2

Abstract Graphs and Set Systems

We introduce basic concepts and notation related to graphs, posets, abstract simplicial complexes, and matroids. In Section 2.1, we discuss graphs, digraphs, and hypergraphs. Section 2.2 is devoted to posets and lattices. We proceed with abstract simplicial complexes in Section 2.3 and conclude the chapter with some matroid theory in Section 2.4 and a few words about integer partitions in Section 2.5.

Basic Notation

In the below definitions, n and k are nonnegative integers, x is a real number, and S is a finite set.

$|x|$ is the absolute value of x; $|x| = x$ if $x \geq 0$ and $|x| = -x$ if $x < 0$. $\lfloor x \rfloor$ is the largest integer less than or equal to x, whereas $\lceil x \rceil$ is the smallest integer greater than or equal to x. For $n \geq 1$ and every integer a, $a \bmod n$ is the unique integer b in the set $\{0, \dots, n-1\}$ such that $(b-a)/n$ is a integer.

$\mathbb{Q}$ and $\mathbb{R}$ are the fields of rational and real numbers, respectively, whereas $\mathbb{Z}$ is the ring of integers. Define $\mathbb{Z}_n = \mathbb{Z}/n\mathbb{Z}$; this is the ring of integers modulo n. If n is a prime, then $\mathbb{Z}_n$ is a field.

We denote the empty set by $\emptyset$. 2^S is the family of all subsets of the set S, including S itself and $\emptyset$. $|S|$ is the cardinality (size) of the set S. Let $\binom{S}{k}$ be the family of all subsets T of S satisfying $|T| = k$; clearly, $|\binom{S}{k}| = \binom{|S|}{k}$. $\mathfrak{S}_S$ denotes the symmetric group on the set S, i.e., the group of permutations (bijections) $\pi : S \to S$. Multiplication is defined by $(\pi\pi')(x) = \pi(\pi'(x))$. Finally, we define $[k, n] = \{m \in \mathbb{Z} : k \leq m \leq n\}$ and $[n] = [1, n] = \{1, \dots, n\}$.

2.1 Graphs, Hypergraphs, and Digraphs

We present standard graph-theoretic concepts.

2.1.1 Graphs

A (simple) *graph* $G = (V, E)$ consists of a finite set V of *vertices* and a family E of subsets of V of size two called *edges*; $E \subseteq \binom{V}{2}$. An edge should be thought of as a line connecting the two vertices in it. A graph being *simple* means that there is at most one edge between any two vertices; E is not a multiset. The edge between the two vertices a and b is denoted as ab or $\{a, b\}$. Two vertices a and b are *adjacent* in G if $ab \in E$.

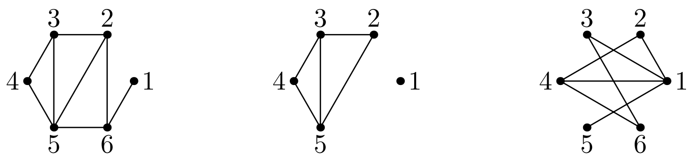

Fig. 2.1. The graph $G = ([6], \{16, 23, 25, 26, 34, 35, 45, 56\})$ to the left, the induced subgraph $G([5])$ in the middle, and the complement of G to the right. We have that $N_G(6) = \{1, 2, 5\}$ and $\deg_G(6) = 3$. The vertex set $\{1, 2, 4\}$ is a stable set in G, whereas $\{2, 3, 5\}$ is a clique. The edge set $\{16, 25, 34\}$ forms a perfect matching contained in G. We obtain a proper 3-coloring $\gamma : [6] \to [3]$ of G by defining $\gamma^{-1}(1) = \{1, 2, 4\}$, $\gamma^{-1}(2) = \{3, 6\}$, and $\gamma^{-1}(3) = \{5\}$.

For $v \in V$, the *neighborhood* of v is the set $N_G(v) = \{w \in V \setminus \{v\} : vw \in E\}$. The *degree* of v is $\deg_G(v) = |N_G(v)|$. For $W \subseteq V$, define the *induced subgraph* $G(W)$ of G on the vertex set W as the pair $(W, E \cap \binom{W}{2})$.

A *matching* on a vertex set V is a graph $G = (V, E)$ such that each vertex $v \in V$ is adjacent to at most one other vertex in G. A matching is *perfect* if each vertex is adjacent to exactly one other vertex.

A vertex set U in G is *stable* if no edge in G is a subset of U; no two vertices in U are adjacent. Some authors refer to stable sets as *independent*. A vertex set W is a *clique* in G if $\binom{W}{2} \subseteq E$; every two vertices in W are adjacent. The *complement* of a graph $G = (V, E)$ is the graph $\bar{G} = (V, \binom{V}{2} \setminus E)$. Note that U is a clique in G if and only if U is an stable set in $\bar{G}$.

A *t-coloring* of a graph $G = (V, e)$ is a function $\gamma : V \to [t]$. A coloring γ is *proper* if $\gamma(v) \neq \gamma(w)$ whenever $vw \in E$. A graph $G = (V, E)$ is *t-colorable* if there is a proper t-coloring of G.

For $n \geq 1$, K_n denotes the *complete graph* on n vertices containing all $\binom{n}{2}$ possible edges. 2^{K_n} is the family of all graphs on n vertices.

Some of the concepts introduced in this section are illustrated in Figure 2.1.

2.1.2 Paths, Components and Cycles

A *path* in a graph $G = (V, E)$ is a sequence $(\rho_1, \ldots, \rho_r)$ of not necessarily distinct vertices from V such that $\rho_i \rho_{i+1} \in E$ for $1 \leq i \leq r - 1$. If $\rho_1, \ldots, \rho_r$

are all distinct, then the path is *simple*. We obtain an equivalence relation on V by letting v and w be equivalent if and only if there is a (simple) path $(\rho_1, \ldots, \rho_r)$ in G with $\rho_1 = v$ and $\rho_r = w$. The equivalence classes under this relation are the *connected components* of G. We will typically identify the connected components $W_1, \ldots, W_k$ with the corresponding induced subgraphs $G(W_1), \ldots, G(W_k)$. A graph G is *disconnected* if G contains at least two connected components; otherwise, G is *connected*. A vertex v is *isolated* in G if the connected component containing v equals $\{v\}$.

A vertex set W in a graph $G = (V, E)$ is a *cut set* if $G(V \setminus W)$ is disconnected. If $W = \{w\}$, then w is a *cut point*. For $1 \leq k \leq |V|$, we say that G is k-*connected* if G does not contain any cut set of size less than k. For example, G being 1-connected means that G is connected.

A path $(\rho_1, \ldots, \rho_r)$ in a graph G is a *cycle* if $\rho_r \rho_1 \in G$. The cycle is *simple* if it is simple as a path. G contains a cycle if and only if G contains a simple cycle. A *forest* is a cycle-free graph. A *tree* is a forest such that all non-isolated vertices belong to the same connected component. A *spanning tree* is a tree with one single connected component.

A simple path containing all vertices in a graph is a *Hamiltonian path*; a simple cycle containing all vertices is a *Hamiltonian cycle*. A graph is *Hamiltonian* if it contains at least one Hamiltonian cycle and *non-Hamiltonian* otherwise.

2.1.3 Bipartite Graphs

A graph G is *bipartite* if G is 2-colorable. Equivalently, the vertex set of G is the disjoint union of two stable vertex sets U and W; we say that (U, W) is a *bipartition* of G and refer to U and W as the *blocks* of G. Note that the blocks are not uniquely determined unless G is connected. For $m, n \geq 1$, $K_{m,n}$ denotes the *complete bipartite graph* on a vertex set $U \cup W$ such that $U \cap W = \emptyset$, $|U| = m$, and $|W| = n$; this graph contains all mn possible edges uw such that $u \in U$ and $w \in W$.

2.1.4 Digraphs

A (simple and loopless) *digraph* $D = (V, A)$ consists of a finite set V of *vertices* and a set A of ordered pairs $vw = (v, w)$ such that $v \neq w$; $A \subseteq V \times V \setminus \{(v, v) : v \in V\}$. the elements in A are called *directed edges*. The edge vw is *directed* from v to w; v is the *tail* and w is the *head*. For $n \geq 1$, $K_n^{\rightarrow}$ denotes the *complete digraph* on n vertices containing all $n(n-1)$ possible edges.

2.1.5 Directed Paths and Cycles

A *directed path* in a digraph D is a sequence $(\rho_1, \ldots, \rho_r)$ of not necessarily distinct vertices in V such that $\rho_i \rho_{i+1} \in A$ for $1 \leq i \leq r - 1$. A directed path

$(\rho_1, \dots, \rho_r)$ is a *directed cycle* if $\rho_r \rho_1 \in A$. In a *simple* directed path or cycle, we require all vertices to be distinct. A *directed Hamiltonian path* is a simple directed path containing all vertices; *directed Hamiltonian cycles* are defined analogously. A digraph D is *acyclic* if D does not contain any directed cycles. A digraph is *Hamiltonian* if it contains at least one directed Hamiltonian cycle and *non-Hamiltonian* otherwise. A digraph D is *strongly connected* if every pair of vertices in D are contained in a directed cycle; the cycle need not be simple.

D is a *directed forest* if D is acyclic and each vertex is the head of at most one edge.[1] A *directed tree* is a directed forest such that all non-isolated vertices belong to the same connected component. A *spanning directed tree* is a directed tree with one single connected component. In such a tree, there is a unique element – the *root* – that is not the head of any edge.

2.1.6 Hypergraphs

A (simple) *hypergraph* $H = (V, E)$ consists of a finite set V of vertices and a family E of nonempty subsets of V called *edges*. We denote the edge $\{a_1, a_2, \dots, a_r\}$ as $a_1 a_2 \dots a_r$. For a set S of positive integers, H is an *S-hypergraph* if $|e| \in S$ for every $e \in E$. If H is an *$\{r\}$-hypergraph* (i.e., all edges have the same size r), then H is *r-uniform*. For example, ordinary graphs are 2-uniform. For $W \subseteq V$, define the *induced subhypergraph* $G(W)$ of G with respect to the vertex set W as the pair $(W, E \cap 2^W)$; only edges contained in W remain.

2.1.7 General Terminology

Let $G = (V, E)$ be a graph, hypergraph, or digraph. G is *empty* if $E = \emptyset$ and *nonempty* otherwise. A vertex is *covered* in G if the vertex is contained in some edge in G and *uncovered* otherwise. For hypergraphs, the terms "uncovered" and "isolated" (see Section 2.1.2) are not equivalent. Specifically, if the only edge in G containing a given vertex v is the singleton edge $\{v\}$, then v is isolated but not uncovered. Whenever the underlying vertex set V is fixed, we identify G with its set of edges; $e \in G$ means that $e \in E$. For an edge e, we will write $G - e = (V, E \setminus \{e\})$ and $G + e = (V, E \cup \{e\})$. We let $|G|$ denote the size of the edge set of G. Whenever we refer to "the family of all graphs on n vertices with a given property P", we mean to first fix a vertex set V of size n and then consider the family of all graphs G on the vertex set V with property P.

[1] Some authors prefer to define directed forests in terms of the dual requirement that each vertex is the *tail* of at most one edge.

2.2 Posets and Lattices

A finite *partially ordered set* or *poset* is a pair $P = (X, \leq)$, where X is a finite set and $\leq$ is a binary relation on X satisfying the following conditions for all $x, y, z \in X$:

- $x \leq x$.
- If $x \leq y$ and $y \leq x$, then $x = y$.
- If $x \leq y$ and $y \leq z$, then $x \leq z$.

An element x is an *atom* in P if $y \not\leq x$ whenever $y \neq x$. Two elements x and y form a *covering relation* in P if $x < y$ (i.e., $x \leq y$ and $x \neq y$) and no element z in X satisfies $x < z < y$. The *direct product* of two posets $P = (X, \leq_P)$ and $Q = (Y, \leq_Q)$ is the poset $P \times Q = (X \times Y, \leq_{P\times Q})$, where $(x, y) \leq_{P\times Q} (x', y')$ if and only if $x \leq_P x'$ and $y \leq_Q y'$. An (order-preserving) *poset map* between two posets $P = (X, \leq_P)$ and $Q = (Y, \leq_Q)$ is a function $f : X \to Y$ such that $f(x) \leq_Q f(y)$ whenever $x \leq_P y$. We will often write $f : P \to Q$.

A *chain* is a set $\{x_1, \ldots, x_r\}$ of elements in X such that $x_1 < x_2 < \cdots < x_r$. A poset is *ranked* of *rank* d if every maximal chain has size d. The *rank* of an element x is the size of a largest chain in which x is the maximal element. It is often useful to introduce a minimal element $\hat{0}$ with rank 0 and a maximal element $\hat{1}$ of rank $d + 1$. $\hat{0}$ is smaller and $\hat{1}$ is larger than all elements in X.

A finite *lattice* is a finite poset $L = (X, \leq_L)$ such that the following hold:

- There are elements $\hat{0}, \hat{1} \in X$ such that $\hat{0} \leq_L x$ and $x \leq_L \hat{1}$ for all $x \in X$.
- Any two elements $x, y \in X$ have a unique greatest lower bound. Thus there exists an element $z \leq_L x, y$ such that $w \leq_L z$ whenever $w \leq_L x, y$.

These conditions imply that any two elements have a unique least upper bound. The *proper part* of a lattice L, denoted $\overline{L}$, is the poset obtained by removing the top element $\hat{1}$ and the bottom element $\hat{0}$ from L.

A *partition* of a finite set V is a family $\{U_1, \ldots, U_k\}$ of nonempty sets such that V is the disjoint union of $U_1, \ldots, U_k$. The *partition lattice* Π_V is the poset of partitions of V ordered under refinement; $\{W_1, \ldots, W_m\}$ is a refinement of – and hence smaller than – $\{U_1, \ldots, U_k\}$ if every W_i is a subset of some U_j. The partition lattice is indeed a lattice [133]. We write $\Pi_n = \Pi_{[n]}$.

Unless otherwise specified, whenever a family Δ of subsets of a set X is referred to as a poset, the underlying order $\leq$ is given by set inclusion;

$$A \leq B \Longleftrightarrow A \subseteq B.$$

2.3 Abstract Simplicial Complexes

We introduce set-theoretic concepts and notation related to abstract simplicial complexes. Throughout the section, all sets and families are finite. Whenever appropriate, we extend our definitions to arbitrary families of sets rather than restricting to the special case of simplicial complexes.

2.3.1 Basic Definitions

An (abstract) *simplicial complex* Δ on a finite set X is a family of subsets of X closed under deletion of elements. We refer to the singleton sets $\{x\}$ in Δ as 0-*cells* or *vertices*. We do *not* require that $\{x\} \in \Delta$ for all $x \in X$. For the purposes of this book, we adopt the convention that the *void complex* $\emptyset$ is a simplicial complex. For geometric reasons, many authors refer to the complex $\{\emptyset\}$, which is different from the void complex, as the *empty complex*. To avoid any confusion, we will consistently refer to any empty family $\emptyset$ as "void" rather than "empty". Members of a simplicial complex Δ are called *faces*. For a face σ and an element $x \in X$, we write $\sigma - x = \sigma \setminus \{x\}$ and $\sigma + x = \sigma \cup \{x\}$. For two simplicial complexes Δ_1 and Δ_2, $\Delta_1 \cong \Delta_2$ means that Δ_1 and Δ_2 are *combinatorially equivalent*. Assuming that X and Y are the vertex sets of Δ_1 and Δ_2, respectively, this means that there exists a bijection $\varphi : X \to Y$ such that $\sigma \in \Delta_1$ if and only if $\varphi(\sigma) \in \Delta_2$ for each set $\sigma \subseteq X$. Note that the same symbol $\cong$ also denotes homeomorphism between topological spaces. Whenever we use the symbol, it will be clear from context how to interpret it. The simplicial complex *generated* by a family $\mathcal{M}$ of sets is the complex of all subsets of sets in $\mathcal{M}$, including $\mathcal{M}$ itself.

2.3.2 Dimension

Define the *dimension* of a set σ as $|\sigma| - 1$. One sometimes refers to a set of dimension d as a d-*face* or d-*cell*. The dimension of a nonvoid family Δ is the maximum dimension among faces of Δ. The *(reduced) Euler characteristic* of Δ is defined as the integer

$$\tilde{\chi}(\Delta) = \sum_{\sigma \in \Delta} (-1)^{\dim \sigma}.$$

For $d \geq -1$, the d-*skeleton* of a family is the family of all sets of dimension at most d. A family is *pure* if all maximal faces (with respect to inclusion) have the same dimension. For a set σ, we refer to the family 2^σ as the *full simplex* on σ. Writing $d = \dim \sigma = |\sigma| - 1$, we say that 2^σ is a d-*simplex*. Note that the (-1)-simplex contains the empty set and nothing else. We sometimes refer to the 0-simplex as a *point*. We obtain the *boundary* $\partial 2^\sigma$ of the d-simplex 2^σ by removing the maximal face σ.

2.3.3 Collapses

A simplicial complex Δ is obtained from another simplicial complex Δ' via an *elementary collapse* if $\Delta' \setminus \Delta = \{\sigma, \tau\}$ and $\sigma \subsetneq \tau$. This means that τ is the only face in Δ' properly containing σ. If Δ can be obtained from Δ' via a sequence of elementary collapses, then Δ' can be *collapsed* to Δ. If Δ' is void or can be collapsed to a 0-simplex $\{\emptyset, \{v\}\}$, then Δ' is *collapsible (to a point)*.

2.3.4 Joins, Cones, Suspensions, and Wedges

The *join* of two families Δ and Γ (assumed to be defined on disjoint ground sets) is the family $\Delta * \Gamma = \{\sigma \cup \tau : \sigma \in \Delta, \tau \in \Gamma\}$. Note that $\Delta * \emptyset = \emptyset$ and $\Delta * \{\emptyset\} = \Delta$. Let x be a 0-cell not in Δ. The *cone* $\mathsf{Cone}(\Delta) = \mathsf{Cone}_x(\Delta)$ over Δ with *cone point* x is the join of Δ with the 0-simplex $\{\emptyset, \{x\}\}$. Cones over simplicial complexes are collapsible. Let y be another 0-cell not in Δ. The *suspension* $\mathrm{Susp}(\Delta) = \mathrm{Susp}_{x,y}(\Delta)$ of Δ with respect to the pair $\{x, y\}$ is the join of Δ with $\{\emptyset, \{x\}, \{y\}\}$. Note that $\mathrm{Susp}_{x,y}(\Delta) = \mathsf{Cone}_x(\Delta) \cup \mathsf{Cone}_y(\Delta)$. We obtain the (one-point) *wedge* $\Delta \vee \Gamma$ of two simplicial complexes Δ and Γ with respect to 0-cells $x \in \Delta$, $y \in \Gamma$ by taking the disjoint union of Δ and Γ and then identifying x and y.

2.3.5 Alexander Duals

For a simplicial complex Δ on a set X, the *Alexander dual* of Δ with respect to X is the simplicial complex $\Delta_X^* = \{\sigma \subseteq X : X \setminus \sigma \notin \Delta\}$. If there is no reference to any underlying set X, it is assumed that X is the set of 0-cells in Δ.

2.3.6 Links and Deletions

For a family Δ of sets and a set σ, the *link* $\mathrm{lk}_\Delta(\sigma)$ is the family of all $\tau \in \Delta$ such that $\tau \cap \sigma = \emptyset$ and $\tau \cup \sigma \in \Delta$. The *deletion* $\mathrm{del}_\Delta(\sigma)$ is the family of all $\tau \in \Delta$ such that $\tau \cap \sigma = \emptyset$. We define the *face-deletion* $\mathrm{fdel}_\Delta(\sigma)$ as the family of all $\tau \in \Delta$ such that $\sigma \not\subseteq \tau$. The link, deletion, and face-deletion of a simplicial complex are all simplicial complexes.

2.3.7 Lifted Complexes

For the purposes of this book, a family Σ of sets is a *lifted complex* over a set σ if Σ is of the form $\Delta * \{\sigma\}$, where Δ is a simplicial complex and σ is a finite set disjoint from all sets in Δ. Any simplicial complex is also a lifted complex; σ may be the empty set.

Given a lifted complex Σ and disjoint sets I and E, define

$$\Sigma(I, E) = \{I\} * \mathrm{lk}_{\mathrm{del}_\Sigma(E)}(I) = \{\tau \in \Sigma : I \subseteq \tau, E \cap \tau = \emptyset\}.$$

If Σ is a lifted complex over σ, then $\Sigma(I, E)$ is a lifted complex over $\sigma \cup I$. Note that $\Sigma(\emptyset, E) = \mathrm{del}_\Sigma(E)$.

2.3.8 Order Complexes and Face Posets

The *order complex* $\Delta(P)$ of a poset $P = (X, \le)$ is the simplicial complex of all chains in P; a set $A \subseteq X$ belongs to $\Delta(P)$ if and only if $a \le b$ or $b \le a$ for all

$a, b \in A$. Whenever we say that a poset P has a certain topological property (e.g., a certain homotopy type), we mean that $\Delta(P)$ has the property. The *face poset* $P(\Delta)$ of a simplicial complex Δ is the poset of *nonempty* faces of Δ ordered by inclusion. $\mathrm{sd}(\Delta) = \Delta(P(\Delta))$ is the *(first) barycentric subdivision* of Δ.

2.3.9 Graph, Digraph, and Hypergraph Complexes and Properties

A *graph complex* on a finite vertex set V is a family Σ of simple graphs on the vertex set V such that Σ is closed under deletion of edges; if $G \in \Sigma$ and $e \in G$, then $G - e \in \Sigma$. Identifying $G = (V, E) \in \Sigma$ with the edge set E, we may interpret Σ as a simplicial complex. Analogously, a *digraph complex* on V is a family of simple and loopless digraphs on V closed under deletion of edges, whereas a *hypergraph complex* on V is a family of simple hypergraphs on V, again closed under deletion of edges. The restriction to *simple* graphs, digraphs, and hypergraphs is for the purposes of this book.

For a graph complex Σ on a vertex set V and a graph $G = (V, E)$, define $\Sigma(G)$ as the graph complex consisting of all graphs H in Σ such that H is a subgraph of G. We refer to $\Sigma(G)$ as the *induced (graph) subcomplex* of Σ. We adopt the same terminology for digraph and hypergraph complexes.

We refer to a digraph complex $\hat{\Delta}$ as the *trivial extension* of a graph complex Δ if the following holds:

- A digraph D is a maximal face of $\hat{\Delta}$ if and only if D equals $\{ab, ba : ab \in G\}$ for some maximal face G of Δ.

For example, the property of being a disconnected digraph is the trivial extension of the property of being a disconnected undirected graph.

A *graph property* is a family Σ of simple graphs on a finite vertex set V such that Σ is closed under permutations of the vertex set V; if $\sigma := \{a_1b_1, \ldots, a_rb_r\} \in \Sigma$ and $\pi \in \mathfrak{S}_V$, then

$$\pi(\sigma) := \{\pi(a_1)\pi(b_1), \ldots, \pi(a_r)\pi(b_r)\} \in \Sigma.$$

We refer to this action as the *natural action* of $\mathfrak{S}_V$ on Δ.

A *digraph property* is a family Σ of simple and loopless digraphs on a finite vertex set V such that Σ is closed under permutations of the vertex set V. Analogously, a *hypergraph property* is a family of hypergraphs, again on a fixed vertex set, that is closed under permutations of the underlying vertex set.

A graph, digraph, or hypergraph property Σ is *monotone* if Σ is closed under deletion of edges. Equivalently, Σ is a simplicial complex.

2.4 Matroids

A finite *matroid* M is a pair (E, F), where E is a finite set and $\mathsf{F} = \mathsf{F}(M) \subseteq 2^E$ is a nonvoid simplicial complex satisfying the following property:

- If $\sigma, \tau \in \mathsf{F}$ and $|\sigma| < |\tau|$, then there is an element $x \in \tau \setminus \sigma$ such that $\sigma + x \in \mathsf{F}$.

$\mathsf{F}(M)$ is the *independence complex* or *matroid complex* of M. The sets in $\mathsf{F}(M)$ are the *independent sets* in M. Note that F is a pure complex; all maximal faces have the same size. Define the *rank* of M as this size. A *basis* is a maximal independent set. A *circuit* is a minimal dependent set, i.e., a minimal nonface of $\mathsf{F}(M)$.

For a subset τ of E, let $M(\tau)$ denote the pair $(\tau, \mathsf{F} \cap 2^\tau)$. This is a matroid, and we refer to it as the *induced submatroid* of M on the set τ. Define the rank $\rho_M(\tau)$ of τ as the rank of the matroid $M(\tau)$. A set τ is a *flat* in M if the rank of $\tau + x$ exceeds the rank of τ for each x in $E \setminus \tau$. If a flat τ has rank $\rho(E) - 1$, then τ is a *cocircuit* in M.

For $e \in E$, $M - e$ is the pair $(E - e, \mathrm{del}_{\mathsf{F}}(e))$; $M - e$ is the *deletion* of M with respect to e. M/e is the pair $(E - e, \mathrm{lk}_{\mathsf{F}}(e))$; M/e is the *contraction* of M with respect to e. The rank function of M/e satisfies the identity

$$\rho_{M/e}(\sigma) = \rho_M(\sigma + e) - \rho_M(\{e\}).$$

The *dual* of M is the matroid M^* on the same ground set E with the property that the rank function ρ^* satisfies

$$\rho^*(\sigma) = |\sigma| + \rho(E \setminus \sigma) - \rho(E). \tag{2.1}$$

Equivalently, σ is a basis of M^* if and only if $E \setminus \sigma$ is a basis of M.

We refer the reader to Oxley [105] or Welsh [147] for more information about matroids.

2.4.1 Graphic Matroids

For a graph $G = ([n], E)$, define $M_n(G)$ to be the pair $(E, \mathsf{F}_n(G))$, where $\mathsf{F}_n(G)$ is the complex of forests contained in G. This is well-known to be a matroid, and the rank function is given by $\rho(H) = n - c(H)$, where $c(H)$ is the number of connected components in H. We refer to $M_n(G)$ as the *graphic matroid* on G. Write $M_n = M_n(K_n)$.

Another matroid that we may associate to G is the *(one-step) truncation* of $M_n(G)$ obtained by redefining the rank function as $\rho(H) = \min\{\rho(H), n - 2\} = n - \max\{2, c(H)\}$. The independent sets in this matroid are exactly all disconnected forests in G. One may pursue this construction further, considering the "k-step" truncation with rank function $\rho(H) = n - \max\{k, c(H)\}$, but we will confine ourselves to the one-step construction.

For a digraph D, let $M_n(D)$ be the matroid with the property that a set of edges is independent if and only if there are no multiple edges or cycles in the underlying undirected graph. The former condition means exactly that $\{ij, ji\}$ is *not* independent. We refer to $M_n(D)$ as the *digraphic* matroid on D. Write $M_n^{\rightarrow} = M_n(K_n^{\rightarrow})$.

2.5 Integer Partitions

For a sequence $\lambda = (\lambda_1, \ldots, \lambda_r)$, define $|\lambda| = \sum_{i=1}^{r} \lambda_i$. Say that λ is a *partition* of n if $\lambda_1 \geq \cdots \geq \lambda_r \geq 1$ and $|\lambda| = n$; we write this as $\lambda \vdash n$. By convention, we set λ_i equal to 0 whenever $i > r$. One may interpret λ as the set $\{(i,j) : 1 \leq j \leq \lambda_i\}$ of lattice points, where (i,j) is the lattice point in the i^{th} row and j^{th} column. Write $D_\lambda = \{(i,i) : \lambda_i \geq i\}$; this is the *diagonal* of λ. Points (i,j) such that $i < j$ are *above* the diagonal, whereas points (i,j) such that $i > j$ are *below* the diagonal.

Given two partitions λ and μ of n, we say that λ *dominates* μ if

$$\sum_{i=1}^{k} \lambda_i \geq \sum_{i=1}^{k} \mu_i$$

for all $i \geq 1$. The *conjugate* λ^T of a partition $\lambda = (\lambda_1, \ldots, \lambda_r)$ is the sequence $(\mu_1, \ldots, \mu_{\lambda_1})$ with the property that μ_j is the largest m such that $\lambda_m \geq j$. Equivalently, the length of the j^{th} row in λ^T equals the length of the j^{th} column in λ for each j. λ is *self-conjugate* if $\lambda = \lambda^T$.

Springer
http://www.springer.com/978-3-540-75858-7

Simplicial Complexes of Graphs
Jonsson, J.
2008, XIV, 382 p. 34 illus., Softcover
ISBN: 978-3-540-75858-7