, then the sequence of coefficients is log-concave. ∎

12.2.34. Example. The rank generating function for $\mathbf{2}^n$ is $(1+x)^n$, so Theorem 12.2.33 applies. However, for M^e it is $\prod_{i=1}^{n}(1 + x + \cdots + x^{e_i - 1})$, which has complex roots when $\max e_i > 2$. Log-concavity follows because the convolution of log-concave sequences is log-concave. ∎

EXERCISES 12.2

12.2.1. (−) Use Sperner's Theorem and Dilworth's Theorem to prove (weaker than Theorem 12.2.14) that the number of antichains in $\mathbf{2}^n$ is at most $(n+1)^{\binom{n}{\lceil n/2 \rceil}}$. (Gilbert [1954])

12.2.2. (−) Let P be a rank-symmetric LYM order of rank n. Let S be a down-set in P. Prove that the elements of S have average rank at most $n/2$.

12.2.3. (−) Show that the sequence of rank sizes of a product poset may be log-concave even though the sequences for the factors are not both log-concave.

12.2.4. (−) Prove that the minimum number of chains in a regular covering of an LYM order is lcm N_k. (Hint: Blow up each element x to (lcm N_k)/$N_{r(x)}$ copies and use the argument of Theorem 12.2.21.) (West–Harper–Daykin [1983])

12.2.5. (−) Prove that $L(2, n)$ (Example 12.2.13) is a symmetric chain order.

12.2.6. (−) An **automorphism** of a poset is an order-preserving permutation of the elements. Let P be a graded poset such that whenever x and y have the same rank in P, some automorphism maps x to y. Prove that P is an LYM order.

12.2.7. Prove $w(\mathbf{n} \times \mathbf{n} \times \mathbf{n}) = \lfloor (3n^2 + 1)/4 \rfloor$.

12.2.8. Construct a symmetric chain decomposition of $L(3, n)$ for odd n.

12.2.9. Let P be a finite poset whose diagram is connected.

(a) Show that P may have a maximal chain that is not a longest chain even when every element in P lies in a longest chain of P.

(b) Suppose that, whenever y covers x in P, some longest chain contains both x and y. Prove that P is a graded poset. (Stanley [1971, p. 19–20])

12.2.10. *k-families in chain products.*

(a) For $i, j, k \in \mathbb{N}$, find a formula without summations for the maximum size of a k-family in the chain-product $\mathbf{i} \times \mathbf{j}$. (Hint: Consider three cases for $\{i, j, k\}$.)

(b) Consider three sets of parallel lines in the plane, forming equilateral triangles. Within the sets, the lines need not be equally spaced. Let the sets have sizes r, s, t. Determine the maximum number of points that occur as the intersection of three lines, one from each set. (Matsko–West–Wetzel [2001])

12.2.11. (◇) A **semiantichain** in $P \times Q$ is a subset S having $(u, v) < (u', v')$ for two elements of S only if $u < u'$ and $v < v'$. A product poset is the **2-part Sperner** when some single rank is a largest semiantichain.

(a) Prove that the product of two symmetric chain orders is 2-part Sperner.

(b) Use part (a) to show that for any 2-coloring of the elements of $[n]$, the largest subposet of $\mathbf{2}^n$ having no related pair with monochromatic difference consists of the middle rank. (Kleitman [1965], Katona [1966])

12.2.12. *Bracketing decomposition of M^e.* View $a \in M^e$ as a list with $0 \le a_i < e_i$. Dedicate $e_i - 1$ positions for the ith coordinate. In these positions, put a_i right parentheses followed by $e_i - 1 - a_i$ left parentheses. Match parentheses as in Example 12.2.10. The chain containing a consists of all elements with the same bracketing as a. (Below is a chain from $M^{(5,4)}$ with one matched pair; it is the second chain in the figure for Theorem 12.2.9. Prove that these chains form a symmetric chain decomposition of M^e and that it is the same as the decomposition in Theorem 12.2.9. (Greene–Kleitman [1976b], Leeb [unpub.])

$$\begin{array}{ccccccccccc} ((((\ )((& &)(((\ )((& &))((\ )((& &)))(\ )((& &)))(\ ))(& &)))(\ ))) \\ (0,1) & < & (1,1) & < & (2,1) & < \ (& (3,1) & < & (3,2) & < & (3,3) \end{array}$$

12.2.13. A finite set of integers is *balanced* if the numbers of even and odd elements differ by at most 1. Let F be a family of subsets of $[n]$ such that whenever $A, B \in F$ with $A \subseteq B$, the set $B - A$ is not balanced. Prove that $|F| \le \binom{n}{\lceil n/2 \rceil}$ and that this bound is best possible. (Greene–Kleitman [1976b])

12.2.14. (◇) A **skew chain order** is a poset having a rank function in which every minimal element has rank 0 and a decomposition into skipless chains starting at rank 0. Describe (with proof) the largest semiantichain (see Exercise 12.2.11) in the product of two skew chain orders. Let P_k be the inclusion order on the set of intervals in $\mathbb{R}$ with endpoints in $[k]$. Give a geometric description of the maximum semiantichain in $P_m \times P_n$. (West–Kleitman [1979])

12.2.15. Use a chain decomposition of $\mathbf{2}^n$ to construct a spanning subgraph of the hypercube graph Q_n that has diameter n and has only $2^n + \binom{n}{\lceil n/2 \rceil} - 2$ edges. (Comment: This subgraph has the same diameter as Q_n while keeping only a vanishing fraction of the edges.) (Graham–Harary [1992])

12.2.16. ($\diamond$) A **universal subset list** on an alphabet S is a word having every subset of S as a consecutive substring. For example, 1231 is such a list on [3], and 123421341 is such a list on [4]. Listing all 2^n subsets successively yields a universal subset list on $[n]$ with length $n2^{n-1}$, since the average size is $n/2$.

(a) For even n, use symmetric chain decompositions of two copies of $\mathbf{2}^{n/2}$ to construct a universal subset list on $[n]$ with length asymptotically at most $(4/\pi)2^n$. (Hint: Use Stirling's formula to approximate $\binom{k}{\lfloor k/2 \rfloor}$.)

(b) Show that top element on chains in a symmetric chain decomposition of $\mathbf{2}^{[k]}$ has average size $k/2 + O(k^{1/2})$. Use this to reduce the bound in (a) by a factor of 2. (Lipski [1978])

(c) Prove that these upper bounds are within a factor of $c\sqrt{n}$ of being optimal.

12.2.17. ($\diamond$) *The Littlewood–Offord Problem.* Let $a_1, \dots, a_n$ be vectors in $\mathbb{R}^d$, each having length at least 1. Let $R_1, \dots, R_k$ be regions in $\mathbb{R}^d$, each having diameter less than 1 (i.e., contained in a sphere of diameter 1), and let R be their union. Let $d_k(n) = \sum_{i=r}^{s} \binom{n}{i}$, where $r = \lfloor (n-k+1)/2 \rfloor$ and $s = \lfloor (n+k-1)/2 \rfloor$ (this counts the k middle ranks in $\mathbf{2}^n$).

(a) Prove that $d_k(n) = d_{k+1}(n-1) + d_{k-1}(n-1)$.

(b) Prove that the number of 0,1-vectors x such that $\sum x_i a_i \in R$ is at most $d_k(n)$. (Hint: To apply part (a) in an inductive proof, one must group these vectors x into sets corresponding to two problems of the same type with $n-1$ vectors, one having $k+1$ regions and one having $k-1$ regions.) (Kleitman [1970])

12.2.18. Let G be an X,Y-bigraph with $X = \{x_1, \dots, x_m\}$ and $Y = \{y_1, \dots, y_n\}$. Define the *greedy matching* of X into Y as follows: having processed each x_i with $i < r$, match x_r to its least-indexed available neighbor in Y, if any is available. Prove that the greedy matchings of X into Y and Y into X are the same.

12.2.19. ($\diamond$) For $X \subseteq [n]$, let $X = \{x_1, \dots, x_k\}$ with $x_1 < \cdots < x_k$, and set $x_0 = 0$. Let t be the largest nonnegative index i minimizing $x_i - 2i$. If $x_t < n$, then form $f(X)$ from X by adding the element $1 + x_t$, but leave $f(X)$ undefined if $x_t = n$. Prove that the pairs $(X, f(X))$ form the chains in the bracketing decomposition of $\mathbf{2}^n$. (For example, applications of f in $\mathbf{2}^4$ starting with $\{3\}$ yield $\{1,3\}$, then $\{1,3,4\}$, then nothing. Other examples are $f(\{1,3,4,7\}) = \{1,3,4,5,7\}$ with $t = 3$ and $f(\{3,5,7,9\}) = \{1,3,5,7,9\}$ with $t = 0$.) (White–Williamson [1977])

12.2.20. Prove that Aigner's lexicographically-generated chains (Example 12.2.12) are the same as the chains in Exercise 12.2.19.

12.2.21. ($\diamond$) *Orthogonal chain partitions.* (Kleitman–Shearer [1979])
Elements are *related* in a chain partition if they lie on the same chain. Two partitions are **orthogonal** if no two elements are related in both partitions.

(a) Define the **antibracketing decomposition** of $\mathbf{2}^n$ from the bracketing decomposition by replacing each set with its complement and reversing the chains. Prove that a slight modification of the antibracketing decomposition yields a chain partition orthogonal to the bracketing decomposition, when $n \ge 4$.

(b) Construct a pair of orthogonal Dilworth decompositions for $\mathbf{2}^2$ and $\mathbf{2}^3$, and construct three pairwise orthogonal Dilworth decompositions for $\mathbf{2}^4$. Prove that $\mathbf{2}^n$ has at most $\lceil (n+1)/2 \rceil$ pairwise orthogonal Dilworth decompositions.

(c) Prove that there are $n!$ Dilworth decompositions of $\mathbf{2}^n$ such that no two elements appear on the same chain in more than the fraction $\lceil (n+1)/2 \rceil^{-1}$ of the decompositions.

12.2.22. (+) Prove that the number of symmetric chain decompositions of $\mathbf{2}^n$ is at least $n \cdot 3!4! \cdots \lceil (n+1)/2 \rceil !$. (Anderson [1985])

12.2.23. ($\diamond$) Let G be an X,Y-bigraph with $|Y| \ge |X|$ and no isolated vertices. The *deficiency* $\phi(A)$ of a set $A \subseteq V(G)$ is $|A| - |N(A)|$. Say that G has the **strong Hall property** if $\phi(A) + \phi(B) \le |Y| - |X|$ when $A, B \subseteq Y$ with $|A| + |B| \le |Y|$. Let H be a graph formed from $G + G$ by adding a matching joining the copies of Y.

(a) Prove that if G satisfies the strong Hall property, then H has a perfect matching.

(b) Let P be a self-dual rank-symmetric rank-unimodal poset. Prove that if the bipartite graphs joining consecutive levels all have the strong Hall property, then P has a symmetric chain decomposition. (Lu–Wang–Wong [1998])

12.2.24. The **Weak Order** W_n on the permutations of $[n]$ is defined by letting σ cover τ if σ is obtained from τ by transposing two consecutive elements to make them an inversion. The poset is graded: $r(\sigma)$ is the number of inversions in σ, so the rank generating function is $1(1+x)(1+x+x^2)\cdots(1+x+\cdots+x^{n-1})$, with unimodal coefficients (Example 12.2.34).

(a) Show that W_n is a symmetric chain order for $n \le 4$. (Comment: Also W_5 is a symmetric chain order by using Exercise 12.2.23, but for $n > 5$ the answer is unknown.)

(b) Show that W_4 is not an LYM order by considering the permutation 2143. Generalize the argument to prove that W_{2m} is not an LYM order for $m \ge 2$.

12.2.25. Prove that $L(m,n)$ is an LYM order if and only if $m = n = 3$ or $\min\{m,n\} \le 2$. (Hint: Find the lowest violation of the normalized matching property in $L(3,4)$ and generalize this for $m \ge 3$ and $n \ge 4$. Exception: a different violation is needed for $L(4,4)$; consider the element 1114.)

12.2.26. Let Π_n denote the poset of partitions of $[n]$, ordered by $\sigma \le \tau$ if σ is a union of partitions of the blocks of τ. An element of rank k in Π_n has $n-k$ blocks, so $N_k(\Pi_n) = S(n, n-k)$ (the Stirling number). For even n, let X be the set of partitions of $[n]$ into two blocks of size $n/2$. Use X and its shadow to prove that Π_n is not an LYM order when $n \ge 20$. (Spencer [1974]) (Comment: Rota asked whether Π_n always has the Sperner property; Canfield [1978] showed that it doesn't. Subsequently, Shearer [1979] and Jichang–Kleitman [1984] reduced the least n such that Π_n is not Sperner to 4×10^9 and then to 3.4×10^6.)

12.2.27. ($\diamond$) Let F be a k-family in $\mathbf{2}^n$.

(a) Prove that in each circular permutation of $[n]$, at most nk members of F appear as (consecutive) substrings.

(b) Use part (a) to prove $\sum_{x \in F} N_{r(x)}^{-1} \le k$.

(c) Use part (b) to prove the strong Sperner property for $\mathbf{2}^n$. (Füredi–Katona)

12.2.28. Without using regular coverings or the LYM Inequality, prove that the dual of a graded poset with the normalized matching property also has the property. That is, given that G is an X,Y-bigraph with $|N(S)|/|Y| \ge |S|/|X|$ for all $S \subseteq X$, prove the analogous statement for all $T \subseteq Y$.

12.2.29. A **minimal LYM order** is an LYM order such that deleting any cover relation destroys the LYM property.

(a) Prove that the relations between adjacent ranks of a minimal LYM order form a forest, connected when the rank sizes are relatively prime. Construct an example to show that the forest may or may not be connected when the rank sizes are not relatively prime. (Hint: Consider the Graham–Harper version of the normalized matching condition.)

(b) Prove that every LYM order P has a regular covering using at most $|P| - r(P)$ distinct chains. Prove that this is sharp. (West–Harper–Daykin [1983])

12.2.30. Let $\left[{m+n \atop m}\right]_q = \prod_{j=0}^{m-1} \frac{q^{n+m-j}-1}{q^{m-j}-1} = \sum a_k q^k$. Using generating functions, prove that a_k is the number of partitions of k with at most m parts and largest part at most n (this equals $N_k(L(m,n))$). Prove also that a_k is the number of m-subsets of $[m+n]$ whose elements sum to $k + \binom{m+1}{2}$. (Comment: The expression $\left[{m+n \atop m}\right]_q$ is called a **Gaussian polynomial**).

12.2.31. Let $L_n(q)$ be the inclusion order on the subspaces of an n-dimensional vector space over a field with q elements.

(a) Count the bases in a subspace of dimension k.

(b) Prove that $N_k(L_n(q)) = \dfrac{(q^n-1)(q^{n-1}-1)\cdots(q^{n-k+1}-1)}{(q^k-1)(q^{k-1}-1)\cdots(q^1-1)}$.

(c) Prove that each element of rank k covers $(q^k-1)/(q-1)$ elements of rank $k-1$. Conclude that $L_n(q)$ is an LYM order.

12.2.32. A **regular poset** is a graded poset in which all elements of rank k are covered by the same number of elements at rank $k+1$ and cover the same number of elements at rank $k-1$. A graded poset has the **strict Sperner property** if every maximum antichain consists of a single rank. A regular covering is **exhaustive** if every two comparable elements are together on some chain in the covering.

(a) Show that the product of two chains of differing lengths does not have the strict Sperner property (even though it is an LYM order).

(b) Prove that a poset with an exhaustive regular covering has the strict Sperner property if and only if for every pair of maximum-sized ranks, the bipartite graph of relations between them is connected. (Broline)

(c) Prove that a regular poset has an exhaustive regular covering, and conclude that $\mathbf{2}^n$ has the strict Sperner property.

12.2.33. Let $\langle a \rangle$ and $\langle b \rangle$ be log-concave sequences.

(a) Prove that $a_i a_j \ge a_{i+1}a_{j-1}$ whenever $i \ge j$.

(b) Prove that $\sum_i \sum_j (a_i a_j - a_{i+1}a_{j-1})(b_{k-i}b_{k-j} - b_{k-1-i}b_{k+1-j}) \ge 0$.

(c) (+) Use part (b) to prove that the product of two graded posets with log-concave rank sizes has log-concave rank sizes.

12.2.34. Suppose that $\{a_k\}$ and $\{b_k\}$ are log-concave sequences. Let $g_k(r,s) = \sum_{i=r}^{s} a_i b_{k-i}$. Prove that $g_{k+1}(r, s+1)/g_k(r,s)$ is increasing in r and decreasing in s. (Comment: This completes the proof of Theorem 12.2.31.)

12.2.35. *Derived LYM posets.*

(a) Let P be an LYM order of rank r. For $S \subseteq [r]$, let Q be the subposet consisting of all elements whose rank in P belongs to S. Prove that Q is an LYM order.

(b) Let $\langle a_k \rangle$ be a log-concave sequence. Prove that $a_k^2 \ge a_{k+j}a_{k-j}$ for all k, j.

(c) Let $A_1, \ldots, A_r$ be a partition of $[n]$ into r blocks, and for each $1 \le i \le r$ let S_i be an arithmetic progression. Let $P = \{X \subset [n] \colon |X \cap A_i| \in S_i \text{ for } 1 \le i \le r\}$. Prove that the inclusion order on P is an LYM order. (Griggs)

12.2.36. *Bollobás' Inequality.*

(a) Let $A_1, \ldots, A_m$ and $B_1, \ldots, B_m$ be subsets of $[n]$ such that $A_i \cap B_j = \varnothing$ if and only if $i = j$. Prove that $\sum \binom{|A_i|+|B_i|}{|A_i|}^{-1} \le 1$. (Hint: Consider instances of B_k completely after A_k in permutations of $[n]$.) (Bollobás [1965])

(b) Use part (a) to prove the LYM property for $\mathbf{2}^n$.

(c) Use part (a) to prove that $\binom{n-k}{\lfloor (n-k)/2 \rfloor}$ is the maximum t such that $\mathbf{2}^n$ contains $(k+1)$-chains $C_1, \ldots, C_t$ such that every member of each chain is incomparable to all members of all the other chains in the list. (Griggs–Stahl–Trotter [1984])

12.2.37. Let $\mathcal{F}$ be a family of finite sets $A_1, \ldots, A_m$. Let $\tau(\mathcal{F})$ denote the **transversal number** of $\mathcal{F}$ (the least size of a set intersecting all A_i). A family $\mathcal{F}$ is **τ-critical** if $\tau(\mathcal{F} - \{A_i\}) < \tau(\mathcal{F})$ for all i.

(a) For τ-critical $\mathcal{F}$ with $\tau(\mathcal{F}) = s+1$, prove that $\sum_{i=1}^{m} \binom{|A_i|+s}{s}^{-1} \le 1$. (Hint: Use Exercise 12.2.36a.)

(b) Let G be an n-vertex graph such that addition of any edge of $\overline{G}$ completes an r-clique. Prove that the minimum number of edges in G is $(r-2)(n-(r-3)/2)$.

12.3. Linear Extensions & Dimension

Few partial orders are chains, but chains are useful in understanding more complicated posets. We can describe a poset using chains that are consistent with the order relation, called “linear extensions”.

ORDER DIMENSION

When purchasing a new car, a buyer considers many models. Between any two, the buyer may prefer one or be undecided. Assuming that the preference relation is a partial order, let P be the resulting poset. The buyer may try to encode P by rating the cars on criteria such as price, reliability, mileage, roominess, color, styling, etc. Assume that cars are ranked linearly on each scale. The scales “realize” P when car x is preferred to car y in P if and only if x is preferred to y on each scale. We try to realize a poset using a small number of linear orders.

12.3.1. Definition. An **extension** of a poset P is a partial order on the elements of P that contains all the relations of P. A **linear extension** is an extension that is a chain. The **intersection** of partial orders on a given set is the set of relations that appears in each of them.

12.3.2. Proposition. A finite poset is the intersection of its linear extensions.

Proof: A linear extension iteratively lists (and deletes) a minimal unlisted element. All linear extensions arise in this way. For every down-set F in P, this procedure can produce a linear extension that lists all of F before all of $P-F$. If x and y are incomparable elements of P, then $x \notin D[y]$ and $y \notin D[x]$. We thus have a linear extension listing x after y (and all of $D(y)$) and another listing y after x (and all of $D(x)$). Hence $x < y$ in P if and only if x precedes y in every linear extension. Thus P is the intersection of all its linear extensions. ■

In the algorithmic literature, linear extensions are called **topological orderings**. Proposition 12.3.2 holds also for infinite posets but is more subtle; one applies Zorn’s Lemma to the family of all extensions, ordered by inclusion (Szpilrajn [1930]). We consider only finite posets.

12.3.3. Definition. A **realizer** of P is a set of extensions whose intersection is P. The **(order) dimension** $\dim P$ is the minimum number of extensions in a realizer of P by linear extensions. Let $I(P)$ denote the set of ordered pairs of incomparable elements in P. For $(x, y) \in I(P)$, an extension L **establishes** the pair (x, y) if $x < y$ in L.

A set of extensions realizes P if and only if for every $(x, y) \in I(P)$, some extension in the set has $x < y$. That is, every incomparable pair must be established.

12.3.4. Example. A poset has dimension 1 if and only if it is a chain. Antichains have dimension 2 (list the elements in some order and the reverse order). The poset below also has dimension 2. With 123456 as one extension, the other must have $4 < 1$ and $6 < 3$. The extension 412563 completes a realizer. Not every linear extension is part of a realizer of size 2 (see Exercise 3). For example, starting with 124356, we cannot have $4 < 1$, $3 < 4$, and $6 < 3$ in a single additional extension.

We next characterize posets having dimension 2. ■

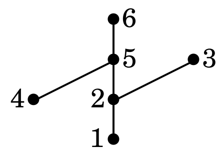

12.3.5. Theorem. (Dushnik–Miller [1941]) A poset P with comparability graph G has dimension at most 2 if and only if $\overline{G}$ is also a comparability graph.

Proof: *Necessity.* Let L_1 and L_2 be two linear extensions of P with intersection P. We have $x \| y$ if and only if x and y appear in opposite order on L_1 and L_2. Hence $x \| y$ if and only if x and y are in the same order on L_1 and L_2', where L_2' is the reverse of L_2. Now $\overline{G}$ is the comparability graph of the intersection of L_1 and L_2'.

Sufficiency. Let F be the comparability digraph of P. Let $\hat{F}$ be a transitive orientation of $\overline{G}$, and let $\hat{F}'$ be its reverse (also transitive). Both $F \cup \hat{F}$ and $F \cup \hat{F}'$ are orientations of K_n, where $n = |P|$. A non-transitive orientation of K_n has a directed 3-cycle, but this would violate the transitivity of F or $\hat{F}$ or $\hat{F}'$. Hence $F \cup \hat{F}$ and $F \cup \hat{F}'$ are transitive orientations of K_n.

A transitive orientation of K_n linearly orders its vertices by outdegree (this holds inductively by deleting a source). Hence there are chains L and L' on $V(G)$ whose comparability digraphs are $F \cup \hat{F}$ and $F \cup \hat{F}'$. The intersection of these two chains is P, since two elements are ordered in the same way on L and L' if and only if they are adjacent in F, which is the comparability digraph of P. ■

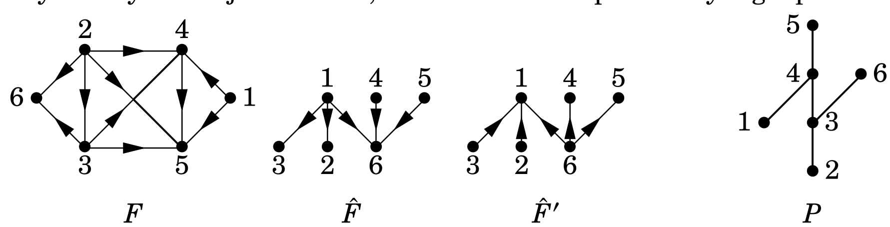

$F \cup \hat{F}$ yields $1 < 2 < 3 < 4 < 5 < 6$.

$F \cup \hat{F}'$ yields $2 < 3 < 6 < 1 < 4 < 5$.

On the other hand, dimension can be arbitrarily large.

12.3.6. Example. (Dushnik–Miller [1941]) The **standard example** S_n is the subposet of $\mathbf{2}^{[n]}$ induced by the singletons and their complements, denoted by i and $\overline{i}$ for $i \in [n]$ (S_4 appears below). In a realizer of S_n, for each i there must be a linear extension in which i appears above $\overline{i}$, establishing the incomparable pair

$(\bar{i}, i)$. This forces the other singletons to appear below this pair and the other sets of size $n-1$ to appear above them. Thus n distinct extensions are needed. Any n extensions establishing these pairs also establish all others. ∎

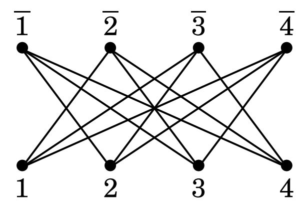

One motivation for studying dimension is compact encoding of n-element posets. A $0,1$-matrix for the order relation takes n^2 bits, testing $x < y$ in unit time by checking an entry. A realizer of size k uses $kn \log_2 n$ bits. This takes less storage when $\dim P \le O(n/\log_2 n)$, although testing $x < y$ takes k comparisons.

Encoding each element by its heights on the k extensions of a realizer embeds a poset in $\mathbb{R}^k$ under the product ordering. Indeed, the least such k is another definition of dimension, often attributed to Ore but given earlier by Hiraguchi.

12.3.7. Theorem. (Hiraguchi [1955], Ore [1962]) A partial order P has a realizer of size k if and only if it embeds in the product of k chains.

Proof: Given a realizer, we obtain such an embedding using the heights on the extensions to map each element of P to an element of the product. Conversely, given an embedding that maps $x \in P$ to $(x_1, \ldots, x_k)$, we seek a realizer by placing the elements on the ith extension in the order of their ith coordinates.

These values need not be distinct; there may be "ties" in coordinate i. Let $S = \{x \in P\colon x_i = t\}$. To break the ties, expand this position on the ith extension into any linear extension of the subposet S. This preserves all the relations of P, so we have obtained linear extensions.

To show that this is a realizer, consider $x, y \in S$ with $x \| y$. Because the encoding embeds P, there are coordinates j and k such that $x_j < y_j$ and $x_k > y_k$. The extensions corresponding to these coordinates have $x < y$ and $y < x$, respectively, which is not affected by breaking ties in another coordinate. ∎

Theorem 12.3.7 yields $\dim \mathbf{2}^{[n]} \le n$, since $\mathbf{2}^{[n]}$ is a product of n chains. Komm [1948] showed that $\dim \mathbf{2}^{[n]} = n$. To see that $\mathbf{2}^{[n]}$ cannot embed in a product of fewer chains, we use S_n and the next observation.

12.3.8. Corollary. If Q is a subposet of P, then $\dim Q \le \dim P$.

Proof: A subposet of P embeds wherever P embeds. Also, dropping $P - Q$ from the extensions in a realizer of P yields a realizer of Q. ∎

By Example 12.3.6, Theorem 12.3.7, and Corollary 12.3.8, products of n nontrivial chains have dimension n (Ore [1962]). We next prove a more general result. The upper bound was observed by Hiraguchi [1951]. The sufficient condition for equality was proved by Baker [1961] using completion of lattices We present the more explicit proof by Kelly [1981] using realizers. A poset is **bounded** if it has a unique minimal element $\hat{0}$ and a unique maximal element $\hat{1}$.

12.3.9. Theorem. If P and Q are posets, then $\dim(P\times Q) \le \dim P + \dim Q$, with equality if P and Q are bounded posets of size at least 2.

Proof: Let $f\colon P \to \mathbb{R}^m$ and $g\colon Q \to \mathbb{R}^n$ be optimal embeddings. Define $h\colon P \times Q \to \mathbb{R}^{m+n}$ by letting $h((p,q))$ be the concatenation of $f(p)$ and $g(q)$. This is an embedding of $P \times Q$, so $\dim P \times Q \le m + n$.

Given bounded posets P and Q with size at least 2, let $L_1, \dots, L_t$ be a realizer of $P \times Q$, indexed so $L_1, \dots, L_m$ are the extensions with $(\hat{0}, \hat{1}) < (\hat{1}, \hat{0})$. Using $(\hat{0}, \hat{1}) \| (\hat{1}, \hat{0})$, we force a realizer of $P \times \{\hat{0}\}$ among the m extensions with $(\hat{0}, \hat{1}) < (\hat{1}, \hat{0})$ and a realizer of $\{\hat{0}\} \times Q$ among the $t - m$ extensions with $(\hat{0}, \hat{1}) > (\hat{1}, \hat{0})$. Thus $m \ge \dim P$ and $t - m \ge \dim Q$, so $\dim(P \times Q) = t \ge \dim P + \dim Q$.

If P is a chain, then already $m \ge \dim P$, so we may choose $(a,b) \in I(P)$. Since $(a, \hat{1}) \| (b, \hat{0})$, we must have $(a, \hat{1}) < (b, \hat{0})$ on some extension. Since $(\hat{0}, \hat{1}) < (a, \hat{1})$ and $(b, \hat{0}) < (\hat{1}, \hat{0})$, such an extension will have $(\hat{0}, \hat{1}) < (\hat{1}, \hat{0})$ and must be one of $L_1, \dots, L_m$. Furthermore, since $(a, \hat{0}) < (a, \hat{1})$, this extension has $(a, \hat{0}) < (b, \hat{0})$. Since $(a,b) \in I(P)$ was arbitrary, $L_1, \dots, L_m$ contains a realizer of the copy $P \times \{\hat{0}\}$ of P, so $m \ge \dim P$. By analogous reasoning, the extensions with $(\hat{0}, \hat{1}) > (\hat{1}, \hat{0})$ realize $\{\hat{0}\} \times Q$, so there are at least $\dim Q$ of those. ■

To see that $\dim P \times Q$ can be less than $\dim P + \dim Q$, consider the simplest product involving a non-bounded poset: $\dim(\bullet\ \bullet)\times(\bullet) = 2 = \dim(\bullet\ \bullet) + \dim(\bullet) - 1$. However, it seems that it cannot be much less.

12.3.10. Conjecture. (Kelly–Trotter [1982]) If P and Q are posets, then $\dim P \times Q \ge \dim P + \dim Q - 2$. ■

Little is known about this conjecture. Trotter [1985] proved $\dim S_n \times S_n = 2n-2$, and Reuter [1989a] extended this to $\dim S_m \times S_n = m+n-2$ (he also proved $\dim P \times P \ge 4$ when $\dim P = 3$).

COMPUTATION AND BOUNDS

Since a set of extensions realizes P if and only if each ordered pair $(x,y) \in I(P)$ appears in some extension, computing dimension is equivalent to covering of $I(P)$ by the fewest sets of pairs that can appear in one extension. This expresses dimension as hypergraph coloring (recall that the chromatic number of a hypergraph is the minimum size of a vertex partition into sets containing no edge).

12.3.11. Definition. An **alternating cycle** of incomparable pairs in P is a set $\{(x_i, y_i)\}_{i=1}^k$ in $I(P)$ such that $y_i \le x_{i+1}$ in P for all i (indices modulo k).

If $x \| y$, then (x,y) and (y,x) together form an alternating cycle.

12.3.12. Lemma. The dimension of P is the minimum number of classes covering $I(P)$ such that no class contains an alternating cycle.

Proof: A single extension cannot establish all pairs in an alternating cycle, because the relations $x_i < y_i$ together with $y_i \le x_{i+1}$ violate transitivity. Thus the number of classes needed is a lower bound on $\dim P$.

For the opposite inequality, let S be a subset of $I(P)$ containing no alternating cycle. The digraph having $y \to x$ whenever $x < y$ in P or $(x, y) \in S$ is acyclic. A linear extension of the transitive closure of this digraph is a linear extension of P that contains all relations in S. Thus $\dim P$ is at most the number of classes in such a covering. ■

Thus $\dim P$ is the chromatic number of a hypergraph where the vertices are $I(P)$, the edges are the alternating cycles, and the colors are the linear extensions. We do not need all of this hypergraph to compute $\dim P$. In computing $\dim S_n$ it was enough to establish the incomparable pairs $(\bar{i}, i)$; the others were then necessarily also established. In general, it suffices to establish the "crucial" incomparable pairs.

12.3.13. Definition. Among ordered incomparable pairs, (a, b) **forces** (c, d) if $c \le a$ and $b \le d$ (every extension with $a < b$ has $c < d$). A pair $(x, y) \in I(P)$ is **unforced** if $x < y$ is not implied by adding any other pair from $I(P)$ to P. Let $C(P)$ denote the set of unforced pairs.

In the literature, when (x, y) is an unforced pair, the ordered pair (y, x) is called a **critical pair**. Realizing a poset requires *reversing* the critical pairs, which is equivalent to *establishing* the unforced pairs.

12.3.14. Proposition. An ordered incomparable pair (x, y) is an unforced pair if and only if $D(y) \subseteq D(x)$ and $U(x) \subseteq U(y)$.

Proof: The pair (x, y) fails to be unforced if and only if there is an incomparable pair (a, b) other than (x, y) such that adding the relation $a < b$ forces $x < y$. Such forcing occurs if and only if $x \le a$ and $b \le y$. Hence it will occur if any element of $U[x]$ is incomparable to any element of $D[y]$. Since x and y must remain incomparable, the forcing fails if and only if all of $U[x]$ is above all of $D[y]$, except for the pair (x, y) itself. This is equivalent to $U(x) \subseteq U(y)$ and $D(y) \subseteq D(x)$. ■

12.3.15. Theorem. The dimension of P is the minimum number of linear extensions establishing the unforced pairs of P. This equals the chromatic number of the hypergraph $H(P)$ with vertex set $C(P)$ whose edges are the minimal alternating cycles consisting of unforced pairs.

Proof: By Definition 12.3.13, extensions that establish $x < y$ for all $(x, y) \in C(P)$ also establish all incomparable pairs. ■

The hypergraph $H(P)$ is useful for computing dimension in special classes because the lower bound can be established by exhibiting any subgraph of $H(P)$ with the desired chromatic number. The upper bound is then verified by exhibiting a realizer that has all the unforced pairs, rather than by verifying that every edge of $H(P)$ is properly colored. For the standard example, $C(S_n) = \{(\bar{i}, i)\colon 1 \le i \le n\}$. The minimal alternating cycles are precisely the sets of two unforced pairs so $H(S_n) = K_n$ and $\chi(H(S_n)) = n$.

12.3.16.* Remark. Since computing $\dim P$ is a hypergraph coloring problem, it is not surprising that testing $\dim P \le 3$ is NP-complete. Yannakakis [1982] showed this by constructing from any graph a poset that has dimension 3 if and only the graph is 3-colorable. Recognition of 2-dimensional posets runs in time linear in the number of comparabilities (McConnell–Spinrad [1999]). Structural descriptions of 2-dimensional posets allow problems that are hard in general to run quickly on this class (see Möhring [1985], Spinrad [1982, 2003]).

We can also seek a forbidden subposet characterization of d-dimensional posets. A poset P is **irreducible** if deleting any element reduces its dimension; it is **k-irreducible** if also $\dim P = k$. The only 2-irreducible poset is the 2-element antichain. Kelly [1977] and Trotter–Moore [1976b] independently found all 3-irreducible posets (see the dimension survey in Kelly–Trotter [1982]). The list includes seven infinite families and ten small examples. ■

Although exact computation of $\dim P$ is difficult, there are bounds in terms of other parameters. For S_n, the dimension equals the width and is half the size. Hiraguchi proved that this is extremal for both parameters. The proof of $\dim P \leq |P|/2$ developed here is due independently to Kimble [1973] and Trotter [1975].

12.3.17. Definition. A linear extension L **puts** Y **over** X if X and Y are disjoint subposets and y is above x in L whenever $x \| y$, $x \in X$, and $y \in Y$. For $Q \subseteq P$, an **upper extension** of Q is a linear extension that puts $P - Q$ over Q; a **lower extension** puts Q over $P - Q$.

12.3.18. Lemma. Every chain in a poset has upper and lower extensions.

Proof: By symmetry, it suffices to find upper extensions. Let C be a chain in a poset P consisting of $x_1, \dots, x_k$ from bottom to top. Let L be a linear extension of $P - C$. For $1 \leq i \leq k-1$, let Y_i be the set of elements of $P - C$ that are less than x_{i+1} but not less than x_i; also let $Y_0 = D(x_1)$ and $Y_k = P - D[C]$. Form a linear ordering L' of the elements of P by inserting between x_i and x_{i+1} all elements of Y_i in the order that they have on L. Similarly insert Y_0 before x_1 and Y_k after x_k.

If $y' < y$ for $y \in Y_i$ and $y' \notin C$, then transitivity yields $y' \in Y_j$ for some $j \leq i$, so L' puts y' and y in the right order. Also x_i and y appear in the right order when they are comparable. If $x_i \| y$, then $y \in Y_j$ for some $j \geq i$. Thus L' is an upper extension of C. ■

Not all subposets have upper or lower extensions. Rabinovitch [1978] determined when there is an extension putting Y over X (Exercise 17).

12.3.19. Theorem. (Dilworth [1950], Hiraguchi [1955]) $\dim P \leq w(P)$.

Proof: Start with a Dilworth decomposition of P (a partition into $w(P)$ chains), and take an upper extension of each chain. The resulting extensions form a realizer, since incomparable elements appear on different chains in the original partition $\mathbf{C}$. If $x \| y$, then x appears above y on the extension arising from the chain of $\mathbf{C}$ containing y, and y appears above x on the extension arising from the chain of $\mathbf{C}$ containing x. ■

In addition to the notations $U(x)$ and $D(x)$ for the sets of elements above and below x, we also use $I(x) = \{y \in P\colon y \| x\}$. Thus $P - x = U(x) \cup D(x) \cup I(x)$.

12.3.20. Theorem. (**One-Point Removal Theorem**; Hiraguchi [1951]) If x is an element of a poset P, then $\dim P \le 1 + \dim(P - x)$.

Proof: We construct a realizer of P from a realizer $\mathbf{L}$ of $P-x$. Extract from P the subsets $U(x)$, $D(x)$, $U(x) \cup I(x)$, and $D(x) \cup I(x)$. Using in each of these subposets the order given by an extension L in $\mathbf{L}$, form the two extensions $U(x) \cup I(x) > x > D(x)$ and $U(x) > x > D(x) \cup I(x)$. Replace L in $\mathbf{L}$ by these two extensions, and insert x anywhere between $U(x)$ and $D(x)$ on the other extensions.

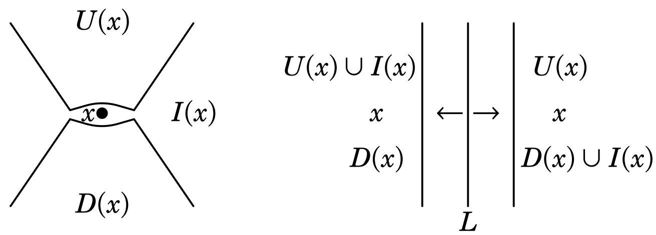

We claim that this yields $1 + \dim(P - x)$ extensions realizing P. The two new extensions establish the incomparable pairs involving x. For incomparable pairs not involving x, we need only worry about pairs y, z with $y \in I(x)$ and $z \in U(x) \cup D(x)$ that were established by L. Since the two new extensions have $U(x) \cup I(x)$ and $D(x) \cup I(x)$ in order as on L, the pair y, z appears as it did in L in one of the new extensions. ∎

12.3.21. Theorem. If A is an antichain of P, then $\dim P \le \max\{2, |P - A|\}$.

Proof: We want $\dim P \le 2$ when $|P - A| \le 2$. This reduces by easy remarks to the two 2-dimensional cases shown below (Exercise 7). For $|P - A| > 2$, deleting all but two elements of $P - A$ from P leaves a poset like one of the 2-dimensional posets drawn below (A may be larger, and $P - A$ may be below A). Repeatedly using the One-Point Removal Theorem to restore the deleted elements yields $\dim P \le |P - A| - 2 + 2$. ∎

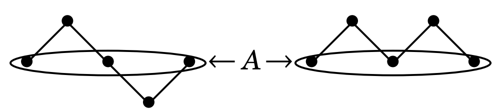

12.3.22. Corollary. (**Hiraguchi's Inequality**; Hiraguchi [1955]) If $|P| \ge 4$, then $\dim P \le |P|/2$.

Proof: To obtain $\dim P \le |P|/2$ when $w(P) > |P|/2$, apply Theorem 12.3.21; otherwise apply Theorem 12.3.19. ∎

The technique of Theorem 12.3.20 yields other *removal theorems*, bounding $\dim P$ in terms of the dimension of a subposet Q. Start with a realizer of Q, modify and/or add extensions appropriately, and show that all the incomparable pairs of P are established (Exercises 18–24).

Theorem 12.3.20 and Corollary 12.3.22 suggest a famous conjecture.

12.3.23. Conjecture. (**Two-Point Removal Conjecture**) Every poset with at least three elements has a **removable pair** of elements; a pair $\{x, y\}$ such that $\dim(P - \{x, y\}) \geq \dim P - 1$. ■

Removable pairs are discussed in Exercises 21–28. Tator [1983] proved a weaker statement than Conjecture 12.3.23: always there exist four points in P whose removal decreases the dimension by at most 2 (Exercise 24).

BIPARTITE POSETS

The standard examples S_n are posets whose elements are all maximal or minimal. We consider more general such posets.

12.3.24. Definition. A **bipartite poset** is a poset having no 3-element chains (bipartite comparability graph). A bipartite poset is **normal** if (1) its comparability graph is connected, and (2) when x and y are on the same level, some element other than x is comparable to x but not to y.

12.3.25. Remark. In a normal bipartite poset, the unforced pairs are the pairs $(x, y) \in I(P)$ such that x is maximal and y is minimal. ■

Various bipartite subposets of $\mathbf{2}^n$ generalize S_n and lead to natural dimension problems. We study those consisting of two ranks.

12.3.26. Definition. Write $\mathbf{2}^n_{l,k}$ for the subposet of $\mathbf{2}^n$ induced by the l-sets and k-sets. Let $d_n(l, k) = \dim \mathbf{2}^n_{l,k}$.

The case of most interest is $d_n(1, k)$. Spencer [1971] showed that $d_n(1, k) \sim c_k \lg\lg n$ when k is constant (we henceforth use lg for $\log_2$ and ln for $\log_e$). Dushnik [1950] computed $d_n(1, k)$ exactly when $k \geq 2\sqrt{n}$; the exact result appears in Theorem 12.3.29. After a slow decline in $d_n(1, k)$ as k decreases from n, the drop becomes rapid for k below $2\sqrt{n}$. The upper and lower bounds when $k \in o(n)$ differ by a factor of $\ln n$.

k	$d_n(1, k)$	reference
$n-1$	n	standard example S_n
$\frac{n-1}{2} \leq k \leq n-2$	$n-1$	Dushnik [1950]
$\frac{n}{3} < k \leq \frac{n}{2} - 1$	$n-2$	Dushnik [1950]
$k = n^\alpha \geq 2\sqrt{n}$	$n - n^{1-\alpha} + O(n^{2-3\alpha})$	Dushnik [1950]
general	$< k(k+1)\ln(ne/k)$	Füredi–Kahn [1986]
$2 \leq k \leq \sqrt{n}$	$\geq \frac{1}{4}k^2$	Exercise 31
constant	$\sim c_k \lg\lg n$	Spencer [1971]
2	$\sim \lg\lg n$	Spencer [1971]

These results use different methods. Dushnik's lower bound uses the pigeonhole principle, Spencer's bounds are by relating the problem to other questions, and the Füredi–Kahn upper bound is probabilistic.

For bipartite posets in general, the computation of dimension can be reduced to constructing an appropriate set of permutations of the minimal elements. By "permutation", here we mean a linear ordering written out as a list of elements in order; we use "permutation" to distinguish a linear ordering of some of the elements from a linear extension of the full poset. An element of a permutation "comes later than" or "follows" all the elements to its left.

12.3.27. Definition. Given a bipartite poset P, let X and Y be the sets of minimal and maximal elements, respectively. For $y \in Y$, let $S_y = \{x \in X\colon x < y\}$. A set $\{L_1, \ldots, L_t\}$ of permutations of X is a **suitable set** for P if whenever $y \| x$ with $y \in Y$ and $x \in X$, some L_i puts x later than all of S_y.

12.3.28. Lemma. If P is a normal bipartite poset, then $\dim P$ equals the minimum size of a suitable set for P.

Proof: Let t be the minimum size of a suitable set for P. Given a suitable set of size t, in each permutation we insert each maximal element y immediately after the last element of S_y. The resulting linear orderings can be viewed as linear extensions of P. In fact, they form a realizer of size t, since the unforced pairs are the max-min pairs (y, x) with $x \notin S_y$. Hence $\dim P \le t$.

For the opposite inequality, consider a smallest realizer of P; this is a set of $\dim P$ linear extensions of P. In each such extension, any maximal element y of P must come later than each element of S_y, since extensions of P preserve the order relation. Hence in an extension where x follows y, also x follows all of S_y. If $x \| y$, then x must follow y in some extension in the realizer. Therefore, deleting the maximal elements from each extension in the realizer yields a set of permutations of the minimal elements that by definition is a suitable set. Hence $t \le \dim P$. ■

In studying $d_n(1, k)$, we thus seek realizers as suitable sets of permutations of $[n]$. Lemma 12.3.28 immediately implies that $d_n(1, k)$ is nondecreasing in k. Dushnik's result thus yields $d_n(1, k)$ exactly for all k with $k \ge 2\sqrt{n}$.

12.3.29. Theorem. (Dushnik [1950]) If $1 \le r \le \sqrt{n}$, then $d_n(1, k) \le n - r$ if and only if $k \le n/r + r - 3$.

Proof: We show that a suitable set of $n - r$ permutations of $[n]$ exists for $\mathbf{2}^n_{1,k}$ if and only if $k \le n/r + r - 3$. Let $t = n - r$.

Necessity. Let $L_1, \ldots, L_t$ be a suitable set for $\mathbf{2}^n_{1,k}$. By symmetry, we may assume that 1 is last in L_1. Thus in L_1, element 1 follows all k-sets omitting 1. Hence in $L_2, \ldots, L_t$ we may move 1 to the beginning. By arguing similarly for $2, \ldots, t$ on $L_2, \ldots, L_t$, we may assume that each L_i ends with i.

Let $R = \{t + 1, \ldots, n\}$. For $x \in R$, let T_x be the set of indices $i \in [t]$ such that x appears last among R in L_i. Let $A_x = (R - \{x\}) \cup T_x$. Note that x does not follow A_x in any L_i (if $i \in T_x$, then i follows x; otherwise, some element of $R - \{x\}$

follows x). Every k-set not containing x must precede x on some L_i, so $|A_x| > k$. Thus $r - 1 + |T_x| > k$, where $r = |R| = n - t$.

On the other hand, since $|R| = r$, the pigeonhole principle guarantees that some $x \in R$ appears last among R at most t/r times. Thus $|T_x| \le t/r$ for some x. Now $k - r + 2 \le \min |T_x| \le t/r$. Using $t = n - r$, we obtain $k \le n/r + r - 3$. Hence this inequality is necessary for $d_n(1, k) \le t$.

Sufficiency. We define permutations $L_1, \ldots, L_t$ of $[n]$ that meet the necessity conditions above and form a suitable set if $k \le n/r + r - 3$ and $r \le \sqrt{n}$. The last element of L_i is i, preceded immediately by R in some order. Each element of R is next-to-last in $\lfloor t/r \rfloor$ or $\lceil t/r \rceil$ of the permutations.

Since $r \le \lfloor \sqrt{n} \rfloor \le \lfloor n/r \rfloor$, we have $r - 1 \le \lfloor (n-r)/r \rfloor = \lfloor t/r \rfloor$. Thus for $x \in R$, we can make each element of $R - \{x\}$ appear immediately before x in one of the permutations where x is next-to-last. Thus for distinct x and y in R, there is a permutation L_j that ends y, x, j.

Let A be a set that does not appear before x in any L_j. If x is next-to-last in L_i, then $i \in A$. If x is third-to-last in L_j, followed by z and j, then A must contain z or j (these pairs are disjoint and omit those i where x is next-to-last on L_i). Thus $|A| \ge \lfloor t/r \rfloor + r - 1 > n/r + r - 3 \ge k$. Hence $L_1, \ldots, L_t$ is a suitable set, and $k \le n/r + r - 3$ yields $d_n(1, k) \le t$. ∎

When $k \ge 2\sqrt{n}$, Theorem 12.3.29 gives $d_n(1, k)$ exactly. For smaller k, it gives no better upper bound than $n - \sqrt{n}$. A technique like that in Theorem 12.3.29 yields a lower bound $d_n(1, k) \ge k^2/4$ when $k \le \sqrt{n}$ (Exercise 31). We present a general upper bound.

12.3.30. Theorem. (Füredi–Kahn [1986]) $d_n(1, k) \le \lceil k(k+1)\ln(ne/k) \rceil$.

Proof: Generate permutations $L_1, \ldots, L_t$ of $[n]$ by selecting each at random from all $n!$ orders. We show that if $t \ge k(k+1)\ln(ne/k)$, then with positive probability these form a suitable set of permutations for $\mathbf{2}^n_{1,k}$. Hence some outcome of the experiment is a realizer of the desired size.

For each k-set S and each $x \in [n]$ with $x \notin S$, the probability is $1/(k+1)$ that x follows all of S on L_j. Hence the probability that x follows all of S in none of the random permutations is $(\frac{k}{k+1})^t$. There are $n\binom{n-1}{k}$ such pairs (S, x). Hence we bound the probability that some pair is not established:

$$\begin{aligned}\mathbb{P}(\text{failure}) &\le n\binom{n-1}{k}\left(\frac{k}{k+1}\right)^t < \binom{n}{k}\left(1 - \frac{1}{k+1}\right)^t \\ &< \left(\frac{ne}{k}\right)^k e^{-t/(k+1)} \le \left(\frac{ne}{k}\right)^k e^{-k\ln(ne/k)} = 1\end{aligned}$$

We used standard inequalities $\binom{n}{k} \le \left(\frac{ne}{k}\right)^k$ and $1 - x \le e^{-x}$ (Chapter 14), slightly weakening the Füredi–Kahn bound to simplify computation. ∎

In the realm of constant k, Spencer's construction of small realizers used special families of sets (he attributed this argument to A. Hajnal).

12.3.31. Definition. A family F of sets is **k-scrambling** if for all choices $S_1, \ldots, S_k$ of k distinct sets in F and all subsets A of the index set $[k]$, the set $(\bigcap_{r\in A} S_r) \cap (\bigcap_{r\notin A} \overline{S}_r)$ is nonempty.

The k-scrambling condition can be stated in several ways. If F is 2-scrambling and $X, Y \in F$, then $X \cap Y$, $X - Y$, $Y - X$, and $\overline{X} \cap \overline{Y}$ are nonempty; for k in general, the condition states that in the Venn diagram on any k sets in F, every cell is nonempty. In terms of elements, the condition is that for $S_1, \ldots, S_k \in F$ and $A \subseteq [k]$, there exists an element x such that $x \in S_r$ for $r \in A$ and $x \notin S_r$ for $r \notin A$.

We will use k-scrambling families to prove an upper bound on $d_n(1,k)$.

12.3.32. Example. The four permutations of the elements 0 through 7 shown below form a suitable set for $\mathbf{2}^8_{1,2}$, proving $d_8(1,2) \le 4$.

$$\begin{aligned} L_1&: 0,1,2,3,4,5,6,7\\ L_2&: 3,2,1,0,7,6,5,4\\ L_3&: 5,4,7,6,1,0,3,2\\ L_4&: 6,7,4,5,2,3,0,1 \end{aligned}$$

It is easy to show directly that this is a suitable set. An element at the end of a permutation follows all pairs among the other elements in that permutation. This takes care of the singletons $1,2,4,7$. The elements $0,3,5,6$ appear next-to-last. When r is next-to-last in a permutation, on that permutation it follows all pairs of other elements except the pairs involving the last element s. In each case, there are two permutations where s immediately precedes r, and each of the other six elements precedes s in one of those permutations. ■

These carefully structured permutations arose as a special case of a general construction using k-scrambling sets.

12.3.33. Lemma. Given a k-scrambling family $S_1, \ldots, S_m$ of subsets of $[t]$, there exist permutations $L_1, \ldots, L_t$ of the numbers 0 through $2^m - 1$ such that if $a < b$ and j is the leftmost position where the m-bit binary expansions of a and b differ, then b follows a on L_i if and only if $i \in S_j$.

Proof: Begin with all the m-bit binary integers in one "group". To produce L_i, perform Steps 1 through m in order as follows. On Step j, each current group splits into two smaller groups. In a current group X, let X_r be the numbers whose expansion has r in coordinate j, for $r \in \{0,1\}$. Replace X with X_1 after X_0 if $i \in S_j$; otherwise put X_0 after X_1.

After Step j, the groups all have size 2^{m-j}; thus an explicit ordering L_i is produced after Step m. Furthermore, the relative ordering between a and b on L_i is determined in Step j, where j is the leftmost position where the expansions of a and b differ (a has 0 there; b has 1). The procedure explicitly puts b in a group after a if and only if $i \in S_j$. ■

Example 12.3.32 arises from Lemma 12.3.33 using the family given by $S_1 = \{1,2\}$, $S_2 = \{1,3\}$, and $S_3 = \{1,4\}$. This is a 2-scrambling family of subsets of $[4]$; here $m = 3$. For each L_i, Step 1 splits the elements into the two groups $\{0,1,2,3\}$ and $\{4,5,6,7\}$. The lower group goes first when $i \in S_1$, which holds for $i \in \{1,2\}$; in L_3 and L_4 the lower group goes last. Step 2 splits the lower group into $\{0,1\}$ and $\{2,3\}$ and the upper group into $\{4,5\}$ and $\{6,7\}$. Within each group, the lower subgroup goes first when $i \in S_2$, which holds for $i \in \{1,3\}$. Step 3 finishes the job, deciding which goes first in each pair.

We will use this construction for the upper bound in the next theorem. Given m sets forming a k-scrambling family of subsets of $[t]$, we will obtain $d_n(1,2) \le t$ when $n = 2^m$. Hence we seek a large k-scrambling family.

Let $M(t,k)$ be the maximum size of a k-scrambling family in $\mathbf{2}^t$. We are particularly interested in $k = 2$. The family of all $\lfloor t/2 \rfloor$-sets in $[t]$ that contain the element 1 is 2-scrambling, so $M(t,2) \ge \binom{t-1}{\lfloor t/2 \rfloor - 1} > 2^t/\sqrt{2\pi t}$. For $k \in \mathbb{N}$, there is a constant c_k such that $M(t,k) \ge c_k^t$ (Exercise 33).

12.3.34. Theorem. (Spencer [1971]) If c is a constant such that $[t]$ has a k-scrambling family of size greater than c^t, then

$$\lg\lg(n-1) < d_n(1,k) < \tfrac{1}{\lg c}\lg\lg n.$$

Proof: *Upper bound.* We prove $d_n(1,k) \le t$ for $n = 2^{M(t,k)}$. Since $M(t,k) > c^t$, this yields $d_n(1,k) < \frac{1}{\lg c}\lg\lg n$. For convenience, let $m = M(t,k)$. Let $\{S_1,\dots,S_m\}$ be a largest k-scrambling family of subsets of $[t]$. Lemma 12.3.33 provides orderings $L_1,\dots,L_t$ of 0 through $n-1$ such that if $a < b$ and j is the leftmost position where a and b differ as vectors, then b follows a on L_i if and only if $i \in S_j$.

We claim that $\{L_1,\dots,L_t\}$ is a suitable set for $\mathbf{2}^n_{1,k}$. Viewing the elements as binary m-vectors, consider a vector b and vectors $a_1,\dots,a_k$ other than b. For $1 \le r \le k$, let j_r be the first coordinate where a_r and b differ. Let $A = \{r\colon b_{j_r} = 1\}$, so $A \subseteq [k]$. Since $\{S_1,\dots,S_m\}$ is a k-scrambling family and $\{j_1,\dots,j_k\}$ is a set of at most k indices, there is a value $i \in [t]$ such that $i \in S_{j_r}$ for $r \in A$ and $i \notin S_{j_r}$ for $r \notin A$. We claim that b occurs after all of $a_1,\dots,a_k$ on L_i.

In constructing L_i, element b is compared with a_r when processing coordinate j_r. If $b_{j_r} = 1$, then $r \in A$ and $i \in S_{j_r}$. If $b_{j_r} = 0$, then $r \notin A$ and $i \notin S_{j_r}$. In either case, b is placed after a_r at stage j. Hence b follows each of $a_1,\dots,a_k$ in L_i. We conclude that $L_1,\dots,L_t$ is a suitable set.

Lower bound. A suitable set of permutations for $\mathbf{2}^n_{1,k}$ is also a suitable set for $\mathbf{2}^n_{1,k-1}$. Hence by Lemma 12.3.28 it suffices to prove the lower bound for $k = 2$. We prove that if $n \ge 2^{2^t} + 1$, then any t orderings of $[n]$ yield a triple that appears monotonically (increasing or decreasing) in each ordering. For such a triple $\{x,y,z\}$ with $x < y < z$, we have y between x and z in each permutation. Hence y never appears after $\{x,z\}$. We conclude that a suitable set must have more than t permutations, so $d_n(1,2) > t$ if $t \le \lg\lg(n-1)$.

We prove the claim by induction on t using the Erdős–Szekeres Theorem (Theorem 10.1.13): in every list of m^2+1 distinct numbers some $m+1$ numbers appear monotonically (Exercise 12.1.13 requests a proof using Dilworth's Theorem). For $t = 1$, five elements suffice to guarantee a monotone triple. For $t > 1$, consider orderings $L_1,\dots,L_t$ on $[2^{2^t}+1]$. The Erdős–Szekeres Theorem yields a set S of size $2^{2^{t-1}}+1$ that appears monotonically in L_t. Within S, the induction hypothesis yields a triple $\{x,y,z\}$ that appears monotonically in each of $L_1,\dots,L_{t-1}$. By the choice of S, this triple is also monotone in L_t. ∎

12.3.35. Corollary. (Spencer [1971])

$$\lg\lg n \le d_n(1,2) < \lg\lg n + \tfrac{1}{2}\lg\lg\lg n + O(1).$$

Proof: We have $M(t,2) \ge \binom{t-1}{\lfloor t/2 \rfloor -1} > 2^t/\sqrt{2\pi t}$ (using Stirling's Formula). By the argument in Theorem 12.3.34, $d_n(1,k) \le t$ when $\lg n = M(t,k)$. Solving for t in $\lg n = 2^t/\sqrt{2\pi t}$ yields the more precise upper bound. ■

Note the effect of the double exponential. Since $M(t,2)$ is the number of $\lfloor t/2 \rfloor$-subsets of $[t]$ containing element 1, for $t = 6$ we have $M(t,2) = 10$. The lower and upper bounds thus yield $4 \le d_{1024}(1,2) \le 6$.

For $d_n(1,2)$, Corollary 12.3.35 establishes the asymptotic behavior. We provided the next term of the upper bound because this is in fact sharp. Indeed, there are four natural related problems whose answers in terms of n are $\lg\lg n + (\frac{1}{2} + o(1)) \lg\lg\lg n$. Besides $d_n(1,2)$, these are the least t such that $\mathbf{2}^t$ has at least n antichains, the chromatic number of a special graph G'_n called the "double shift graph", and a fourth problem presented in Exercise 12.4.8.

Füredi–P.Hajnal–Rödl–Trotter [1992] proved that $\chi(G'_n)$ equals the least t such that $\mathbf{2}^t$ has n antichains, while $\chi(G'_n)$ provides a lower bound on $d_n(1,2)$ that is matched by Corollary 12.3.35. For exact values of n where $d_n(1,2)$ increases, see Hosten–Morris [1999].

12.3.36. Definition. The **shift graph** G_n is the graph with vertex set $\binom{[n]}{2}$ and edges defined by $ij \leftrightarrow jk$ if and only if $i < j < k$ (disjoint pairs are nonadjacent). The **double shift graph** G'_n is the graph with vertex set $\binom{[n]}{3}$ and edges defined by $ijk \leftrightarrow jkl$ if $i < j < k < l$.

12.3.37. Lemma. (A. Hajnal) The chromatic number of the shift graph G_n is $\lceil \lg n \rceil$; this is the least t such that $\mathbf{2}^t$ has at least n elements.

Proof: Given that $\mathbf{2}^t$ has at least n elements, we properly color G_n using $[t]$ as colors. Let $A_1 < A_2 < \cdots < A_n < \cdots$ be a linear extension of $\mathbf{2}^t$. For each pair $ij \in V(G_n)$ with $i < j$, color ij with some element of $A_j - A_i$. Since the ordering is an extension of $\mathbf{2}^t$, such an element exists. Since no element of $A_j - A_i$ can belong to $A_k - A_j$, the coloring is proper.

Conversely, suppose that $\chi(G_n) = t$; we show that $[t]$ has at least n subsets. Consider a proper coloring of G_n using $[t]$ as colors. For each $i \in [n]$, let S_i be the set of colors appearing on vertices of the form ij with $j > i$. If $S_i = S_j$ with $j > i$, then the color c that appears on ij must also appear on jk for some $k > j$. This is impossible in a proper coloring, since $ij \leftrightarrow jk$ in G_n. Hence $S_1, \ldots, S_n$ are distinct. ■

12.3.38. Lemma. The chromatic number of the double shift graph G'_n is the least t such that there are at least n antichains in $\mathbf{2}^t$.

Proof: For the upper bound, we give a proper coloring. Instead of antichains, consider the down-sets they generate. Let $D_1, D_2, \ldots$ be a linear extension of the poset of down-sets in $\mathbf{2}^t$, ordered by inclusion. Associate with each pair $\{p,q\} \in \binom{[n]}{2}$ having $p < q$ a set $A_{pq} \in D_q - D_p$. Given $i < j < k$, color $ijk \in V(G'_n)$ with an element of $A_{jk} - A_{ij}$. Such an element exists, since otherwise $A_{jk} \subseteq A_{ij}$, but the set A_{jk} not contained in D_j cannot be a subset of the set A_{ij} contained in the down-set D_j. Now suppose that ijk and jkl both have color c. The first requires $c \in A_{jk}$, and the second requires $c \notin A_{jk}$. Thus the coloring is proper.

Conversely, let $t = \chi(G'_n)$ and consider an optimal coloring; we show that $\mathbf{2}^t$ has at least n down-sets. For each pair $i, j \in [n]$, let S_{ij} be the set of colors appearing on vertices of the form ijk with $k > j$, and let $B_i = \{S_{ij}\colon j > i\}$. Let D_i be the down-set consisting of all subsets of $[t]$ contained in elements of B_i. We must show that these down-sets are distinct. Note that $D_n = \varnothing$ and $D_{n-1} = \{\varnothing\}$; the other down-sets contain non-empty sets.

Consider $D_i = D_j$ with $i < j \le n-2$. Since S_{ij} is nonempty and contained in D_i, and D_j is generated by B_j, the condition $D_i = D_j$ requires a set of the form S_{jk} that contains S_{ij}. This contradicts the coloring of G'_n, because the color in S_{ij} that is used on ijk cannot appear in S_{jk}. ■

12.3.39. Corollary. (Erdős–A.Hajnal) The chromatic number of the double shift graph G'_n is $\lg\lg n + (\frac{1}{2} + o(1))\lg\lg\lg n$.

Proof: This follows from the Kleitman–Markovsky expression $2^{\binom{n}{\lfloor n/2\rfloor}(1+o(1))}$ for the number of down-sets in $\mathbf{2}^n$ (see Section 11.3). ■

12.3.40. Corollary. The dimension of the subposet of $\mathbf{2}^n$ induced by the sets of sizes 1 and 2 is $\lg\lg n + (\frac{1}{2} + o(1))\lg\lg\lg n$.

Proof: Corollary 12.3.35 establishes the upper bound. We prove that $\chi(G'_n)$ is a lower bound. Let $L_1, \dots, L_t$ be a realizer. For ijk with $i < j < k$, choose $c \in [t]$ such that the singleton j is above the doubleton ik in L_c. Let c be the color of ijk. We cannot also give color c to jkl with $k < l$, because this would place $ik < j < jl < k < ik$ on L_c. Hence this defines a proper t-coloring of G'_n, and $\chi(G'_n)$ is a lower bound on the dimension. ■

Although we have emphasized $d_n(1, k)$, the dimension of the subposet of k-sets and $(n-k)$-sets in $\mathbf{2}^n$ is also of interest. Füredi [1994] proved $d_n(k, n-k) \ge n - 2k + 2$ for $n > 2k$ (Exercise 16.2.18).

EXERCISES 12.3

12.3.1. (−) Prove that the intersection of two order relations on the same set is an order relation.

12.3.2. (−) Prove that a set of elements forms a down-set in a poset P if and only if it is an initial segment of some linear extension of P.

12.3.3. (−) List all linear extensions of the poset in Example 12.3.4. Determine which belong to realizers of size 2.

12.3.4. (−) Construct an explicit minimal realizer for the product of k chains.

12.3.5. (−) Let Q be a subposet of P, and let Q' be a (partial) extension of Q. Prove that P has a (partial) extension P' so that the subposet formed by the elements of Q is Q'. Conclude that every linear extension of a subposet of P occurs in some linear extension of P.

12.3.6. (−) Describe all realizers of size 2 for $\mathbf{m} + \mathbf{n}$ and for the "fence" poset F_n on $x_1, \dots, x_{2m-1}$ having relations $x_{2i} > x_{2i-1}$ and $x_{2i} > x_{2i+1}$ for $1 \le i \le n-1$.

12.3.7. (−) Complete the proof that $\dim P \le \max\{2, P - A\}$ by proving that $\dim P \le 2$ if P has an antichain A such that $|P - A| = 2$.

12.3.8. (−) Use Theorem 12.3.5 to prove that each condition below characterizes when a graph G with vertices $v_1, \dots, v_n$ is the comparability graph of a 2-dimensional poset.

(a) There is a permutation σ such that $v_i v_j \in E(G)$ for $i < j$ if and only if $\sigma_i > \sigma_j$.

(b) There are real intervals $I_1, \dots, I_n$ such that $v_i v_j \in E(G)$ if and only if $I_i \cap I_j \neq \varnothing$ and neither of I_i and I_j contains the other.

12.3.9. (−) Seven posets have cover graph C_6; compute the dimension of each.

12.3.10. Prove that a graph G is the complement of a comparability graph if and only if G is the intersection graph of the curves representing a set of continuous real-valued functions on $[0, 1]$. (Golumbic–Rotem–Urrutia [1983])

12.3.11. Prove that the posets below have dimension 3. List four techniques that can be used for the lower bound.

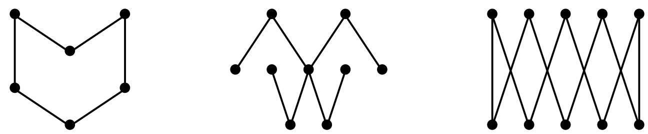

12.3.12. Prove that the poset below has dimension 3.

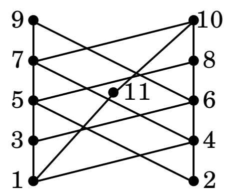

12.3.13. Let P be the graded poset with rank sizes $a_0, \dots, a_k$ such that elements are incomparable if and only if they have the same rank. Describe and count the minimal alternating cycles of incomparable pairs. Prove $\dim P = 2$.

12.3.14. Let $\{(x_i, y_i): 1 \le i \le k\}$ be a minimal alternating cycle of incomparable pairs in a poset P. Prove that $\{x_1, \dots, x_k\}$ and $\{y_1, \dots, y_k\}$ are antichains in P.

12.3.15. The **composition** (also called **lexicographic product**) $Q[P_1, \dots, P_k]$ is formed from a poset Q of size k by expanding each $x_i \in Q$ to a copy of P_i; elements expanded from x_i and x_j are related as x_i and x_j are related in Q. Prove that if $P = P_0[P_1, \dots, P_k]$, then $\dim P = \max_{i \ge 0} \dim P_i$. (Hiraguchi [1951])

12.3.16. (◊) Given a connected graph G, let P be the poset of subsets of $V(G)$ that induce connected subgraphs of G, ordered by inclusion. Prove that $\dim P$ is the number of vertices of G that are not cut-vertices. (Hint: For the upper bound, use distance from non-cut-vertices to partition the ordered incomparable pairs into classes that avoid alternating cycles.) (Trotter–Moore [1976a])

12.3.17. (◊) Let X and Y be disjoint subposets of a poset P. Prove that P has a linear extension putting Y over X if and only if P contains no copy of $\mathbf{2}+\mathbf{2}$ with minimal elements in Y and maximal elements in X. (Hint: Use induction on the number of incomparable pairs (x, y) with $x \in X$ and $y \in Y$.) (Rabinovitch [1978])

12.3.18. In a poset P that is not an antichain, let C be a chain, M be the antichain of maximal elements, and A be an antichain. Prove the following inequalities.

(a) $\dim P \le 2 + \dim(P - C)$. (Hiraguchi p1951])

(b) $\dim P \le 1 + w(P - M)$. (Trotter [1975])

(c) $\dim P \le 1 + 2w(P - A)$. (Trotter [1975])

12.3.19. Prove that the posets below do not contain the standard example S_n and yet have dimension n and width n and are irreducible. (Trotter [1975])

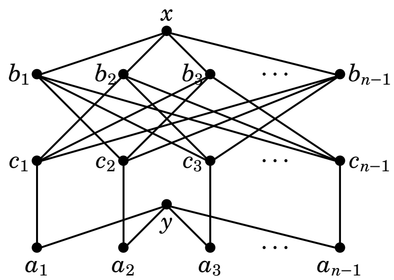

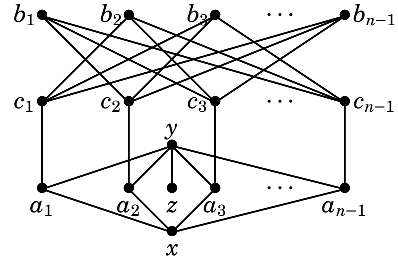

12.3.20. (+) Let C be a chain in a poset P such that each element of $P-C$ is incomparable to at most one element of C. Prove that $\dim P \le 1 + \dim(P-C)$. (Bogart–Trotter [1973])

12.3.21. ($\diamond$) Suppose that a is a maximal element in P, b is a minimal element in P, and $a\|b$. Prove that $\dim P \le 1 + \dim(P - \{a,b\})$.

12.3.22. ($\diamond$) Suppose that $a < b$ in a poset P. Let $r(a,b)$ be the number of ordered pairs $(x,y) \in I(P)$ such that $a < x$ and $y < b$. Prove that (a,b) is a removable pair if $r(a,b) \le \dim P - 3$. (Hint: In a realizer of $P - \{a,b\}$, replace a well-chosen extension with two others to obtain a realizer of P.) (Hiraguchi [1951])

12.3.23. ($\diamond$) Let P be a poset with elements a and b such that $a\|b$. Let $r(a,b)$ be the number of ordered pairs $(x,y) \in I(P)$ such that x is comparable to both a and b and y is incomparable to both a and b. Prove that (a,b) is a removable pair if $r(a,b) \le \dim P - 3$. (Hint: In a realizer of $P - \{a,b\}$, replace a well-chosen extension with two others to obtain a realizer of P.) (Kelly–Trotter [1982])

12.3.24. ($\diamond$) *Four-Point Removal Theorem.* Let P be a poset.

(a) Let C and D be chains in P such that $x\|y$ for all $x \in C$ and $y \in D$. Prove that P has a linear extension that puts $P-C-D$ over C and puts D over $P-C-D$. (Hint: Partition $P-C-D$ into the set X_1 below the top of C, the set X_2 above the bottom of D, and the remainder X_3. Combine linear extensions of $X_1 \cup C$, $X_2 \cup D$, and X_3.) (Hiraguchi [1955])

(b) For C and D as in part (a), prove $\dim P \le 2 + \dim(P-C-D)$. (Hiraguchi [1955])

(c) For x and y maximal in P, with $D(x) \subseteq D(y)$, prove $\dim P \le 1 + \dim(P-x-y)$.

(d) Given $|P| \ge 4$, use parts (b) and (c) to prove that $x,y,z,w \in P$ exist such that $\dim P \le 2 + \dim(P - \{x,y,z,w\})$. (Tator [1983])

12.3.25. ($\diamond$) The k**-dimension** $\dim_k(P)$ of a poset P is the minimum t such that P embeds in $\mathbf{k}^t$. Prove that if P decomposes into t chains of size less than k, then $\dim_k P \le t$. (Comment: Thus k-dimension is well defined.) (Trotter [1976])

12.3.26. ($\diamond$) Prove that the maximum size of a minimal realizer of an n-element antichain is $\lfloor n^2/4 \rfloor$ (for $n \ge 4$). (Maurer–Rabinovitch [1977])

12.3.27. In the poset below, prove that (x,y) is an unforced pair whose removal decreases the dimension by 2. (Reuter [1989b])

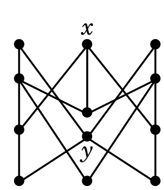

12.3.28. For $n \ge 5$, we construct a poset P_n with $4n-4$ elements that has dimension n and has an unforced pair (y, x) such that $\dim(P_n - \{x, y\}) = n-2$. Begin with disjoint copies A and B of S_{n-2} plus four elements x, y, z, w. Let A_1 and A_2 denote the sets of minimal and maximal elements in A, respectively, and similarly for B. The cover relations involving x, y, z, w are $B_1 \prec x \prec A_2$ and $A_1 \cup B_1 \prec y$ and $\{z, w\} \prec A_2 \cup B_2$ and $w \prec y$. Prove that $\dim P_n = n$, that (y, x) is an unforced pair, and that $\dim(P_n - \{x, y\}) = n - 2$. Which of these properties fails for $n = 4$? The poset P_5 is shown below. (Kierstead–Trotter [1991])

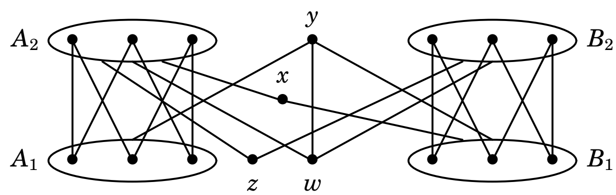

12.3.29. Prove that the n-dimensional standard example S_n is n-irreducible.

12.3.30. The **generalized crown** S_n^k is a bipartite poset with minimal elements i and maximal elements $\bar{i}$ for $1 \le i \le n + k$. They satisfy $i \| \bar{j}$ when j is congruent to one of $\{i, \dots, i + k\}$ modulo $(n + k)$ and otherwise $i < \bar{j}$. (Trotter [1974])

(a) Show that $\overline{k+1} < k+1 < \cdots < \bar{1} < 1$ and $\overline{k+2} < 2 < \cdots < \overline{2k+2} < k+2$ together establish all the unforced pairs involving $\{1, \dots, k+2\}$. Use these and similar extensions to prove that $\dim S_n^k \le \lceil 2(n+k)/(k+2) \rceil$.

(b) Use alternating cycles of length 2 to prove that $\dim S_n^k \ge \lceil 2(n+k)/(k+2) \rceil$.

12.3.31. For $2 \le k \le \sqrt{n}$, prove $d_n(1, k) \ge \frac{1}{4}k^2$. (Hint: Use the idea of Theorem 12.3.29.)

12.3.32. For $k < l$, prove that $d_n(k, l) \ge n - k$ if $l > (n + k - 1)/2$, and $d_n(k, l) \ge l$ if $l \ge 2k$. (Hint: Use Dushnik's Theorem.) (Füredi [1994])

12.3.33. *k-scrambling sets.*

(a) Prove that the orderings on 2-sets in the proof of Theorem 12.3.34 are transitive.

(b) Prove that the $\lfloor t/2 \rfloor$-sets containing a fixed element form a maximum 2-scrambling family in $\mathbf{2}^t$. (Hint: For odd t, use the Erdős–Ko–Rado Theorem (Theorem 11.2.18).)

(c) Prove existence of a k-scrambling family in $\mathbf{2}^t$ with size at least $\frac{1}{2}[(1 - 2^{-k})^{-1/k}]^t$. (Hint: Generate m subsets of $[t]$ independently at random, and prove that the probability they are not k-scrambling is less than 1 when m is smaller than desired.) (Spencer [1971])

12.3.34. ($\diamond$) Let $f(r, k)$ be the maximum size of a k-scrambling family of subsets of $[r]$. Let $\kappa(n, m)$ be the minimum size of a set of vertices in the hypercube Q_n intersecting every m-dimensional subcube. (Graham–Harary–Livingston–Stout [1993])

(a) Determine $\kappa(4, 2)$.

(b) Prove that $\kappa(n, m) = \min\{r\colon f(r, n-m) \ge n\}$. (Hint: Using an incidence matrix, establish a bijection between n-tuples of subsets of $[r]$ that are $(n - m)$-scrambling and r-tuples of vertices (binary vectors) that break all m-cubes in the n-dimensional hypercube.)

12.4. Special Families of Posets

In this section we first consider "chain-like" posets, which maintain some aspects of linear orders as tools of measurement and comparison. Subsequently, we consider posets with special algebraic properties.

SEMIORDERS AND INTERVAL ORDERS

A **ranking** or **weak order** is a partial order whose elements occur in ranks $P_1, \dots, P_k$ such that elements are incomparable if and only if they belong to the same rank. Rankings are used in voting theory; one seeks a consensus ranking among all voters. Paradoxes abound, and Arrow [1951] proved that no function producing a consensus ranking can satisfy a particular set of four natural axioms.

Rankings are too limited to provide a rich generalization of chains. They have dimensions 2, and their comparability graphs are complete multipartite graphs. Most importantly, they are unrealistic as models of preference, since they require transitivity of indifference.

A person asked to choose between cups of coffee with different amounts of sugar is likely to be indifferent when the amount differs by one or two grains, but a large enough difference will yield a preference. A difference of a few dollars in the price of a house won't affect one's attitude toward it, but thousands of dollars will. Luce [1956] introduced a model for "just-noticeable" difference.

12.4.1. Definition. A **semiorder** is a poset that can be represented by a real-valued function f and fixed threshold $\delta \geq 0$ such that $x < y$ if and only if $f(y) - f(x) > \delta$ (this is a **semiorder representation**).

The rankings are the posets having semiorder representations with $\delta = 0$. By scaling, we may assume that the threshold δ is 1 when it is nonzero. The terms *weak order* and *semiorder* suggest weakening the conditions for an order relation, but these are quite restricted posets; what they weaken is the condition of total order. We will characterize semiorders as the posets not containing $\mathbf{1}+\mathbf{3}$ or $\mathbf{2}+\mathbf{2}$.

There may be uncertainty not only in comparison of elements, but also in assignment of values. For example, the skill of a tennis player may vary from day to day, leading us to use an interval $[a_x, b_x]$ to represent a player x. We might then conclude that x beats y if the interval for x is wholly above the interval for y. When $a_x > b_y$, we expect x to win.

12.4.2. Definition. An **interval order** is a poset that can be represented by assigning a real interval to each element such that $x < y$ if and only if the interval for x ends before the interval for y begins (the assignment is an **interval representation**).

12.4.3. Example. Among topics to be discussed in a committee, we set $x < y$ if topic x must be settled before topic y is discussed. On the other hand, we set $x \| y$ if x and y will be available for discussion at the same time. A schedule assigns an interval of time to each topic during which it is available for discussion. A suitable schedule exists if and only if these constraints form an interval order. ■

12.4.4. Remark. *Every semiorder is an interval order.* From a semiorder representation f, letting $[a_x, b_x] = [f(x) - \frac{\delta}{2}, f(x) + \frac{\delta}{2}]$ yields an interval representation. The incomparability graph of an interval order is an interval graph, because elements are incomparable when their intervals in an interval representation intersect. The incomparability graph of a semiorder has an interval representation using intervals of length 1. ■

We will characterize interval orders as the posets not having $\mathbf{2}+\mathbf{2}$ as a subposet. The table below compares the four classes we have discussed.

class	representing function(s)	forbidden subposet
chain	distinct values	$\mathbf{1}+\mathbf{1}$
ranking	$x<y$ if $f(x)<f(y)$	$\mathbf{1}+\mathbf{2}$
semiorder	$x<y$ if $f(x)<f(y)-\delta$	$\mathbf{1}+\mathbf{3}$ and $\mathbf{2}+\mathbf{2}$
interval order	$x<y$ if $f(x)\le g(x)<f(y)\le g(y)$	$\mathbf{2}+\mathbf{2}$

In specifying an interval order, the functions f and g above give the left and right endpoints of the corresponding interval, respectively.

The characterizations of interval orders and semiorders by forbidden subposets can be used to construct representations. Semiorders were characterized much earlier, but it is convenient to characterize interval orders first and then characterize semiorders among them.

12.4.5. Theorem. (Fishburn [1970], Mirkin [1972]) A poset is an interval order if and only if it does not contain $\mathbf{2}+\mathbf{2}$ as a subposet.

Proof: (Balof–Bogart [2003]) Suppose that P has an interval representation and that x,y,z,w are four elements inducing $\mathbf{2}+\mathbf{2}$ with $x<y$ and $z<w$. Let $[a_i,b_i]$ be the interval representing $i\in\{x,y,z,w\}$. Because $x<y$ and $z<w$, we have $b_x<a_y$ and $b_z<a_w$. From $x\|w$ and $z\|y$, we have $b_x\ge a_w$ and $b_z\ge a_y$. This yields the contradiction $b_x<a_y\le b_z<a_w\le b_x$. Thus the condition is necessary.

For the converse, we use induction on $|P|$ to produce an interval representation for a poset P without $\mathbf{2}+\mathbf{2}$. Choose $x\in P$ to maximize $|D(x)|$; note that x is a maximal element. Let $P'=P-x$. Since P' has no copy of $\mathbf{2}+\mathbf{2}$, the induction hypothesis yields an interval representation of P'.

The key claim is that elements of $I(x)$ are maximal. If z is incomparable to x but not maximal, then choose $w\in U(z)$. Since $|D(x)|\ge|D(w)|$ and $z\in D(w)-D(x)$, there exists $y\in D(x)-D(w)$. Since $z\|x$, also y is incomparable to z and w, and the subposet formed by $\{x,y,z,w\}$ is $\mathbf{2}+\mathbf{2}$.

Conversely, maximal elements are incomparable to x, so $I(x)$ is the set of maximal elements in P other than x. In an interval representation of P', extend the intervals for all maximal elements rightward to a common endpoint a. Since the elements are maximal, this changes no intersections, and a is now the rightmost point in the representation. Among these maximal elements, extend the intervals for those that are also maximal in P further rightward to b. Adding an interval for x that starts at $(a+b)/2$ completes an interval representation of P. ■

The earlier proof of Theorem 12.4.5 in Bogart [1993] produced an explicit interval representation with the minimum number of endpoints (Exercise 6). We build on Theorem 12.4.5 to characterize semiorders, as did Balof–Bogart [2003].

12.4.6. Theorem. (**Scott–Suppes Theorem**; Scott–Suppes [1958]) A poset is a semiorder if and only if neither $\mathbf{2}+\mathbf{2}$ nor $\mathbf{3}+\mathbf{1}$ is a subposet.

Proof: (Bogart–West [1999]) A semiorder is an interval order, so $\mathbf{2}+\mathbf{2}$ is forbidden. Given a semiorder representation f with threshold 1, if $\mathbf{3}+\mathbf{1}$ is a subposet with $x<y<z$ and incomparable element w, then $f(w)\ge f(y)$ contradicts $x\|w$ and $f(w)\le f(y)$ contradicts $z\|w$. Hence the condition is necessary.

For sufficiency, suppose that P has no $\mathbf{2}+\mathbf{2}$ or $\mathbf{3}+\mathbf{1}$. By Theorem 12.4.5, P is an interval order, representable as in that proof. We convert this representation to one whose intervals have the same length. Since P has no $\mathbf{3}+\mathbf{1}$, there is no pair $x, y \in P$ such that (1) $I_y \subset I_x$ and (2) I_x intersects intervals to the left and right of I_y that do not intersect I_y. This enables us to alter the representation so that no interval properly contains another. If $I_x = [a, b]$ and $I_y = [c, d]$ with $a < c \le d < b$, then we know that $[a, c]$ or $[d, b]$ contains no endpoint of an interval not intersecting I_y. Hence we can extend I_y past the end of I_x on one end.

Doing this until no more pairs of intervals are related by inclusion yields a proper interval representation. From this we obtain a representation using intervals of length 1: a semiorder representation. When no interval properly contains another, the left ends appear in the same order as the right ends. We process the representation from left to right, adjusting all intervals to have length 1.

Of the remaining unadjusted intervals, let I_x be one with leftmost left endpoint, with $I_x = [a, b]$. If some interval has right end in $[a, b)$, then its left end is before a, and the interval already has length 1. If this occurs, then let α be the largest such right end; otherwise, let $\alpha = a$. In either case, $\alpha \in [a, a+1)$.

Now, adjust the portion of the representation in $[a, \infty)$ by shrinking or expanding $[\alpha, b]$ to $[\alpha, a+1]$ and translating $[b, \infty)$ to $[a+1, \infty)$. The order of endpoints does not change, intervals that begin before a still have length 1, and I_x also now has length 1. Iterating produces the desired representation. ∎

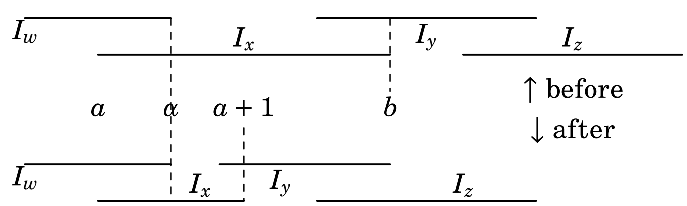

Counting interval orders is hard (Hanlon [1982]), but semiorders behave much better. There are $\frac{1}{n+1}\binom{2n}{n}$ isomorphism classes of n-element semiorders, proved bijectively using the Catalan numbers (Dean–Keller [1968], Exercise 11).

Mitas [1994] studied semiorders representable by open intervals of length k with integer endpoints. For each k, she obtained a forbidden subposet characterization where the number of forbidden subposets is the kth Catalan number!

LATTICES

Some special posets admit algebraic operations that generalize the notions of intersection/union for subsets and gcd/lcm for divisibility.

12.4.7. Definition. If $x \le y$, then x is a **lower bound** for y and y is an **upper bound** for x. A poset L is a **lattice** if for all $x, y \in L$ there is a unique maximal common lower bound (the **meet** $x \wedge y$) and a unique minimal common upper bound (the **join** $x \vee y$). Again, a poset P is **bounded** if it has one minimal element $\hat{0}$ and one maximal element $\hat{1}$.

"Unique maximal common lower bound" means that $x \wedge y$ is an upper bound for every common lower bound of x and y; that is, the subposet of common lower bounds has a unique maximal element. The definition immediately implies that the meet and join operations are commutative and associative.

12.4.8. Example. *Subsets and divisors.* The poset $\mathbf{2}^n$ and the divisibility poset on divisors of n are lattices. To be contained in sets x and y, a set must be contained in their intersection, and $x \cap y$ is above (that is, contains) every such set. Hence $x \cap y$ is the unique maximal common lower bound. Similarly, $x \cup y$ is the unique minimal common upper bound. Hence meets and joins exist, and $\mathbf{2}^n$ is a lattice.

In a divisibility poset, meet and join are least common multiple and greatest common divisor, respectively. ∎

12.4.9. Proposition. A product of lattices is a lattice.

Proof: Let $P = L_1 \times \cdots \times L_n$. A lower bound for an element must be a lower bound coordinate by coordinate, and similarly for upper bounds. Thus meets and joins arise componentwise, with $(x \wedge y)_i = x_i \wedge_{L_i} y_i$ and $(x \vee y)_i = x_i \vee_{L_i} y_i$. ∎

The divisibility poset on divisors of an integer and the containment poset on multisets with bounded multiplicities can both be expressed as products of chains. We can write an element a in a product of n chains as $(a_1, \ldots, a_k)$, where a_i is the height of the ith coordinate of a on its chain. The meet and join operations then have simple formulas: $(a \wedge b)_i = \min\{a_i, b_i\}$ and $(a \vee b)_i = \max\{a_i, b_i\}$.

12.4.10. Example. Π_n*: The lattice of partitions of* $[n]$. A **refinement** of a partition replaces each block with a partition of that block. Put $\sigma < \tau$ in Π_n when σ is a refinement of τ (Π_3 and Π_4 appear below). Rank is given by $r(\sigma) = n - b(\sigma)$, where $b(\sigma)$ is the number of blocks in σ.

For partitions π and τ, the meet $\pi \wedge \tau$ is the common refinement of π and τ with the fewest blocks, and $\pi \vee \tau$ is the partition with the most blocks that does not "split" any block of π or τ. ∎

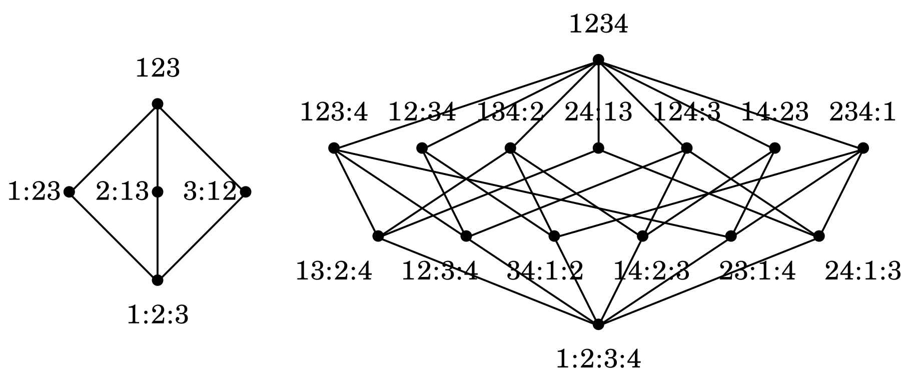

The argument for multiset lattices in Example 12.4.8 generalizes.

12.4.11. Definition. A subposet P of a lattice L is a **sublattice** of L if $x \wedge_L y$ and $x \vee_L y$ are in P for all $x, y \in P$. Equivalently, a lattice L is a sublattice of a lattice M if there is an embedding $f\colon L \to M$ such that $f(x \wedge_L y) = f(x) \wedge_M f(y)$ and similarly for join.

12.4.12. Example. *$J(P)$: The poset of down-sets.* Let $J(P)$ denote the containment poset on the family of down-sets in P. Each down-set is a subposet of P, so $J(P) \subseteq \mathbf{2}^{|P|}$. Since $J(P)$ is a containment poset, common lower bounds of elements x and y in $J(P)$ are contained in $x \cap y$, and common upper bounds contain $x \cup y$. Since the intersection and union of down-sets in P are down-sets in P, we have $x \wedge y = x \cap y$ and $x \vee y = x \cup y$. Hence meet and join in $J(P)$ agree with meet and join when viewed in all of $\mathbf{2}^P$. Thus $J(P)$ is a sublattice of $\mathbf{2}^{|P|}$. The poset is bounded, with $\hat{1} = P$ and $\hat{0} = \varnothing$. Indeed, $J(P)$ is graded, with $r(I) = |I|$.

Antichains in P correspond to down-sets in P. Consider antichains A and B generating down-sets $D[A]$ and $D[B]$. We have $D[A] \subseteq D[B]$ if and only if, for every $x \in A$, there exists $y \in B$ such that $x \le y$. This condition just rephrases the containment condition on down-sets, so $J(P)$ and the resulting lattice on antichains are isomorphic. The maximum antichains induce a sublattice (Exercise 30), which was used in the original proof of the Greene–Kleitman Theorem. ∎

12.4.13. Example. *$L(m, n)$, again.* The poset $L(m, n)$ of Example 12.4.13 arises as a lattice of down-sets: $L(m, n) \cong J(\mathbf{m} \times \mathbf{n})$ (Exercise 18). The isomorphism maps $a \in L(m, n)$ to the down-set of $\mathbf{m} \times \mathbf{n}$ generated by $\{(m + 1 - i, a_i)\colon a_i > 0\}$. The down-set in $\mathbf{4} \times \mathbf{5}$ corresponding to $(0, 1, 5, 5) \in L(4, 5)$ is shown below.

Another proof that $L(m, n)$ is a lattice is by applying Definition 12.4.11 to the lattice $(\mathbf{n + 1})^m$ that contains it (Exercise 19). ∎

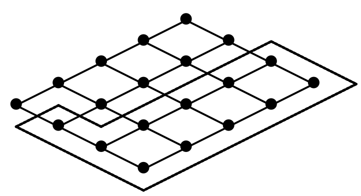

We develop several properties that hold for all lattices.

12.4.14. Lemma. For elements x, y, z of a lattice L,

(a) $x \wedge y \le x \le x \vee z$.

(b) If $x \le z$, then $x \wedge y \le z \wedge y$ and $x \vee y \le z \vee y$.

(c) (4-point Lemma) If $z, w \le x, y$, then $z \vee w \le x \wedge y$.

(d) $(x \wedge y) \vee (x \wedge z) \le x \wedge (y \vee z)$.

(e) If $z \le x$, then $x \wedge (y \vee z) \ge (x \wedge y) \vee z$.

Proof: (a): This holds by definition.

(b): Since $x \wedge y$ is a common lower bound for y and z, it lies below the unique greatest lower bound (the second conclusion is symmetric).

(c): Since both z and w are lower bounds for each of x and y, they are lower bounds for $x \wedge y$. Hence $x \wedge y$ is a common upper bound for z and w, which yields $x \wedge y \ge z \vee w$ (see figure below).

(d): By statement (a), x is an upper bound for $x \wedge y$ and $x \wedge z$, and $y \vee z$ is an upper bound for both y and z and hence for $x \wedge y$ and $x \wedge z$. Hence (d) follows from (c) by using x and $y \vee z$ as $\{x, w\}$ in (c) and $x \wedge y$ and $x \wedge z$ as $\{x, y\}$ in (c).

(e): When $z \le x$, we have $x \wedge z = z$, so (e) follows immediately from (d). ∎

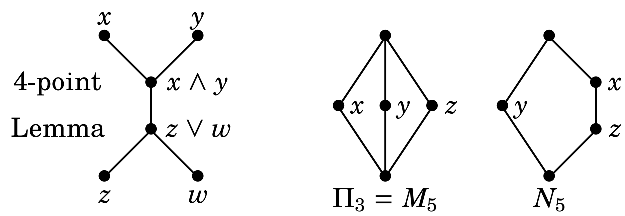

12.4.15. Example. Examples of strict inequality in Lemma 12.4.14(d) and Lemma 12.4.14(e) occur in in the two lattices M_5 and N_5 shown on the right above; we will explain their names later. Note that M_5 also equals Π_3.

12.4.16. Definition. For elements $x, y \in P$ with $x \leq y$, the **interval** $[x, y]$ is $\{z \in P\colon x \leq z \leq y\}$. A poset is **locally finite** if every interval is finite.

Every interval $[x, y]$ in $\mathbf{2}^n$ with $r(x) = k$ and $r(y) = l$ is isomorphic to $\mathbf{2}^{l-k}$. Also, every interval in a chain-product is a chain-product. Meets and joins of elements in an interval also lie in the interval.

All intervals in lattices are sublattices, but there may be sublattices that are not intervals. For example, three of the nine sublattices of $\mathbf{2}^3$ isomorphic to $\mathbf{2}^2$ are not intervals in $\mathbf{2}^3$. Also, a subposet of P that is a lattice need not be a sublattice of P. For example, N_5 is a subposet of $\mathbf{2}^{[3]}$ but is not a sublattice of $\mathbf{2}^{[3]}$.

The next lemma saves some work in proving that posets are lattices. If the poset has an upper bound (or a lower bound, by symmetry), then it is not necessary to construct both meets and joins. The argument is reminiscent of the 4-point Lemma, but in the 4-point Lemma we start with a lattice, and here we are finding joins to show that we have a lattice.

12.4.17. Lemma. If P is locally finite, has an upper bound (the element $\hat{1}$), and has a well-defined meet operation, then P is a lattice.

Proof: It suffices to prove that joins exist. Consider $x, y \in P$. The upper bound 1 is a common upper bound for x and y. Since P is locally finite, the interval from $x \wedge y$ to $\hat{1}$ is finite. Thus we can consider the minimal elements among the set of common upper bounds of x and y; they lie in the interval $[x \wedge y, \hat{1}]$.

Let u and v be minimal common upper bounds for x and y. Since x and y are common lower bounds for u and v, we have $x < u \wedge v$ and $y < u \wedge v$. Thus $u \wedge v$ is a common upper bound for x and y. Since u and v are minimal such elements, $u \wedge v \in \{u, v\}$. Thus u and v are comparable. Since they are minimal elements in a subposet of P, they must therefore be equal. Hence there is a unique minimal common upper bound for x and y. ■

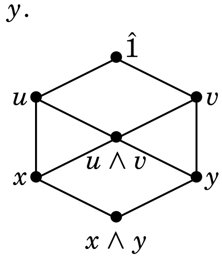

DISTRIBUTIVE LATTICES

Lemma 12.4.14(d) states that $(x \wedge y) \vee (x \wedge z) \leq x \wedge (y \vee z)$ for all x, y, z in any lattice. In the subset lattice, equality holds. We study the class of lattices where equality always holds, seeking characterizations of the class and properties of the subset lattice that extend to such lattices.

12.4.18. Definition. A lattice L is **distributive** if meet distributes over join in L; that is, $x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$ for all $x, y, z \in L$.

Exercise 35 requests a direct proof that a lattice L is distributive if and only if its dual L^* is distributive, by showing that the defining condition for distributivity is equivalent to the property that $x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z)$ for all $x, y, z \in L$. We will also see other ways to prove the equivalence.

12.4.19. Example. *Subset and divisor lattices.* Distributivity for the subset lattice can be seen by marking $x \wedge (y \vee z)$ and $(x \wedge y) \vee (x \wedge z)$ in a Venn diagram.

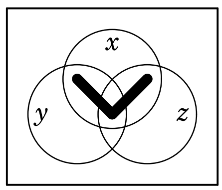

For the divisor lattice, we can argue directly about divisors using gcd and lcm to show distributivity. Alternatively, since every chain is a distributive lattice (min distributes over max for integers), the conclusion that M^e is distributive follows immediately from the next lemma. ■

12.4.20. Lemma. A product of lattices is distributive if and only if each factor is distributive.

Proof: The order relation in a product is defined componentwise, so $(x \vee y)_i = x_i \vee y_i$ and $(x \wedge y)_i = x_i \wedge y_i$. Thus distributivity holds for the full lattice if and only if it holds in each factor. ■

Recall that P is a sublattice of a lattice L if and only if P is closed under the taking of meets and joins in L.

12.4.21. Lemma. Every sublattice of a distributive lattice is distributive.

Proof: The sublattice inherits the distributivity condition from the full lattice, since meets and joins are computed in the full lattice. ■

Lemma 12.4.21 implies that there is a forbidden sublattice characterization of distributive lattices. The two 5-element lattices M_5 and N_5 in Example 12.4.15 are not distributive; they violate the condition with x, y, z as illustrated. Forbidding M_5 and N_5 as sublattices is thus necessary for distributivity. It is also sufficient, but we will not prove this.

12.4.22. Theorem. A lattice is distributive if and only if it does not have M_5 or N_5 of Example 12.4.15 as a sublattice. ∎

Theorem 12.4.22 is quite strong in proving properties of distributive lattices. For example, since M_5 and N_5 are self-dual, Theorem 12.4.22 implies immediately that a lattice L is distributive if and only if L^* is distributive. Thus interchanging meet and join in Definition 12.4.18 yields an equivalent condition.

Lemma 12.4.21 and Example 12.4.19 imply that every sublattice of a subset lattice (or of a chain product) is distributive. For example, since $J(P)$ is a sublattice of $\mathbf{2}^{|P|}$, always $J(P)$ is distributive, for any poset P. Our main objective in this discussion of distributive lattices will be a proof that every finite distributive lattice L can be expressed as $J(P)$ for an appropriate poset P.

12.4.23. Definition. An element p of a lattice L is **join-irreducible** if it is nonminimal and is not the join of two other elements; equivalently, $p = x \vee y$ implies $p \in \{x, y\}$. Similarly, p is **meet-irreducible** if it is not maximal and $p = x \wedge y$ implies $p \in \{x, y\}$. In a lattice L, we write $P(L)$ and $Q(L)$ for the subposets formed by the join-irreducible elements and the meet-irreducible elements, respectively.

12.4.24. Example. In a finite lattice, the join-irreducible elements are those covering exactly one element. In the subset lattice these are the 1-sets. In the divisor lattice these are the powers of primes.

In the lattice L on the left below, the elements labeled by single letters are the join-irreducible elements. The resulting subposet $P(L)$ is on the right. The label for each $x \in L$ is the set of join-irreducible elements whose join is x. Uniqueness of such expressions is our next objective. The minimal element of L is not considered join-irreducible, and the maximal element is not meet-irreducible. ∎

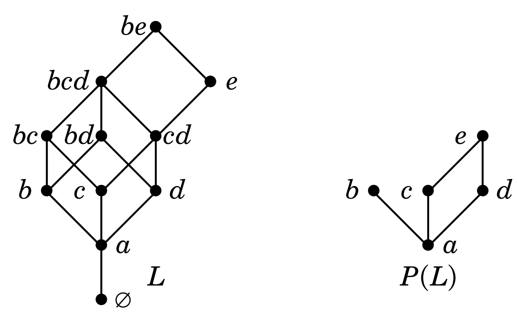

As we will see in Theorem 12.4.36, distributive lattices are the appropriate general setting for many results on $\mathbf{2}^n$ and M^e. It helps to keep the divisor lattice in mind when discussing them. The proof that always $L \cong J(P(L))$ when L is distributive generalizes the proof of the Fundamental Theorem of Arithmetic about unique factorization into primes. In this discussion, we abbreviate "join-irreducible" to **irreducible**.

By induction on the size of the set, in a lattice every finite set has a unique least common upper bound, so we write joins of finite sets without parentheses. Induction also shows that meet distributes over a join of a finite set of elements.

12.4.25. Lemma. If $p \le a_1 \vee \cdots \vee a_k$ for some irreducible element p in a distributive lattice L, then $p \le a_i$ for some i.

Proof: The definition of meet yeilds $x \le y$ if and only if $x \wedge y = x$. With $x = p$ and $y = a_1 \vee \cdots \vee a_k$, we have

$$p = p \wedge (a_1 \vee \cdots \vee a_k) = (p \wedge a_1) \vee \cdots \vee (p \wedge a_k).$$

This expresses p as the join of several elements. Since p is irreducible, it must equal one of them. Thus $p = p \wedge a_i$ for some i, which in turn yields $p \le a_i$. ∎

12.4.26. Definition. An **irredundant representation** of a is a minimal expression of a as a join of a set irreducible elements; that is, $a = p_1 \vee \cdots \vee p_k$ and no proper subset of $\{p_1, \ldots, p_k\}$ has join a.

The irreducible elements in an irredundant representation are incomparable, since if $p > q$ then q is redundant in any join of a set containing p. To prove existence and uniqueness of irredundant representations, we need a unique minimal element and a finiteness condition for inductive arguments. For clarity and simplicity, we restrict to finite lattices our discussion of this lemma and its application to characterize distributivity. The results can be extended to infinite lattices satisfying appropriate local finiteness conditions (such as the divisibility order on $\mathbb{N}$), but we will not discuss this.

12.4.27. Lemma. In a finite distributive lattice L having a lower bound $\hat{0}$, every element has a unique irredundant representation.

Proof: We first prove existence. Note that the identity element for join is $\hat{0}$; following our usual convention, $\hat{0}$ is thus the join of the empty set of irreducible elements. Now, if some element has no irredundant representation, then there is a minimal such element x, and x is the join of two lower elements. By minimality, those elements have irredundant representations, and x is the join of the union of those two sets. Deleting the non-maximal elements in the union yields an irredundant representation.

Now let $p_1 \vee \cdots \vee p_k$ and $q_1 \vee \cdots \vee q_l$ be irredundant representations of a. Since $p_i \le a$, Lemma 12.4.25 implies that each p_r satisfies $p_r \le q_s$ for some s. Similarly each q_s satisfies $q_s \le p_t$ for some t. This yields a relation $p_r \le p_t$, which forces $p_r = q_s = p_t$ since $p_1, \ldots, p_k$ form an antichain. Hence each p_i belongs to $\{q_j\}$, and similarly each q_j belongs to $\{p_i\}$, and the sets are the same. ∎

12.4.28. Definition. The unique irredundant representation of an element x in a finite distributive lattice is the **factorization** of x. For a lattice L, the **ideal map** ϕ: $L \to J(P(L))$ assigns to each $x \in L$ the down-set of join-irreducibles defined by $\phi(x) = \{p \in P(L) \colon p \le x\}$.

We use the word "factorization" because for a divisor lattice, the factorization of an element is the set of prime powers in its numerical prime factorization.

12.4.29. Lemma. For x in a finite distributive lattice L, the factorization of x is the antichain of maximal elements in $\phi(x)$, where ϕ is the ideal map on L.

Proof: Let A be the antichain of elements in the factorization of x. Since $x = \bigvee A$ and $A \subseteq P(L)$, we have $A \subseteq \phi(x)$. Hence $\bigvee A \leq \bigvee B$, where B is the set of maximal elements in $\phi(x)$, since every element of A is bounded above by some element of B. Now $\bigvee B \leq x$, since $B \subseteq \phi(x)$ and x is an upper bound for $\phi(x)$. We have proved $x = \bigvee A \leq \bigvee B \leq x$, and hence $\bigvee A = \bigvee B = x$. By Lemma 12.4.27, $A = B$. ■

12.4.30. Theorem. (Birkhoff [1935]) A finite lattice L is distributive if and only if $L \cong J(P(L))$ (and hence L is a sublattice of $\mathbf{2}^{P(L)}$).

Proof: Because every sublattice of a distributive lattice is distributive, the condition is sufficient. For necessity, suppose that L is distributive. We prove that the ideal map $\phi\colon L \to J(P(L))$ is a lattice isomorphism.

Since the elements in the factorization of x form the antichain of maximal elements in $\phi(x)$, we have $\bigvee \phi(x) = x$. Hence ϕ is injective.

For surjectivity, let D be a down-set in $P(L)$, and let $x = \bigvee D$. Each element of D is bounded above by x, so $D \subseteq \phi(x)$. Since $x = \bigvee D$, we also have $x = \bigvee A$, where A is the antichain of maximal elements in D. An expression of x as a join of an antichain of irreducible elements is irredundant, so Lemma 12.4.27 and Lemma 12.4.29 imply that $D = \phi(x)$.

If $x \leq y$, then $\phi(x) \subseteq \phi(y)$, and $\phi(x) \subseteq \phi(y)$ implies $x = \bigvee \phi(x) \leq \bigvee \phi(y) = y$, so ϕ is a poset isomorphism. For a lattice isomorphism, it remains only to check that ϕ preserves meets and joins. For meets,

$$\begin{aligned}\phi(x) \wedge_{J(P(L))} \phi(y) &= \phi(x) \cap \phi(y) = \{z \in P(L)\colon\ z \leq x \text{ and } z \leq y\}\\ &= \{z \in P(L)\colon\ z \leq x \wedge_L y\} = \phi(x \wedge_L y)\end{aligned}$$

The computation for joins is analogous. ■

12.4.31. Corollary. (1) Every distributive lattice L is graded, with $r(L) = |P(L)|$.
(2) Under ϕ, the irreducible elements of L map to members of $J(P(L))$ generated by single elements.

Proof: (1) $J(P(L))$ is ranked by cardinality of the down-sets in $P(L)$ (for example, compare L and $P(L)$ in the drawing of Example 12.4.24).

(2) The factorization of an irreducible element $p \in P(L)$ is $\{p\}$. ■

Writing this discussion using meets instead of joins would prove $P(L) \cong Q(L)$; Exercise 38 obtains this statement from Theorem 12.4.30.

12.4.32.* Remark. Like our characterization of semiorders by first characterizing interval orders, Theorem 12.4.22 characterizing distributive lattices by forbidden sublattices is proved by first characterizing a larger class.

A lattice is **modular** if equality always holds in the inequality of Lemma 12.4.14(e). This fails for N_5 with x, y, z as in Example 12.4.15. On the other hand, it holds for M_5. This is the source of the notation: M_5 is modular, and N_5 is non-modular. Using various properties of modular lattices, one can show that a lattice is modular if and only if it does not have N_5 as a sublattice.

Every distributive lattice is modular: when $z \leq x$ in a distributive lattice, $x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z) = (x \wedge y) \vee z$. Theorem 12.4.22 is proved by showing that a modular lattice is distributive if and only if it does not have M_5 as a sublattice.

Modular lattices lie in a still larger class. A lattice is **semimodular** if $x \vee y$ covers y whenever x covers $x \wedge y$. Semimodular lattices are characterized by having submodular rank functions, meaning $r(x \wedge y) + r(x \vee y) \le r(x) + r(y)$. This links them closely to matroids (Section 11.3); in any matroid, the inclusion order on the family of closed sets is a semimodular lattice. ∎

CORRELATIONAL INEQUALITIES

A natural probability space arises from linear extensions of posets. We view a poset Q as partial information about an underlying linear order on the elements. This indeed is the setting when we have partially sorted numbers via pairwise comparisons. The set S of possible outcomes is the set of linear extensions of Q.

We assume that each linear extension of Q is equally likely to be the true ordering. An **event** is a subset of the linear extensions. When we specify an event A by a condition (like "$x < y$"), we mean the set of linear extensions in which the condition occurs. The probability $\mathbb{P}(A)$ of an event A is then $|A|/|S|$.

Events A and B are **independent** if $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B)$. Events A and B are **positively correlated** if $\mathbb{P}(AB) \ge \mathbb{P}(A)\mathbb{P}(B)$. Our goal is the "XYZ Inequality" (Theorem 12.4.39): for elements x, y, z in any poset Q, the events "$x < y$" and "$x < z$" are positively correlated.

Sampling randomly from a poset raises similar questions. When an element is chosen randomly from P, positive correlation for "membership in F" and "membership in G" means $\frac{|F \cap G|}{|P|} > \frac{|F|}{|P|}\frac{|G|}{|P|}$. Kleitman [1966] proved this for down-sets F and G in $\mathbf{2}^n$, by proving equivalently that membership in a down-set and an up-set are negatively correlated. He was motivated by an extremal problem: how large can $|F \cap G|$ be for a down-set F and an up-set G of specified sizes?

Later, Anderson [1976] proved that Kleitman's Inequality also holds in the multiset lattice M^e. We present the proof in Daykin–Kleitman–West [1979], because it involves Chebyshev's Inequality (a correlational inequality for real numbers) and illustrates an inductive technique for chain-products.

12.4.33. Lemma. (**Chebyshev's Inequality**) If $x_1, \ldots, x_m$ and $y_1, \ldots, y_m$ are both nonincreasing or both nondecreasing sequences, then

$$\frac{\sum x_i y_i}{m} \ge \frac{\sum x_i}{m}\frac{\sum y_i}{m}.$$

Furthermore, if μ is a nonnegative weight function on $[m]$, then

$$\frac{\sum x_i y_i \mu(i)}{\sum \mu(i)} \ge \frac{\sum x_i \mu(i)}{\sum \mu(i)}\frac{\sum y_i \mu(i)}{\sum \mu(i)}.$$

Proof: Form the double sum $\sum_{i,j}(x_i - x_j)(y_i - y_j)$. Since both sequences are monotone, both factors in any term have the same sign (or 0), so each term is nonnegative. Multiplying out and moving the terms expressed with minus signs to the other side yields $2m \sum x_i y_i \ge 2 \sum x_i \sum y_j$. For the weighted version, start with $\sum_{i,j}(x_i - x_j)(y_i - y_j)\mu(i)\mu(j)$ and proceed in the same way. ∎

We use this to prove Kleitman's Inequality for chain-products.

12.4.34. Theorem. (Anderson [1976], Daykin–Kleitman–West [1979]) If F and G are down-sets in L, a product of n chains, then

$$|L|\,|F \cap G| \geq |F|\,|G|.$$

Proof: Let $e_1, \ldots, e_n$ be the chain sizes. We use induction on n. If $n = 1$, then $|F \cap G| = \min\{|F|, |G|\}$ and $|L| \geq \max\{|F|, |G|\}$.

For $n > 1$, partition these sets using the last coordinate. Each $x \in L$ is a vector $(x_1, \ldots, x_n)$ with $0 \leq x_j < e_j$ for $j \in [n]$. For $X \subseteq L$ and $0 \leq i \leq e_n - 1$, let $X_i = \{x \in X : x_n = i\}$. Note that L_i is isomorphic to the product L' of the first $n - 1$ chains. If also X is a down-set in L, then X_i is a down-set in L_i, and also $|X_0| \geq \cdots \geq |X_{e_n-1}|$.

Since F, G, and $F \cap G$ are down-sets, applying this chain of inequalities yields nonincreasing lists for the sizes of their "slices" $|F_i|$, $|G_i|$, and $|(F \cap G)_i|$. Also, $(F \cap G)_i = F_i \cap G_i$. Now we apply Chebyshev's Inequality and then the induction hypothesis (for the subsets of L' obtained by deleting the last coordinate from the elements of F_i, G_i, and $F_i \cap G_i$). The computation is

$$\begin{aligned} |F|\,|G| &= \sum |F_i| \sum |G_i| \leq e_n \sum |F_i|\ |G_i| \leq e_n \sum |L_i|\,|F_i \cap G_i| \\ &= e_n\Big(\prod_{i<n} e_i\Big) \sum |(F \cap G)_i| = |L|\,|F \cap G|. \end{aligned}$$

■

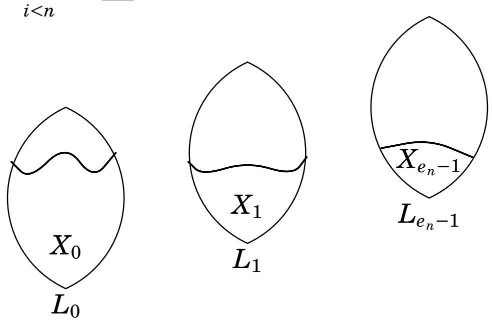

Daykin [1977] proved a more general inequality for more general sets in distributive lattices and thereby characterized distributive lattices (Exercise 47).

As a common extension of Theorem 12.4.34 and Chebyshev's Inequality, we will prove the FKG Inequality for monotone functions on lattices. Discovered jointly by Fortuin, Kasteleyn, and Ginibre, this is the central result about correlational inequalities.

Chebyshev's Inequality can be viewed as a statement about random variables. The weight function $\mu(i)$ gives the probability that the outcome is i. The sequences x and y are monotone functions $f(i)$ and $g(i)$, and we compare the expectation of their product and the product of their expectations. To extend this to distributive lattices, we need a technical condition on the weight function μ.

12.4.35. Definition. A function $f \colon P \to \mathbb{R}$ is **order-preserving** if $x \leq y$ implies $f(x) \leq f(y)$, **order-reversing** if $x \leq y$ implies $f(x) \geq f(y)$. A real-valued function μ on a lattice is **log-supermodular** if $\mu(x \wedge y)\mu(x \vee y) \geq \mu(x)\mu(y)$.

All weight functions on chains are log-supermodular, since in a chain always $\{x \wedge y, x \vee y\} = \{x, y\}$. Chebyshev's Inequality is the special case of the FKG Inequality where the lattice is a chain.

12.4.36. Theorem. (**FKG Inequality**; Fortuin–Kasteleyn–Ginibre [1971]) In a distributive lattice L, functions f and g that are both order-preserving or both order-reversing are positively correlated with respect to each nonnegative log-supermodular weight function μ, meaning that

$$\sum f(x)g(x)\mu(x) \sum \mu(x) \geq \sum f(x)\mu(x) \sum g(x)\mu(x). \qquad \blacksquare$$

The FKG Inequality yields Kleitman's Inequality for all distributive lattices by setting $\mu = 1$ and letting f and g be the (order-reversing) characteristic functions on the down-sets F and G, where the **characteristic function** of a set A is the function χ_A having value 1 on A and value 0 outside A. Nevertheless, generalizations of Kleitman's Inequality continued to appear long after the FKG Inequality, because the FKG Inequality appeared in the literature of statistical mechanics and discrete mathematicians were unaware of it for years.

Later, Ahlswede and Daykin found a generalization of the FKG Inequality having an easier inductive proof. We give this proof, following the presentation of Graham [1982]. The theorem is known both as the **Ahlswede–Daykin Inequality** and as the **Four Function Inequality**. A result intermediate between the FKG Inequality and the Four Function Inequality appeared in Holley [1974].

12.4.37. Definition. For subsets X and Y of a lattice L, define $X \wedge Y = \{x \wedge y \colon x \in X, y \in Y\}$ and $X \vee Y = \{x \vee y \colon x \in X, y \in Y\}$. Also define $f(X) = \sum_{x \in X} f(x)$ when $f \colon L \to \mathbb{R}$.

12.4.38. Theorem. (Ahlswede–Daykin [1978]) If $\alpha, \beta, \gamma, \delta$ are four nonnegative functions on a distributive lattice L such that

$$\alpha(x)\beta(y) \leq \gamma(x \wedge y)\delta(x \vee y) \text{ for all } x, y \in L,$$

then

$$\alpha(X)\beta(Y) \leq \gamma(X \wedge Y)\delta(X \vee Y) \text{ for all } X, Y \subseteq L.$$

Proof: We first reduce to the case of $\mathbf{2}^n$, using that distributive lattices are sublattices of such posets (Theorem 12.4.30). If the claim holds for $\mathbf{2}^n$, then it holds for a sublattice L of $\mathbf{2}^n$ as follows. Given $\alpha, \beta, \gamma, \delta$ defined on L, extend them to $\mathbf{2}^n$ by giving them value 0 on $\mathbf{2}^n - L$. Whenever $\alpha(x)\beta(y) \neq 0$, we have $x, y \in L$. Since L is a sublattice, L also contains $x \wedge y$ and $x \vee y$, and the hypothesis for $\mathbf{2}^n$ follows from its truth for L. Hence the conclusion holds for any $X, Y \subseteq \mathbf{2}^n$, including when $X, Y \subseteq L$.

The proof for $\mathbf{2}^n$ uses induction on n. For $n = 1$, we check several cases. The elements of L are 0 and 1 (representing $\varnothing$ and $[1]$). The hypothesis gives $\alpha(x)\beta(y) \leq \gamma(\min\{x, y\})\delta(\max\{x, y\})$ for the four possibilities of $x, y \in \{0, 1\}$. The conclusion is easy when $|X| = 1$ or $|Y| = 1$. The case $X = Y = \{0, 1\}$ requires a numerical optimization (Exercise 49).

For $n > 1$, consider fixed $X, Y \subseteq \mathbf{2}^n$. Let $L' = \mathbf{2}^{n-1}$. In terms of X and Y, we define $\alpha', \beta', \gamma', \delta'$ on L' so that the desired inequality $\alpha(X)\beta(Y) \le \gamma(X \wedge Y)\delta(X \vee Y)$ will become $\alpha'(L')\beta'(L') \le \gamma'(L' \wedge L')\delta(L' \vee L')$. Note that $L' \wedge L' = L' = L' \vee L'$ in the lattice L'. For $x \in L'$, we define

$$\begin{aligned}
\alpha'(x) &= \alpha(x)\chi_X(x) + \alpha(x \cup \{n\})\chi_X(x \cup \{n\})\\
\beta'(x) &= \beta(x)\chi_Y(x) + \beta(x \cup \{n\})\chi_Y(x \cup \{n\})\\
\gamma'(x) &= \gamma(x)\chi_{X\wedge Y}(x) + \gamma(x \cup \{n\})\chi_{X\wedge Y}(x \cup \{n\})\\
\delta'(x) &= \delta(x)\chi_{X\vee Y}(x) + \delta(x \cup \{n\})\chi_{X\vee Y}(x \cup \{n\}).
\end{aligned}$$

For $t \in \{\alpha, \beta, \gamma, \delta\}$, this definition accumulates in $t'(x)$ the contributions of x and $x \cup \{n\}$ to the "relevant set" Z, which is X, Y, $X \wedge Y$, or $X \vee Y$ when t is α, β, γ, or δ, respectively. Including all contributions, $t'(L') = t(Z)$, and thus the desired inequality becomes $\alpha'(L')\beta'(L') \le \gamma'(L')\delta(L')$.

To obtain this conclusion from the induction hypothesis, we must show that these functions $\alpha', \beta', \gamma', \delta'$ satisfy $\alpha'(x)\beta'(y) \le \gamma'(x \wedge y)\delta'(x \vee y)$ for all $x, y \in \mathbf{2}^{n-1}$. For fixed x and y, each quantity in the desired inequality is computed from two elements of L having the form z and $z \cup \{n\}$. Since each of $\{z, z \cup \{n\}\}$ might or might not belong to the relevant set Z, we could complete the proof by checking 16 cases.

Alternatively, we can reduce the verification to four easier cases by using the induction hypothesis for $n = 1$. We split the contributions to t' by defining new functions on $\mathbf{2}^1$.

$$\begin{array}{lll}
z = x: & \alpha''(\varnothing) = \alpha(z)\chi_X(z) & \alpha''([1]) = \alpha(z \cup \{n\})\chi_X(z \cup \{n\})\\
z = y: & \beta''(\varnothing) = \beta(z)\chi_Y(z) & \beta''([1]) = \beta(z \cup \{n\})\chi_Y(z \cup \{n\})\\
z = x \wedge y: & \gamma''(\varnothing) = \gamma(z)\chi_{X\wedge Y}(z) & \gamma''([1]) = \gamma(z \cup \{n\})\chi_{X\wedge Y}(z \cup \{n\})\\
z = x \vee y: & \delta''(\varnothing) = \delta(z)\chi_{X\vee Y}(z) & \delta''([1]) = \delta(z \cup \{n\})\chi_{X\vee Y}(z \cup \{n\}).
\end{array}$$

For $t \in \{\alpha, \beta, \gamma, \delta\}$, we have defined t'' on $\{\varnothing, [1]\}$ so that $t''(\mathbf{2}^1) = t'(z)$, where z is the "relevant element" of $\mathbf{2}^{n-1}$ as listed above.

If $\alpha'', \beta'', \gamma'', \delta''$ satisfy $\alpha''(u)\beta''(v) \le \gamma''(u \wedge v)\delta''(u \vee v)$ for all $u, v \in \mathbf{2}^1$, then the induction hypothesis for the case $n = 1$ yields $\alpha''(\mathbf{2}^1)\beta''(\mathbf{2}^1) \le \gamma''(\mathbf{2}^1)\delta''(\mathbf{2}^1)$, which is the needed inequality $\alpha'(x)\beta'(y) \le \gamma'(x \wedge y)\delta'(x \vee y)$ for elements of $\mathbf{2}^{n-1}$. Using the definition of each t'', the inequalities for $u, v \in \mathbf{2}^1$ are now statements about elements of L, and we can apply the original hypothesis about $\{\alpha, \beta, \gamma, \delta\}$.

Since $\alpha, \beta, \gamma, \delta$ are nonnegative, in checking the hypothesis on $\mathbf{2}^1$ we may assume that the left side of the needed inequality is positive, which requires that the arguments on the left belong to X and Y, respectively. The corresponding arguments to γ and δ then belong to $X \wedge Y$ and $X \vee Y$, respectively, so all the values of the characteristic functions in the definitions of t'' may be assumed to be 1. For $x, y \in L'$, evaluating t'' at the relevant elements yields the four needed inequalities below.

u	v			
$\varnothing$	$\varnothing$	$\alpha(x)\beta(y)$	$\le$	$\gamma(x \wedge y)\delta(x \vee y)$
$\varnothing$	$[1]$	$\alpha(x)\beta(y \cup \{n\})$	$\le$	$\gamma(x \wedge y)\delta((x \vee y) \cup \{n\})$
$[1]$	$\varnothing$	$\alpha(x \cup \{n\})\beta(y)$	$\le$	$\gamma(x \wedge y)\delta((x \vee y) \cup \{n\})$
$[1]$	$[1]$	$\alpha(x \cup \{n\})\beta(y \cup \{n\})$	$\le$	$\gamma((x \wedge y) \cup \{n\})\delta((x \vee y) \cup \{n\})$

In each nontrivial case, $\alpha''(u)\beta''(v) \le \gamma''(u \wedge v)\delta''(u \vee v)$ thus reduces to an instance of the hypothesis given on $\mathbf{2}^n$. ■

The Ahlswede–Daykin Inequality now yields the FKG Inequality.

Proof of FKG Inequality (Theorem 12.4.36)**:** Set $\alpha = \beta = \gamma = \delta = \mu$. By negating the functions if necessary, we may assume that f and g are order-reversing. Also, adding a constant to one of the function does not affect the inequality, so we may assume that f and g are nonnegative.

We first prove the FKG Inequality when f and g are the characteristic functions of down-sets X and Y. In this case, $\sum_{x\in L} f(x)\mu(x) = \sum_{x\in X} \mu(x) = \mu(X)$, and similarly $\sum_{x\in L} g(x)\mu(x) = \mu(Y)$. The hypothesis that μ is log-supermodular implies the hypothesis for the Ahlswede–Daykin Inequality, and the conclusion of the Ahlswede–Daykin Inequality for X, Y is the desired statement for f and g:

$$\mu(X)\mu(Y) \le \mu(X \wedge Y)\mu(X \vee Y) \le \mu(X \cap Y)\mu(L) = \sum f(x)g(x)\mu(x) \sum \mu(x).$$

To complete the proof, we need only show (1) every nonnegative monotone nonincreasing function on L is a linear combination of at most $|L|$ characteristic functions on down-sets, and (2) the FKG Inequality is preserved by taking a positive linear combination of order-reversing functions f_1 and f_2.

(1) follows by induction on the number of nonzero values of f, with a trivial basis. Let ε be the least nonzero value of f, let $S = \{x \in L\colon\ f(x) \ne 0\}$, and let χ_S be the characteristic function of S. Now $f - \varepsilon\chi_S$ is a nonnegative nonincreasing function on L with fewer nonzero values than f.

For (2), let f be a positive linear combination of f_1 and f_2, which each satisfy the FKG Inequality with g. Term-by-term linearity of real number arithmetic yields the FKG Inequality for f and g. ■

The FKG Inequality was notably applied to prove a conjecture by Rival and Sands about positive correlation in random linear extensions of a poset Q. Recall that $\mathbb{P}(A)$ is the fraction of the linear extensions in which event A occurs. For example, when Q is an antichain, $\mathbb{P}(x < y) = \frac{1}{2}$ for any x and y. If $Q = \mathbf{2}+\mathbf{1}$, with $x < z$ and y unrelated to both, then $\mathbb{P}(x < y) = \mathbb{P}(y < z) = \frac{2}{3}$ and $\mathbb{P}(x < z) = 1$.

12.4.39. Theorem. (**XYZ Inequality**; Shepp [1982]) For any elements x, y, z in any poset Q, the events "$x < y$" and "$x < z$" are positively correlated.

Proof: Roughly speaking, we seek a distributive lattice where the events of interest correspond to down-sets and will be positively correlated. Let the elements of Q be $\{q_1, \dots, q_n\}$, with $(q_1, q_2, q_3) = (x, y, z)$. Henceforth, we use x and y instead as elements of the eventual distributive lattice, as in Theorem 12.4.36.

Let N be a large positive integer, and write $x \in [N]^n$ as $(x_1, \dots, x_n)$. View x as a map $\lambda\colon Q \to [N]$ in which $x_i = \lambda(q_i)$. If the coordinates of x are distinct (λ is injective), then x yields an ordering of Q. There are $\binom{N}{n}$ such vectors x for each ordering. When λ is order-preserving, the ordering is a linear extension.

Using $[N]^n$ rather than the extensions themselves makes it easier to define a distributive lattice. We choose N large to make the effect of non-injective mappings negligible. We return to this detail later; for now we view the order-preserving mappings in $[N]^n$ as corresponding to extensions of Q.

Application of the FKG Inequality needs a weight function; let μ be the characteristic function of the set of order-preserving mappings. So that f and

g capture the desired events, let them be the characteristic functions of F and G, where $F = \{x \in [N]^n \colon x_1 < x_2\}$ and $G = \{x \in [N]^n \colon x_1 < x_3\}$. Roughly,

$$\sum f(x)\mu(x) / \sum \mu(x) \approx \mathbb{P}(q_1 < q_2),$$
$$\sum g(x)\mu(x) / \sum \mu(x) \approx \mathbb{P}(q_1 < q_3),$$
$$\sum f(x)g(x)\mu(x) / \sum \mu(x) \approx \mathbb{P}(q_1 < q_2 \text{ and } q_1 < q_3).$$

To apply the FKG Inequality, we need a partial order on $[N]^n$ so that f and g are order-preserving, μ is log-supermodular, and the poset is a distributive lattice. For $x, y \in [N]^n$, set $x \le y$ when $x_1 \ge y_1$ and $x_i - x_1 \le y_i - y_1$ for $i > 1$.

We have $f(x) = 1$ when $x_1 \le x_2$. If $x \le y$, then also $y_1 \le y_2 - x_2 + x_1 \le y_2$, so $f(y) = 1$. Thus f is monotone increasing, and similarly for g.

To see that $[N]^n$ is a distributive lattice under this ordering, consider the map $\sigma(x) = (-x_1, x_2 - x_1, \ldots, x_n - x_1)$. We have $x \le y$ in our ordering if and only if $\sigma(x) \le \sigma(y)$ in the usual ordering on a product of n chains of size $2N + 1$, indexed from $-N$ to N. Meets and joins are preserved under σ, which follows from

$$(x \wedge y)_i = \min(x_i - x_1, y_i - y_1) + \max(x_1, y_1)$$
$$(x \vee y)_i = \max(x_i - x_1, y_i - y_1) + \min(x_1, y_1).$$

For example, if $i > 1$, then

$$[\sigma(x) \wedge \sigma(y)]_i = \min(x_i - x_1, y_i - y_1) = (x \wedge y)_i - (x \wedge y)_1 = [\sigma(x \wedge y)]_i.$$

Thus our poset is a sublattice of a chain-product and hence a distributive lattice.

To apply the FKG Inequality, we also need μ to be log-supermodular. We defined $\mu(x) = 1$ when x is order-preserving ($x_i \le x_j$ for $q_i \le q_j$). We need only check that $x \wedge y$ and $x \vee y$ are order-preserving when x and y are. This is easy from the formulas given above for $x \wedge y$ and $x \vee y$, and it is the reason for subtracting the first coordinate from the others in defining the order relation on $[N]^n$.

From the FKG Inequality, $\sum f(x)g(x)\mu(x) \sum \mu(x) \ge \sum f(x)\mu(x) \sum g(x)\mu(x)$, where the sums run over $x \in [N]^n$. From our definitions of f, g, μ, this is what we want, except that some vectors are not injective. We show that the relative contribution to each term from non-injective x tends to 0 as $N \to \infty$.

Let $A_1 \subset [N]^n$ be the set of injective n-tuples (distinct coordinate values), and let $A_2 = [N]^n - A_1$. For $i \in \{1, 2\}$, define

$$F_i = \textstyle\sum_{x \in A_i} f(x)\mu(x), \qquad G_i = \sum_{x \in A_i} g(x)\mu(x),$$
$$H_i = \textstyle\sum_{x \in A_i} f(x)g(x)\mu(x), \quad M_i = \sum_{x \in A_i} \mu(x).$$

Note that M_1 is $\binom{N}{n}$ times the number of extensions of Q. There is at least one extension, so $M_1 \ge \binom{N}{n}$. On the other hand, $M_2 \le |A_2| = N^n - N_{(n)}$. With n fixed, we have $M_1 \sim N^n/n!$ and $M_2 = O(N^{n-1})$, so $M_2/M_1 \to 0$. Since F_i and G_i and H_i are all bounded by M_i, also F_2/M_1 and G_2/M_1 and H_2/M_1 all tend to 0.

Since each linear extension is represented equally often in A_1, what we want $H_1 M_1 \ge F_1 G_1$. The FKG Inequality yields

$$(H_1 + H_2)(M_1 + M_2) \ge (F_1 + F_2)(G_1 + G_2).$$

Dividing by M_1^2 and taking the limit as $N \to \infty$ yields $\frac{H_1}{M_1} \ge \frac{F_1}{M_1}\frac{G_1}{M_1}$, which is precisely the desired inequality. ∎

The XYZ Inequality is quite special. Many similar-sounding statements are not true. For example, given $x, y, z, w \in Q$, one might expect that the events $x < y < w$ and $x < z < w$ are positively correlated, but this does not hold for all Q (Exercise 51). Winkler [1986] characterized the very few pairs of subposets that are always positively correlated. Brightwell–Trotter [1999] gives a combinatorial proof of a stronger version of the XYZ Inequality due to Fishburn [1984].

A PROBLEM IN RAMSEY THEORY (optional)

We close this chapter by relating the lattice of down-sets of a poset to a problem in ordered Ramsey theory. No aspects of lattice theory are needed here, just the definition of the poset of down-sets.

12.4.40. Definition. An **ordered hypergraph** is a hypergraph on a linearly ordered vertex set. An ordered hypergraph G occurs as a **subhypergraph** of an ordered hypergraph H if some order-preserving injection from $V(G)$ to $V(H)$ also preserves edges. We then say that H **contains** (a copy of) G.

Let K_n^r denote the complete r-uniform hypergraph with n vertices. By Ramsey's Theorem, when n is sufficiently large every k-coloring of $E(K_n^r)$ contains a monochromatic copy of K_p^r and hence also of every subhypergraph of K_p^r. Because any two complete ordered hypergraphs with the same number of vertices are isomorphic, the same statement holds in the ordered sense, which leads to a well-defined ordered version of Ramsey numbers.

12.4.41. Definition. Let $G_1, \ldots, G_k$ be r-uniform ordered hypergraphs. The k**-color Ramsey number** $R(H_1, \ldots, H_k)$ is the least n such that every k-coloring of an n-vertex complete r-uniform ordered hypergraph contains a copy of H_i in color i for some i. The r**-uniform monotone path** P_t^r is the t-vertex ordered hypergraph whose edges are the sets of r consecutive vertices.

When discussing only ordered hypergraphs, we can use the same notation and terminology as for classical Ramsey numbers. Indeed, when $G_1, \ldots, G_k$ are complete, the Ramsey numbers in the classical and ordered senses are the same.

However, already when $r = 2$ the Ramsey numbers for ordinary paths and for monotone paths differ greatly. For the path with three vertices, the classical k-color Ramsey number is $k + 2$ (among $k + 1$ edges at a vertex, some two must have the same color), but for the 2-uniform monotone path with three vertices the Ramsey number is $2^k + 1$. This generalizes as follows, proved by the same idea as the Erdős–Szekeres Theorem (Theorem 10.1.13).

12.4.42. Proposition. For ordered paths, $R(P_{t_1}^2, \ldots, P_{t_k}^2) = 1 + \prod_{i=1}^{k}(t_i - 1)$.

Proof: Fix a k-coloring of $E(K_n)$. For $x \in V(K_n)$ and $i \in [k]$, let x_i be the maximum number of vertices in a monotone increasing path with color i ending at x. Note that $x_i \geq 1$ for all x and i, and avoiding long paths in color i requires $x_i \leq t_i - 1$.

If y comes before z and edge yz has color i, then $z_i > y_i$, so no two k-tuples can be the same. If there is no sufficiently long path, then only $\prod_{i=1}^{k}(t_i - 1)$ vectors are available. Hence with larger n we have a monotone path in color i for some i.

For the lower bound, associate with the ℓth vertex the ℓth k-tuple in a linear extension of the product of chains with sizes $t_1 - 1, \ldots, t_k - 1$. When y comes before z, there is a coordinate i such that $y_i < z_i$; give the edge yz color i. Since coordinate i must increase along any monotone path in color i, no monotone path in color i has t_i vertices. ∎

This proof is actually the first instance of the induction step in a characterization of $R(P^r_{t_1}, \ldots, P^r_{t_k})$ as the number of down-sets in an appropriate poset, leading to upper and lower bounds that are exponential towers of height $r-2$ for the Ramsey number. Letting $R_k(P^r_t) = R(P^r_{t_1}, \ldots, P^r_{t_k})$, the problem of computing $R_k(P^r_t)$ was introduced in Fox–Pach–Sudakov–Suk [2012] with a geometric application, though it had actually been considered by Duffus–Lefmann-Rödl [1995] in the language of shift graphs (Definition 12.3.36) in order to give a lower bound on classical Ramsey numbers. Recall that $J(Q)$ denotes the lattice of down-sets in a poset Q, ordered by inclusion.

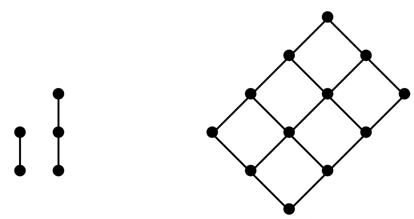

Q $\qquad$ $J(Q)$

12.4.43. Theorem. (Moshkovitz–Shapira [2014]; see also Milans–Stolee–West [2015]) Fix $t_1, \ldots, t_k \in \mathbb{N}$, all larger than r. Let Q be the poset consisting of disjoint chains with sizes $t_1 - r, \ldots, t_k - r$. With $Q_1 = Q$ and $Q_i = J(Q_{i-1})$ for $i > 1$, the Ramsey numbers for r-uniform monotone paths satisfy $R(P^r_{t_1}, \ldots, P^r_{t_k}) = |Q_r| + 1$.

Proof: The case $r = 1$ uses the pigeonhole principle. We have $|Q_1| = \sum_{i=1}^{k}(t_i - 1)$, and any k-coloring of 1-sets chosen from $1 + \sum_{i=1}^{k}(t_i - 1)$ vertices has t_i vertices with color i for some i, forming a monotone copy of $P^1_{t_i}$ in color i. We proceed by induction on r. (The induction step actually simplifies to the proof of Proposition 12.4.42 when $r = 2$.)

Lower Bound. Let $n = |Q_r|$; we construct a k-edge-coloring of K^r_n that avoids $P^r_{t_i}$ in color i for each i. First consider $x, y \in Q_j$, where $2 \le j \le r$. The meaning of $x \not\ge y$ is that x does not contain y when they are viewed as down-sets in Q_{j-1}. In this case, let $f(x, y)$ be a fixed element of the family $y - x$ in Q_{j-1}. For $x, y, z \in Q_j$, if $x \not\ge y$ and $y \not\ge z$, then y (in Q_{j-1}) contains $f(x, y)$ but not $f(y, z)$. Since y is a down-set in Q_{j-1}, we obtain $f(x, y) \not\ge f(y, z)$ in Q_{j-1}.

A list $x_1, \ldots, x_s$ of elements in a poset is *descent-free* if $x_i \not\ge x_{i+1}$ for $1 \le i \le s - 1$. We extend f to descent-free lists in Q_j by setting $f(x_1, \ldots, x_s) = (f(x_1, x_2), \ldots, f(x_{s-1}, x_s))$. Thus the image under f of a descent-free s-list in Q_j is a descent-free $(s-1)$-list in Q_{j-1}. Let f^0 be the identity map. For a descent-free s-list $x_1, \dots, x_s$ in Q_j, where $s, j > d > 1$, define $f^d(x_1, \dots, x_s) = f(f^{d-1}(x_1, \dots, x_s))$; now $f^d(x_1, \dots, x_s)$ is a descent-free $(s-d)$-list in Q_{j-d}.

Let $y_1, \dots, y_n$ be a linear extension of Q_r, so $y_i \le y_j$ implies $i \le j$. Each sublist of a linear extension is descent-free. For a sublist $x_1, \dots, x_r$, note that $f^{r-1}(x_1, \dots, x_r)$ is a single element in Q. Color $\{x_1, \dots, x_r\}$ (as an edge in K_n^r) with the index of the chain in Q that contains $f^{r-1}(x_1, \dots, x_r)$.

Suppose that this coloring has a monotone copy of P_s^r in color i. Let $x_1, \dots, x_s$ be its vertices, in increasing order. Since $x_1, \dots, x_s$ is a sublist of a linear extension, it is descent-free. Since each set of r consecutive vertices from the list $x_1, \dots, x_s$ was give color i, $f^{r-1}(x_1, \dots, x_s)$ is a descent-free $(s-r+1)$-list in the ith chain of Q. A descent-free list in a chain is strictly increasing (because equality is also considered a descent), so $s - r + 1 \le t_i - r$. Thus $s < t_i$. We conclude that the coloring avoids $P_{t_i}^r$ in color i for each i, so $R(P_{t_1}^r, \dots, P_{t_k}^r) > n$.

Upper bound. Given a k-edge-coloring ϕ of $E(K_n^r)$ with vertex set $[n]$ that avoids $P_{t_i}^r$ in color i for each i, it suffices to define an injection from $[n]$ to Q_r. View each vertex subset $Y \subseteq [n]$ as an increasing list, with $Y^- = Y - \max Y$ and $Y^+ = Y - \min Y$.

For $1 \le j \le r < n$, we construct $g_j \colon \binom{[n]}{j} \to Q_{r-j+1}$ such that

$$g_j(Y^-) \not\ge g_j(Y^+) \text{ in } Q_{r-j+1} \text{ when } Y \in \tbinom{[n]}{j+1}. \qquad (*)$$

This suffices, since g_1 will then be the desired injection.

We first define g_r. For $X \in \binom{[n]}{r}$, let $i = \phi(X)$, and let $w_1, \dots, w_{t_i - r}$ be the elements of the ith chain in Q. Set $g_r(X) = w_h$, where h is the largest integer such that some copy of P_{h+r-1}^r in color i has last edge X. Note that $h \le t_i - r$, since ϕ has no copy of $P_{t_i}^r$ in color i. If $\phi(Y^-) = \phi(Y^+)$ for some $Y \in \binom{[n]}{r+1}$, then $g_r(Y^+) > g_r(Y^-)$, and otherwise $g_r(Y^+)$ and $g_r(Y^-)$ are incomparable in Q. In either case, $g_r(Y^-) \not\ge g_r(Y^+)$.

Now consider smaller j, with $g_{j+1} \colon \binom{[n]}{j+1} \to Q_{r-j}$ already defined and satisfying $(*)$. For $X \in \binom{[n]}{j}$, let the *precursors* of X be the $(j+1)$-sets obtained from X by adding an element smaller than all of X; that is, $\{Z \in \binom{[n]}{j+1} \colon Z^+ = X\}$. (Note that sets containing the vertex 1 have no precursors.) Define $g_j(X)$ to be the down-set in Q_{r-j} generated by the set of elements $g_{j+1}(Z)$ such that Z is a precursor of X. Since $Q_{r-j+1} = J(Q_{r-j})$, by definition $g_j(X) \in Q_{r-j+1}$.

To check $(*)$ for g_j, consider $Y \in \binom{[n]}{j+1}$. Since Y is a precursor of Y^+, the definition of g_j yields $g_{j+1}(Y) \in g_j(Y^+)$. If $g_{j+1}(Y)$ also lies in $g_j(Y^-)$, then by the definition of $g_j(Y^-)$ (using $X = Y^-$), the element $g_{j+1}(Y)$ in Q_{r-j} lies below some element $g_{j+1}(Z)$ such that $Z^+ = Y^-$. Letting $W = Z \cup Y$, we have $Z = W^-$ and $Y = W^+$. Thus $g_{j+1}(W^-) \ge g_{j+1}(W^+)$, which contradicts $(*)$ for g_{j+1}.

We conclude $g_{j+1}(Y) \in g_j(Y^+) - g_j(Y^-)$. Hence the down-set $g_j(Y^+)$ in Q_{r-j} does not contain the down-set $g_j(Y^-)$, which means $g_j(Y^-) \not\ge g_j(Y^+)$ in Q_{r-j+1}. ■

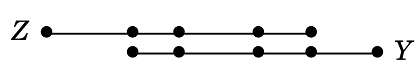

Theorem 12.4.43 inductively yields upper and lower bounds for $R_k(P_t^r)$. For P_4^3, note that Q_2 is the lattice of subsets of $[k]$, and thus computing $R_k(P_4^3)$ is equivalent to the well-known Dedekind's Problem discussed in Theorem 12.2.14.

12.4.44. Corollary. (Moshkovitz–Shapira [2014]) Let $m = t - r + 1$, and let $\mathrm{tow}_h(x)$ be x when $h = 0$ and $2^{\mathrm{tow}_{h-1}(x)}$ when $h \ge 1$. For the r-uniform monotone ordered path P_t^r,

$$\mathrm{tow}_{r-2}(m^{k-1}/2\sqrt{k}) \le R_k(P_t^r) \le \mathrm{tow}_{r-2}(2m^{k-1}),$$

Proof: We prove weaker bounds, still exponential towers of the same height (as in Milans–Stolee-West [2015]). We seek bounds on $|Q_r|$ of Theorem 12.4.43.

As noted above, $|Q_2| = m^k$. Since the elements of Q_j consist of subsets of Q_{j-1}, we have $|Q_j| \le 2^{|Q_{j-1}|}$, and hence $|Q_r| \le \mathrm{tow}_{r-2}(m^k)$.

For the lower bound, recall that $|J(Q)|$ also counts the antichains in Q; each down-set has an antichain of maximal elements. When a poset contains an antichain A, its subsets form $2^{|A|}$ antichains. Since $|Q_1| = k(t-r)$, the sizes of the down-sets in Q_1 range from 0 to $k(t-r)$. Thus Q_2 has fewer than km ranks and has an antichain of size at least m^{k-1}/k. Also, the elements of this antichain have the same size in Q_1. (The chain-product Q_2 is a symmetric chain order, so its middle rank is a largest antichain. Sharper analysis reduces the denominator to $2\sqrt{k}$).

Given an antichain of size M in Q_j, the subsets of size $M/2$ generate pairwise incomparable down-sets and hence an antichain in Q_{j+1}. The resulting lower bound uses an analogue of exponential towers involving middle binomial coefficients. Let $b_0(x) = x$, and for $h \ge 1$ let $b_h(x) = \binom{b_{h-1}(x)}{\lfloor b_{h-1}(x)/2 \rfloor}$. Since $\binom{n}{\lfloor n/2 \rfloor} \sim 2^n/\sqrt{\pi n/2}$ (Example 2.3.10), inductively we have $b_h(x) \ge \mathrm{tow}_h(x - O(\lg x))$.

Let a_j be the maximum size of a family in Q_j consisting of down-sets in Q_{j-1} that have the same size; these sets form an antichain in Q_j. Thus $|Q_{j+1}| \ge 2^{a_j}$, but also $a_{j+1} \ge \binom{a_j}{a_j/2}$. Thus $a_j \ge b_{j-2}(a_2)$. With $a_2 \ge m^{k-1}/k$, we obtain $|Q_r| \ge \mathrm{tow}_{r-2}(m^{k-1}/k - O(k \lg m))$. ■

An application of ordered Ramsey numbers appears in Exercise 53. Another proof of Theorem 12.4.43 follows from a result of Pérez-Giménez–Prałat–West [2018+] for an on-line version of ordered Ramsey numbers.

EXERCISES 12.4

12.4.1. (−) Prove that a poset is a ranking if and only if no subposet is isomorphic to $\mathbf{2} + \mathbf{1}$.

12.4.2. (−) Count the isomorphism classes of lattices with at most five elements.

12.4.3. *A poset is a chain if and only if every subposet is a lattice.* A chain is a lattice, and every subposet of a chain is a chain (pairwise comparable elements). On the other hand, if every subposet is a lattice, then every two-element subposet is a lattice. In particular, this means there is no subposet $\underline{1} + \underline{1}$, because these two elements would not have meet or join in the lattice, so every pair of elements in P is comparable.

12.4.4. (−) Draw the distributive lattice for which the poset below is the poset of join-irreducible elements.

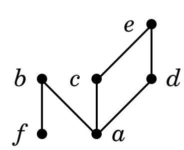

12.4.5. Let P be a poset such that $P - x$ is an interval order. Can relations involving x be added to P to obtain an interval order P' such that $P' - x = P - x$?

12.4.6. *Alternative proof of Theorem 12.4.5.* Let P be a poset without $\mathbf{2} + \mathbf{2}$.

(a) Prove that the sets $U(x)$ for $x \in P$ are linearly ordered by inclusion, and similarly for the sets $D(y)$.

(b) For the 0,1-matrix of the order relation, prove that ordering the rows as $x_1, \ldots, x_n$ in decreasing order of $|U(x_i)|$ and the columns as $y_1, \ldots, y_n$ in decreasing order of $|D(y_j)|$ puts the 1s of the matrix in the positions of a Ferrers diagram.

(c) For $x \in P$, let $u(x)$ count distinct sets of the form $U(y)$ that are proper subsets of $U(x)$, and let $d(x)$ count distinct sets of the form $D(y)$ that are proper subsets of $D(x)$. Let h be the total number of distinct nonempty sets of the form $U(x)$. Prove that assigning each $x \in P$ the interval $[d(x), h - u(x)]$ produces an interval representation of P. (Bogart [1985], motivated by Rabinovitch [1977])

12.4.7. ($\diamond$) Let $\mathbf{I}_n$ be the poset of nontrivial intervals with integer endpoints in $[n]$, ordered by $[a, b] < [c, d]$ if $b < c$. Prove that $\dim \mathbf{I}_n > k$ when n is sufficiently large. (Hint: Given Ramsey's Theorem, use $n \geq R_k(4; 3)$.) (Rabinovitch [1973])

12.4.8. ($\diamond$) Given the interval order $\mathbf{I}_n$ as in Exercise 12.4.7 and the "double-shift graph" G'_n as in Definition 12.3.36, prove $\dim \mathbf{I}_n \geq \chi(G'_n)$. (Comment: Thus $\dim \mathbf{I}_n \geq \lg\lg n + (\frac{1}{2} + o(1)) \lg\lg\lg n$, by Corollary 12.3.39. With some care, Füredi–P.Hajnal–Rödl–Trotter [1992] constructed linear extensions to prove that the lower bound is tight.)

12.4.9. For each $k \in \mathbb{N}$, do the following.

(a) Construct an interval order such that every interval representation uses at least k different lengths of intervals. (Fishburn [1983])

(b) Construct an interval order such that every representation has an interval of length more than k times the length of its shortest interval. (Fishburn–Graham [1985])

12.4.10. A *series-parallel poset* is a poset constructible from isolated elements by using parallel composition (disjoint union) and series composition (combine two posets by putting each element of the first above each element of the second). Prove that a poset is a series-parallel poset if and only if it does not contain the poset $\mathbf{N}$ below.

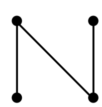

12.4.11. Establish a one-to-one correspondence between the isomorphism classes of n-element semiorders and the ballot lists of length $2n$ (Definition 1.3.16). (Comment: Thus there are $\frac{1}{n+1}\binom{2n}{n}$ isomorphism classes of semiorders on n elements.) (Dean–Keller [1968])

12.4.12. ($\diamond$) Since a semiorder P has no $\mathbf{2} + \mathbf{2}$, by Exercise 12.3.17 every subposet of P has upper and lower extensions. Let Q be the subposet of P consisting of the elements with even rank. Using upper and lower extensions of Q, prove $\dim P \leq 3$. (Rabinovitch [1978])

12.4.13. ($\diamond$) Although interval orders are more general than chains (or semiorders), prove that the standard example S_n cannot be an intersection of fewer than n interval orders.

12.4.14. Let P and Q be bounded posets with $|P|, |Q| \geq 2$. Prove that $P \times Q$ cannot be an intersection of fewer than $\dim P + \dim Q$ interval orders. (Trotter–Bogart [1976])

12.4.15. ($\diamond$) A general binary relation is modeled by a digraph D, with adjacency matrix $A(D)$. A **biorder representation** of D consists of real-valued functions f and g on $V(D)$ such that $uv \in E(D)$ if and only if $f(u) > g(v)$. For a digraph D, prove that the following five conditions are equivalent.

(A) $A(D)$ has no 2-by-2 submatrix that is a permutation matrix.
(B) The successor sets of D are ordered by inclusion.
(C) The predecessor sets of D are ordered by inclusion.
(D) The rows and columns of $A(D)$ can be permuted independently so that every entry below or to the left of a 1 is a 1.
(E) D has a biorder representation.
(Comment: The equivalences are due to various authors, with a short proof in West [1998]. Such relations are called **biorders**, **Ferrers relations**, or **Ferrers digraphs** (Riguet [1951], Wiener [1914]). If also $f(x) \le g(x)$ for all x, then D is an interval order.)

12.4.16. Which of the posets below are lattices? (Stanley [1986, p155])

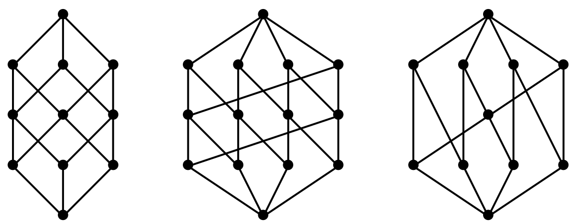

12.4.17. ($\diamond$) Prove that there are at least $n!/2$ copies of $\mathbf{2}^{n-1}$ in Π_n. (Hint: There are $n!/2$ paths with vertex set $[n]$.)

12.4.18. Let $J(P)$ denote the inclusion poset on the down-sets of P. Using the mapping suggested in Example 12.4.12, prove $L(m,n) \cong J(\mathbf{m} \times \mathbf{n}) \cong L(n,m)$,

12.4.19. As a component-wise order on m-tuples from $\{0,\dots,n\}$, the poset $L(m,n)$ is a subposet of $(\mathbf{n+1})^m$. Prove that the meet and join in $(\mathbf{n+1})^m$ of elements of $L(m,n)$ is also in $L(m,n)$, thereby proving by Definition 12.4.11 that $L(m,n)$ is a sublattice of $(\mathbf{n+1})^m$.

12.4.20. ($\diamond$) Write subsets of $[n]$ as strictly decreasing lists, $n \ge a_1 > \cdots > a_k > 0$. Let M_n be the poset of subsets of $[n]$, ordered by $a \le b$ if and only if $a_i \le b_i$ for all i, with trailing 0s added to permit comparison, so a subset is never less than a smaller subset.
(a) Describe the rank function and the cover relation in M_n.
(b) Describe M_{n+1} in terms of M_n.
(c) Prove that $M_{n+1} \cong J^2(\mathbf{2} \times \mathbf{n})$ (here J produces the lattice of down-sets).

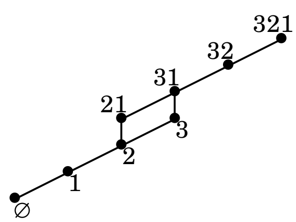

12.4.21. ($\diamond$) A *composition* of an integer n is an ordered list of positive integers summing to n. Let M_n' be the poset of compositions of $n+1$, ordered by $a \le b$ if $\sum_{i=1}^k a_i \ge \sum_{i=1}^k b_i$ for all k (trailing 0s are added as needed to test the order relation). Prove that $M_n' \cong M_n$, where M_n is the poset of Exercise 12.4.20. (The compositions of 4 are $\{4, 31, 22, 211, 13, 121, 112, 1111\}$.)

12.4.22. Let M_n be the poset of Exercise 12.4.20. Find an explicit symmetric chain decomposition for $n \le 5$. (Comment: Lindström conjectured that M_n is always a symmetric chain order. Stanley [1982] observed that M_n satisfies the strong Sperner property.)

12.4.23. (◇) Prove that if (x, y) is an unforced pair in a lattice L, then x is meet-irreducible and y is join-irreducible. Conclude that $\dim L = \dim R$, where R is the subposet of meet-irreducible and join-irreducible elements. (Kelly [1981])

12.4.24. (◇) *The Young lattice.* Let $\mathbf{Y}$ be the poset of all partitions of all integers, ordered by $a \le b$ if and only if $a_i \le b_i$ for all i (trailing zeros appended as needed).

(a) Prove that $\mathbf{Y}$ is a lattice. (Hint: Apply Lemma 12.4.17.)

(b) Prove that every $a \in \mathbf{Y}$ is covered by one more element than it covers.

(c) Use the fact that a product of infinitely many chains is a distributive lattice to prove that $\mathbf{Y}$ is distributive.

(d) Describe the join-irreducible elements of $\mathbf{Y}$.

12.4.25. (◇) Prove that the poset of partitions of n ordered by dominance is a lattice. Here $\mu \le \lambda$ if $\sum_{i=1}^{j} \mu_i \le \sum_{i=1}^{j} \lambda_i$ for all j. (Hint: Apply Lemma 12.4.17.)

12.4.26. (◇) Obtain a necessary and sufficient condition for λ to cover μ in the poset of partitions of n ordered by dominance.

12.4.27. For integers m and k with $0 \le k \le m-1$, let $C(m, k) = \{im + k \colon i \in \mathbb{Z}\}$; these sets are the **congruence classes**. Let P be the poset of all congruence classes, ordered by inclusion, with the empty set added as a minimal element. Prove that P is a lattice, and describe the meet and join of $C(m, k)$ and $C(n, l)$.

12.4.28. The poset $L_n(q)$ is the containment poset on the set of subspaces of an n-dimensional vector space over a q-element field. Prove that $L_n(q)$ is a graded lattice but in general is not a distributive lattice.

12.4.29. Let Λ_n denote the poset of partitions of the integer n, ordered by refinement. That is, for partitions λ and μ of n, we put $\lambda \le \mu$ if the multiset of parts in λ is obtained by replacing each integer part in μ with a partition of that part.

(a) Prove that Λ_n is a graded poset.

(b) Prove that Λ_7 is not a lattice. (Comment: This holds also for $n > 7$.)

12.4.30. (◇) Let A and B be maximum antichains in a poset P, and let A' and B' be the down-sets generated by A and B. Prove that the maximal elements of $A' \cup B'$ form a maximum antichain, and similarly for the maximal elements of $A' \cap B'$. Conclude that the maximum-sized antichains in a poset form a sublattice of the lattice of antichains, ordered by $A \le B$ if and only if for each $x \in A$ there exists $y \in B$ such that $x \le y$. (Dilworth [1960])

12.4.31. Given the result of Exercise 12.4.30, design a polynomial-time algorithm to construct the maximal element of the lattice of maximum-sized antichains. (Hint: Use the relation of Dilworth's Theorem to maximum matching.)

12.4.32. An **automorphism** of a poset permutes elements but preserves the order relation. Prove that the automorphism group of Π_n is isomorphic to the symmetric group $\mathbb{S}_n$.

12.4.33. Given a connected graph G, let $C(G)$ be the poset of the connected induced subgraphs of G, ordered by inclusion. Prove that $C(G)$ is a lattice if and only if every block of G is a complete graph. (Klavžar–Petkovšek [1988])

12.4.34. Given elements x and y in a distributive lattice L, let $F = \{z \in L \colon x \wedge z = y \wedge z\}$. Prove that F is a down-set.

12.4.35. For a lattice L, prove directly that $x \vee (y \wedge z) = (x \vee y) \wedge (x \vee z)$ for all $x, y, z \in L$ if and only if $x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$ for all $x, y, z \in L$.

12.4.36. Prove that in a finite distributive lattice the number of elements that cover exactly k elements equals the number of elements covered by exactly k elements.

12.4.37. Prove that a lattice L is distributive if and only if $(x \wedge y) \vee (y \wedge z) \vee (z \wedge x) = (x \vee y) \wedge (y \vee z) \wedge (z \vee x)$ for all $x, y, z \in L$. (Hint: There are proofs using the definition $x \wedge (y \vee z) = (x \wedge y) \vee (x \wedge z)$, the isomorphism to $J(P(L))$, or the forbidden sublattice characterization; look for a nice one.)

12.4.38. The *dual* R^* of a poset R is obtained by reversing all relations.

(a) Prove that $J(R^*) = J(R)^*$ for every poset R.

(b) Prove that P and Q are isomorphic if $J(P)$ and $J(Q)$ are isomorphic.

(c) Conclude that in a finite distributive lattice L the posets $P(L)$ and $Q(L)$ (join-irreducibles and meet-irreducibles) are isomorphic.

12.4.39. ($\diamond$) *Distributive lattices.* For an element x in a finite lattice L, let $\phi(x) = \{p \in P\colon p \le x\}$, where P is the poset of join-irreducible elements of L, and let $\psi(x) = \{q \in Q\colon q \ge x\}$, where Q is the poset of meet-irreducible elements of L.

(a) Prove that $\phi\colon L \to J(P)$ embeds L as a subposet of $J(P)$.

(b) Prove that the following are equivalent (via A⇒B⇒C⇒D⇒A.)

(A) L is distributive.
(B) L is graded and $|P| = r(L) = |Q|$.
(C) $|P| = |\phi(x)| + |\psi(x)| = |Q|$ for all $x \in L$.
(D) $|\phi(x \vee y)| = |\phi(x) \cup \phi(y)|$ for all $x, y \in L$.

12.4.40. In a lattice, y is a **complement** of x if $x \wedge y = \hat{0}$ and $x \vee y = \hat{1}$.

(a) Using the isomorphism $\phi\colon L \to J(P)$, characterize the elements of a distributive lattice that have complements.

(b) Show that the number of elements of a distributive lattice that have complements is a power of 2.

(c) In a product of chains, which elements have complements?

(d) A lattice is complemented if every element has a unique complement. Show that the only complemented distributive lattices are the Boolean algebras.

12.4.41. Prove that the order dimension (Section 12.3) of a distributive lattice equals the width of its poset of join-irreducible elements. (Comment: The famous "Dilworth's Theorem" was proved as a lemma for this result.) (Dilworth [1950])

12.4.42. Let L be a distributive lattice. Prove that the smallest dimension of a subset lattice in which L appears as a subposet is the number of join-irreducible elements in L.

12.4.43. The rank function of a graded lattice is **submodular** if $r(x \wedge y) + r(x \vee y) \le r(x) + r(y)$ for all elements x, y. Prove that the product of two lattices with submodular rank functions is a lattice with a submodular rank function. Conclude that the rank function of the divisor lattice $D(N)$ is **modular**, meaning that equality always holds in the submodularity inequality.

12.4.44. ($\diamond$) *Semimodularity of the partition lattice Π_n.*

(a) Prove by induction that $r(\pi \wedge \sigma) + r(\pi \vee \sigma) \le r(\pi) + r(\sigma)$ (i.e., that Π_n has a submodular rank function).

(b) Given two partitions π, σ of $[n]$, let $G(\pi, \sigma)$ be the bipartite graph whose parts are the blocks of π and σ, respectively, with vertices adjacent if they intersect as sets. In terms of the number of blocks in $\pi, \sigma, \pi \wedge \sigma, \pi \vee \sigma$, compute the number of vertices, edges, and components of $G(\pi, \sigma)$. Use this to give another proof of (a). (Aigner [1979])

12.4.45. *Applications of Kleitman's Inequality.* An **intersecting family** in $\mathbf{2}^n$ is a set $F \subseteq \mathbf{2}^n$ such that every two members of F have a common element.

(a) Prove that 2^{n-2} is the maximum size of an intersecting family whose complements also form an intersecting family. (Seymour [1973], Schönheim [1974], Daykin–Lovász [1976], Hilton [1976], Anderson [1976], Greene–Kleitman [1978])

(b) Prove that $2^n - 2^{n-k}$ is the maximum size of the union of k intersecting families in $\mathbf{2}^n$. (Kleitman [1966])

12.4.46. *Another application of Kleitman's Inequality.*

(a) Let X be an up-set and Y a down-set in $\mathbf{2}^n$. Let $\overline{X}$ and $\overline{Y}$ be their complements, and let $X - Y$ and $Y - X$ be their differences as sets. Use the up-set/down-set form of Kleitman's inequality to prove $|X - Y|\;|Y - X| \ge |X \cap Y|\;\left|\overline{X} \cap \overline{Y}\right|$, and conclude from this that $|X \cap Y|^{1/2} + \left|\overline{X} \cap \overline{Y}\right|^{1/2} \le 2^{n/2}$.

(b) Let F, G be subsets of $\mathbf{2}^n$ such that no member of either contains a member of the other. Apply (a) to prove that $|F|^{1/2} + |G|^{1/2} \le 2^{n/2}$. (Seymour [1973])

12.4.47. ($\diamond$) Using Theorem 12.4.22 for sufficiency and the Ahlswede–Daykin Inequality for necessity, prove that a lattice L is distributive if and only if $|F|\,|G| \le |F \vee G|\,|F \wedge G|$ for all $F, G \subseteq L$. Conclude Kleitman's Inequality for distributive lattices. (Daykin [1977])

12.4.48. For $X, Y \subseteq \mathbf{2}^n$, let $X - Y = \{A - B\colon A \in X,\ B \in Y\}$. Apply Exercise 12.4.47 to prove that $|X - X| \ge |X|$ for any family X of finite sets. (Comment: Daykin–Lovász [1976] showed in general that every nontrivial boolean function takes at least m distinct values when evaluated over m distinct sets.)

12.4.49. Complete the proofs of the FKG Inequality and Ahlswede–Daykin Inequality by proving the latter for $\mathbf{2}^1$. (Hint: For the case $X = Y = \{\varnothing, [1]\}$, let $w = \alpha(\varnothing)\beta([1])$, $x = \alpha([1])\beta(\varnothing)$, $y = \gamma(\varnothing)\delta([1])$, and $z = \gamma([1])\delta(\varnothing)$. Reduce the problem to inequalities involving these four quantities.)

12.4.50. Suppose that the outcome of matches between tennis players is determined by an unknown linear ordering of ability. Suppose that A and B are two teams with two players each, and that initially we know nothing about the relative abilities of the two players on a team. However, we do know $a_2 < b_1$. Determine whether the events $a_1 < b_1$ and $a_2 < b_2$ are positively correlated, assuming that unknown information is random. (Shepp [1980])

12.4.51. Let $Q = \mathbf{2} + \mathbf{4}$. Let z and y be the top and bottom of the 4-chain, and let w and x be the top and bottom of the 2-chain. Prove that the events $x < y < w$ and $x < z < w$ are not postively correlated on Q. (C. Mallows)

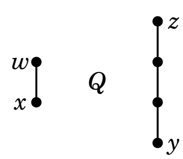

12.4.52. ($\diamond$) For $x \in Q$, let the random variable H_x be the height of x on a random linear extension of Q, and let $h(x) = \mathbb{E}(H_x)$. Let y be incomparable to x. Use the XYZ Inequality to prove $h(x|(x > y)) \ge 1 + h(x|(x < y))$. (Winkler [1982])

12.4.53. (+) The **track number** of a graph G, written $\tau(G)$, is the least t such that G is the union of t interval graphs. The **interval number** of G, written $i(G)$, is the least t such that G is the intersection from of unions of t real intervals. Always $i(G) \le t(G)$. Recall that $L(H)$ denotes the line graph of a graph H.

(a) Prove that $i(G) \le 2$ when G is a line graph.

(b) Prove that if $n \ge R_k(P')$, then $\tau(L(K_n)) > t$, where P' is the ordered hypergraph with vertices $1, 2, 3, 4, 5, 6$ in order obtained from the monotone path P_6^3 by adding two edges: $\{1, 2, 5\}$ and $\{2, 5, 6\}$.

(c) Prove that if $n < R_k(P_4^3)$, then $\tau(L(K_n)) \le t + 2$.

(d) For the complete 3-uniform hypergraph K_6^3, Erdős–Rado [1952] proved $R_k(K_6^3) \le$

$2^{2^{O(t \lg t)}}$ (Conlon–Fox–Sudakov [2010] improved this to $2^{2^{(4+o(1))t \lg t}}$). Use this and the lower bound $R_k(P_t^r) \ge 2^{m^{k-1}/2\sqrt{k}}$ (where $m = t - r + 1$) to prove

$$\Omega(\tfrac{\lg \lg n}{\lg \lg \lg n}) \le \tau(L(K_n)) \le \lg \lg n + \tfrac{1}{2} \lg \lg \lg n + O(\lg \lg \lg \lg n).$$

(Comment: Heldt, Knauer, and Ueckerdt conjectured that $\tau(L(K_n))$ is unbounded, proved here in (b). These results appear in Milans–Stolee–West [2015].)