1

Graphs

Contents

1.1 Graphs and Their Representation

Definitions and Examples

Many real-world situations can conveniently be described by means of a diagram consisting of a set of points together with lines joining certain pairs of these points.

For example, the points could represent people, with lines joining pairs of friends; or the points might be communication centres, with lines representing communication links. Notice that in such diagrams one is mainly interested in whether two given points are joined by a line; the manner in which they are joined is immaterial. A mathematical abstraction of situations of this type gives rise to the concept of a graph.

A *graph* G is an ordered pair $(V(G), E(G))$ consisting of a set $V(G)$ of *vertices* and a set $E(G)$, disjoint from $V(G)$, of *edges*, together with an *incidence function* ψ_G that associates with each edge of G an unordered pair of (not necessarily distinct) vertices of G. If e is an edge and u and v are vertices such that $\psi_G(e) = \{u, v\}$, then e is said to *join* u and v, and the vertices u and v are called the *ends* of e. We denote the numbers of vertices and edges in G by $v(G)$ and $e(G)$; these two basic parameters are called the *order* and *size* of G, respectively.

Two examples of graphs should serve to clarify the definition. For notational simplicity, we write uv for the unordered pair $\{u, v\}$.

Example 1.

$$G = (V(G), E(G))$$

where

$$\begin{aligned} V(G) &= \{u, v, w, x, y\} \\ E(G) &= \{a, b, c, d, e, f, g, h\} \end{aligned}$$

and ψ_G is defined by

$$\begin{array}{llll} \psi_G(a) = uv & \psi_G(b) = uu & \psi_G(c) = vw & \psi_G(d) = wx \\ \psi_G(e) = vx & \psi_G(f) = wx & \psi_G(g) = ux & \psi_G(h) = xy \end{array}$$

Example 2.

$$H = (V(H), E(H))$$

where

$$\begin{aligned} V(H) &= \{v_0, v_1, v_2, v_3, v_4, v_5\} \\ E(H) &= \{e_1, e_2, e_3, e_4, e_5, e_6, e_7, e_8, e_9, e_{10}\} \end{aligned}$$

and ψ_H is defined by

$$\begin{array}{lllll} \psi_H(e_1) = v_1v_2 & \psi_H(e_2) = v_2v_3 & \psi_H(e_3) = v_3v_4 & \psi_H(e_4) = v_4v_5 & \psi_H(e_5) = v_5v_1 \\ \psi_H(e_6) = v_0v_1 & \psi_H(e_7) = v_0v_2 & \psi_H(e_8) = v_0v_3 & \psi_H(e_9) = v_0v_4 & \psi_H(e_{10}) = v_0v_5 \end{array}$$

Drawings of Graphs

Graphs are so named because they can be represented graphically, and it is this graphical representation which helps us understand many of their properties. Each vertex is indicated by a point, and each edge by a line joining the points representing its ends. Diagrams of G and H are shown in Figure 1.1. (For clarity, vertices are represented by small circles.)

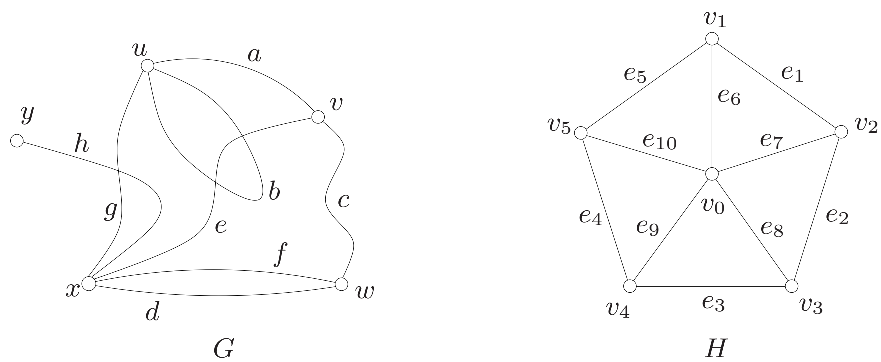

Fig. 1.1. Diagrams of the graphs G and H

There is no single correct way to draw a graph; the relative positions of points representing vertices and the shapes of lines representing edges usually have no significance. In Figure 1.1, the edges of G are depicted by curves, and those of H by straight-line segments. A diagram of a graph merely depicts the incidence relation holding between its vertices and edges. However, we often draw a diagram of a graph and refer to it as the graph itself; in the same spirit, we call its points 'vertices' and its lines 'edges'.

Most of the definitions and concepts in graph theory are suggested by this graphical representation. The ends of an edge are said to be *incident* with the edge, and *vice versa*. Two vertices which are incident with a common edge are *adjacent*, as are two edges which are incident with a common vertex, and two distinct adjacent vertices are *neighbours*. The set of neighbours of a vertex v in a graph G is denoted by $N_G(v)$.

An edge with identical ends is called a *loop*, and an edge with distinct ends a *link*. Two or more links with the same pair of ends are said to be *parallel edges*. In the graph G of Figure 1.1, the edge b is a loop, and all other edges are links; the edges d and f are parallel edges.

Throughout the book, the letter G denotes a graph. Moreover, when there is no scope for ambiguity, we omit the letter G from graph-theoretic symbols and write, for example, V and E instead of $V(G)$ and $E(G)$. In such instances, we denote the numbers of vertices and edges of G by n and m, respectively.

A graph is *finite* if both its vertex set and edge set are finite. In this book, we mainly study finite graphs, and the term 'graph' always means 'finite graph'. The graph with no vertices (and hence no edges) is the *null graph*. Any graph with just one vertex is referred to as *trivial*. All other graphs are *nontrivial*. We admit the null graph solely for mathematical convenience. Thus, unless otherwise specified, all graphs under discussion should be taken to be nonnull.

A graph is *simple* if it has no loops or parallel edges. The graph H in Example 2 is simple, whereas the graph G in Example 1 is not. Much of graph theory is concerned with the study of simple graphs.

A set V, together with a set E of two-element subsets of V, defines a simple graph (V, E), where the ends of an edge uv are precisely the vertices u and v. Indeed, in any simple graph we may dispense with the incidence function ψ by renaming each edge as the unordered pair of its ends. In a diagram of such a graph, the labels of the edges may then be omitted.

Special Families of Graphs

Certain types of graphs play prominent roles in graph theory. A *complete graph* is a simple graph in which any two vertices are adjacent, an *empty graph* one in which no two vertices are adjacent (that is, one whose edge set is empty). A graph is *bipartite* if its vertex set can be partitioned into two subsets X and Y so that every edge has one end in X and one end in Y; such a partition (X, Y) is called a *bipartition* of the graph, and X and Y its *parts*. We denote a bipartite graph G with bipartition (X, Y) by $G[X, Y]$. If $G[X, Y]$ is simple and every vertex in X is joined to every vertex in Y, then G is called a *complete bipartite graph*. A *star* is a complete bipartite graph $G[X, Y]$ with $|X| = 1$ or $|Y| = 1$. Figure 1.2 shows diagrams of a complete graph, a complete bipartite graph, and a star.

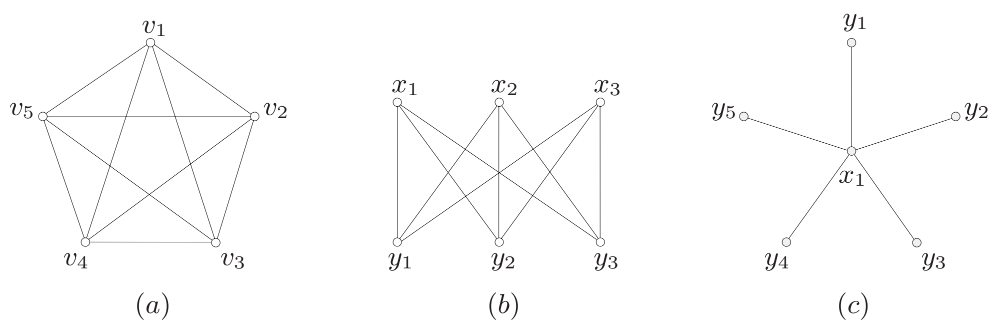

Fig. 1.2. (a) A complete graph, (b) a complete bipartite graph, and (c) a star

A *path* is a simple graph whose vertices can be arranged in a linear sequence in such a way that two vertices are adjacent if they are consecutive in the sequence, and are nonadjacent otherwise. Likewise, a *cycle* on three or more vertices is a simple graph whose vertices can be arranged in a cyclic sequence in such a way that two vertices are adjacent if they are consecutive in the sequence, and are nonadjacent otherwise; a cycle on one vertex consists of a single vertex with a loop, and a cycle on two vertices consists of two vertices joined by a pair of parallel edges. The *length* of a path or a cycle is the number of its edges. A path or cycle of length k is called a k*-path* or k*-cycle*, respectively; the path or cycle is *odd* or *even* according to the parity of k. A 3-cycle is often called a *triangle*, a 4-cycle a *quadrilateral*, a 5-cycle a *pentagon*, a 6-cycle a *hexagon*, and so on. Figure 1.3 depicts a 3-path and a 5-cycle.

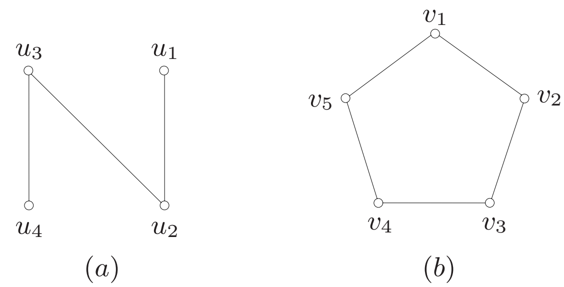

Fig. 1.3. (a) A path of length three, and (b) a cycle of length five

A graph is *connected* if, for every partition of its vertex set into two nonempty sets X and Y, there is an edge with one end in X and one end in Y; otherwise the graph is *disconnected*. In other words, a graph is disconnected if its vertex set can be partitioned into two nonempty subsets X and Y so that no edge has one end in X and one end in Y. (It is instructive to compare this definition with that of a bipartite graph.) Examples of connected and disconnected graphs are displayed in Figure 1.4.

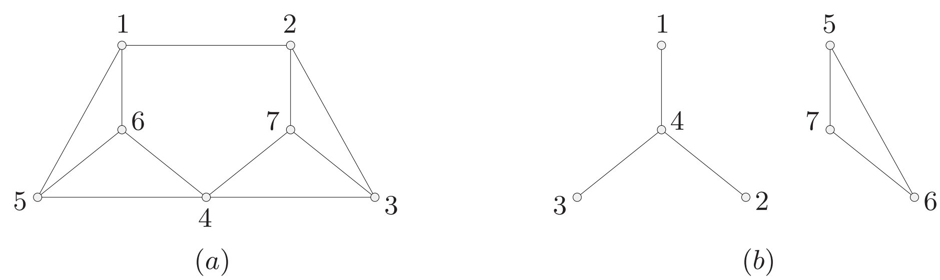

Fig. 1.4. (a) A connected graph, and (b) a disconnected graph

As observed earlier, examples of graphs abound in the real world. Graphs also arise naturally in the study of other mathematical structures such as polyhedra, lattices, and groups. These graphs are generally defined by means of an adjacency rule, prescribing which unordered pairs of vertices are edges and which are not. A number of such examples are given in the exercises at the end of this section and in Section 1.3.

For the sake of clarity, we observe certain conventions in representing graphs by diagrams: we do not allow an edge to intersect itself, nor let an edge pass through a vertex that is not an end of the edge; clearly, this is always possible. However, two edges may intersect at a point that does not correspond to a vertex, as in the drawings of the first two graphs in Figure 1.2. A graph which can be drawn in the plane in such a way that edges meet only at points corresponding to their common ends is called a *planar graph*, and such a drawing is called a *planar embedding* of the graph. For instance, the graphs G and H of Examples 1 and 2 are both planar, even though there are crossing edges in the particular drawing of G shown in Figure 1.1. The first two graphs in Figure 1.2, on the other hand, are not planar, as proved later.

Although not all graphs are planar, every graph can be drawn on some surface so that its edges intersect only at their ends. Such a drawing is called an *embedding* of the graph on the surface. Figure 1.21 provides an example of an embedding of a graph on the torus. Embeddings of graphs on surfaces are discussed in Chapter 3 and, more thoroughly, in Chapter 10.

Incidence and Adjacency Matrices

Although drawings are a convenient means of specifying graphs, they are clearly not suitable for storing graphs in computers, or for applying mathematical methods to study their properties. For these purposes, we consider two matrices associated with a graph, its incidence matrix and its adjacency matrix.

Let G be a graph, with vertex set V and edge set E. The *incidence matrix* of G is the $n \times m$ matrix $\mathbf{M}_G := (m_{ve})$, where m_{ve} is the number of times (0, 1, or 2) that vertex v and edge e are incident. Clearly, the incidence matrix is just another way of specifying the graph.

The *adjacency matrix* of G is the $n \times n$ matrix $\mathbf{A}_G := (a_{uv})$, where a_{uv} is the number of edges joining vertices u and v, each loop counting as two edges. Incidence and adjacency matrices of the graph G of Figure 1.1 are shown in Figure 1.5.

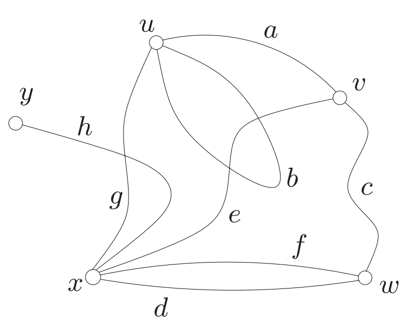

G

	a	b	c	d	e	f	g	h
u	1	2	0	0	0	0	1	0
v	1	0	1	0	1	0	0	0
w	0	0	1	1	0	1	0	0
x	0	0	0	1	1	1	1	1
y	0	0	0	0	0	0	0	1

$\mathbf{M}$

	u	v	w	x	y
u	2	1	0	1	0
v	1	0	1	1	0
w	0	1	0	2	0
x	1	1	2	0	1
y	0	0	0	1	0

$\mathbf{A}$

Fig. 1.5. Incidence and adjacency matrices of a graph

Because most graphs have many more edges than vertices, the adjacency matrix of a graph is generally much smaller than its incidence matrix and thus requires less storage space. When dealing with simple graphs, an even more compact representation is possible. For each vertex v, the neighbours of v are listed in some order. A list $(N(v) : v \in V)$ of these lists is called an *adjacency list* of the graph. Simple graphs are usually stored in computers as adjacency lists.

When G is a bipartite graph, as there are no edges joining pairs of vertices belonging to the same part of its bipartition, a matrix of smaller size than the adjacency matrix may be used to record the numbers of edges joining pairs of vertices. Suppose that $G[X,Y]$ is a bipartite graph, where $X := \{x_1, x_2, \ldots, x_r\}$ and $Y := \{y_1, y_2, \ldots, y_s\}$. We define the *bipartite adjacency matrix* of G to be the $r \times s$ matrix $\mathbf{B}_G = (b_{ij})$, where b_{ij} is the number of edges joining x_i and y_j.

Vertex Degrees

The *degree* of a vertex v in a graph G, denoted by $d_G(v)$, is the number of edges of G incident with v, each loop counting as two edges. In particular, if G is a simple graph, $d_G(v)$ is the number of neighbours of v in G. A vertex of degree zero is called an *isolated vertex*. We denote by $\delta(G)$ and $\Delta(G)$ the minimum and maximum degrees of the vertices of G, and by $d(G)$ their *average degree*, $\frac{1}{n}\sum_{v\in V} d(v)$. The following theorem establishes a fundamental identity relating the degrees of the vertices of a graph and the number of its edges.

Theorem 1.1 *For any graph G,*

$$\sum_{v\in V} d(v) = 2m \tag{1.1}$$

Proof Consider the incidence matrix $\mathbf{M}$ of G. The sum of the entries in the row corresponding to vertex v is precisely $d(v)$. Therefore $\sum_{v\in V} d(v)$ is just the sum of all the entries in $\mathbf{M}$. But this sum is also $2m$, because each of the m column sums of $\mathbf{M}$ is 2, each edge having two ends. □

Corollary 1.2 *In any graph, the number of vertices of odd degree is even.*

Proof Consider equation (1.1) modulo 2. We have

$$d(v) \equiv \begin{cases} 1 \pmod 2 & \text{if } d(v) \text{ is odd,} \\ 0 \pmod 2 & \text{if } d(v) \text{ is even.} \end{cases}$$

Thus, modulo 2, the left-hand side is congruent to the number of vertices of odd degree, and the right-hand side is zero. The number of vertices of odd degree is therefore congruent to zero modulo 2. □

A graph G is *k-regular* if $d(v) = k$ for all $v \in V$; a *regular graph* is one that is k-regular for some k. For instance, the complete graph on n vertices is $(n-1)$-regular, and the complete bipartite graph with k vertices in each part is k-regular. For $k = 0, 1$ and 2, k-regular graphs have very simple structures and are easily characterized (Exercise 1.1.5). By contrast, 3-regular graphs can be remarkably complex. These graphs, also referred to as *cubic* graphs, play a prominent role in graph theory. We present a number of interesting examples of such graphs in the next section.

Proof Technique: Counting in Two Ways

In proving Theorem 1.1, we used a common proof technique in combinatorics, known as *counting in two ways*. It consists of considering a suitable matrix and computing the sum of its entries in two different ways: firstly as the sum of its row sums, and secondly as the sum of its column sums. Equating these two quantities results in an identity. In the case of Theorem 1.1, the matrix we considered was the incidence matrix of G. In order to prove the identity of Exercise 1.1.9a, the appropriate matrix to consider is the bipartite adjacency matrix of the bipartite graph $G[X, Y]$. In both these cases, the choice of the appropriate matrix is fairly obvious. However, in some cases, making the right choice requires ingenuity.

Note that an upper bound on the sum of the column sums of a matrix is clearly also an upper bound on the sum of its row sums (and *vice versa*). The method of counting in two ways may therefore be adapted to establish inequalities. The proof of the following proposition illustrates this idea.

Proposition 1.3 *Let $G[X, Y]$ be a bipartite graph without isolated vertices such that $d(x) \geq d(y)$ for all $xy \in E$, where $x \in X$ and $y \in Y$. Then $|X| \leq |Y|$, with equality if and only if $d(x) = d(y)$ for all $xy \in E$.*

Proof The first assertion follows if we can find a matrix with $|X|$ rows and $|Y|$ columns in which each row sum is one and each column sum is at most one. Such a matrix can be obtained from the bipartite adjacency matrix $\mathbf{B}$ of $G[X, Y]$ by dividing the row corresponding to vertex x by $d(x)$, for each $x \in X$. (This is possible since $d(x) \neq 0$.) Because the sum of the entries of $\mathbf{B}$ in the row corresponding to x is $d(x)$, all row sums of the resulting matrix $\widetilde{\mathbf{B}}$ are equal to one. On the other hand, the sum of the entries in the column of $\widetilde{\mathbf{B}}$ corresponding to vertex y is $\sum 1/d(x)$, the sum being taken over all edges xy incident to y, and this sum is at most one because $1/d(x) \leq 1/d(y)$ for each edge xy, by hypothesis, and because there are $d(y)$ edges incident to y.

The above argument may be expressed more concisely as follows.

$$|X| = \sum_{x \in X} \sum_{\substack{y \in Y \\ xy \in E}} \frac{1}{d(x)} = \sum_{\substack{x \in X \\ y \in Y}} \sum_{xy \in E} \frac{1}{d(x)} \leq \sum_{\substack{x \in X \\ y \in Y}} \sum_{xy \in E} \frac{1}{d(y)} = \sum_{y \in Y} \sum_{\substack{x \in X \\ xy \in E}} \frac{1}{d(y)} = |Y|$$

Furthermore, if $|X| = |Y|$, the middle inequality must be an equality, implying that $d(x) = d(y)$ for all $xy \in E$. $\square$

An application of this proof technique to a problem in set theory about geometric configurations is described in Exercise 1.3.15.

Exercises

1.1.1 Let G be a simple graph. Show that $m \leq \binom{n}{2}$, and determine when equality holds.

1.1.2 Let $G[X, Y]$ be a simple bipartite graph, where $|X| = r$ and $|Y| = s$.

a) Show that $m \leq rs$.
b) Deduce that $m \leq n^2/4$.
c) Describe the simple bipartite graphs G for which equality holds in (b).

$\star$**1.1.3** Show that:

a) every path is bipartite,
b) a cycle is bipartite if and only if its length is even.

1.1.4 Show that, for any graph G, $\delta(G) \leq d(G) \leq \Delta(G)$.

1.1.5 For $k = 0, 1, 2$, characterize the k-regular graphs.

1.1.6

a) Show that, in any group of two or more people, there are always two who have exactly the same number of friends within the group.
b) Describe a group of five people, any two of whom have exactly one friend in common. Can you find a group of four people with this same property?

1.1.7 n-Cube
The *n-cube* Q_n $(n \geq 1)$ is the graph whose vertex set is the set of all n-tuples of 0s and 1s, where two n-tuples are adjacent if they differ in precisely one coordinate.

a) Draw Q_1, Q_2, Q_3, and Q_4.
b) Determine $v(Q_n)$ and $e(Q_n)$.
c) Show that Q_n is bipartite for all $n \geq 1$.

1.1.8 The *boolean lattice* BL_n $(n \geq 1)$ is the graph whose vertex set is the set of all subsets of $\{1, 2, \ldots, n\}$, where two subsets X and Y are adjacent if their symmetric difference has precisely one element.

a) Draw BL_1, BL_2, BL_3, and BL_4.
b) Determine $v(BL_n)$ and $e(BL_n)$.
c) Show that BL_n is bipartite for all $n \geq 1$.

$\star$**1.1.9** Let $G[X, Y]$ be a bipartite graph.

a) Show that $\sum_{v \in X} d(v) = \sum_{v \in Y} d(v)$.
b) Deduce that if G is k-regular, with $k \geq 1$, then $|X| = |Y|$.

⋆**1.1.10** k-Partite Graph
A *k-partite graph* is one whose vertex set can be partitioned into k subsets, or *parts*, in such a way that no edge has both ends in the same part. (Equivalently, one may think of the vertices as being colourable by k colours so that no edge joins two vertices of the same colour.) Let G be a simple k-partite graph with parts of sizes $a_1, a_2, \dots, a_k$. Show that $m \leq \frac{1}{2}\sum_{i=1}^{k} a_i(n - a_i)$.

⋆**1.1.11** Turán Graph
A k-partite graph is *complete* if any two vertices in different parts are adjacent. A simple complete k-partite graph on n vertices whose parts are of equal or almost equal sizes (that is, $\lfloor n/k \rfloor$ or $\lceil n/k \rceil$) is called a *Turán graph* and denoted $T_{k,n}$.

a) Show that $T_{k,n}$ has more edges than any other simple complete k-partite graph on n vertices.
b) Determine $e(T_{k,n})$.

1.1.12

a) Show that if G is simple and $m > \binom{n-1}{2}$, then G is connected.
b) For $n > 1$, find a disconnected simple graph G with $m = \binom{n-1}{2}$.

1.1.13

a) Show that if G is simple and $\delta > \frac{1}{2}(n-2)$, then G is connected.
b) For n even, find a disconnected $\frac{1}{2}(n-2)$-regular simple graph.

1.1.14 For a simple graph G, show that the diagonal entries of both $\mathbf{A}^2$ and $\mathbf{MM}^t$ (where $\mathbf{M}^t$ denotes the transpose of $\mathbf{M}$) are the degrees of the vertices of G.

1.1.15 Show that the rank over $GF(2)$ of the incidence matrix of a graph G is at most $n-1$, with equality if and only if G is connected.

1.1.16 Degree Sequence
If G has vertices $v_1, v_2, \dots, v_n$, the sequence $(d(v_1), d(v_2), \dots, d(v_n))$ is called a *degree sequence* of G. Let $\mathbf{d} := (d_1, d_2, \dots, d_n)$ be a nonincreasing sequence of nonnegative integers, that is, $d_1 \geq d_2 \geq \cdots \geq d_n \geq 0$. Show that:

a) there is a graph with degree sequence $\mathbf{d}$ if and only if $\sum_{i=1}^{n} d_i$ is even,
b) there is a loopless graph with degree sequence $\mathbf{d}$ if and only if $\sum_{i=1}^{n} d_i$ is even and $d_1 \leq \sum_{i=2}^{n} d_i$.

1.1.17 Complement of a Graph
Let G be a simple graph. The *complement* $\overline{G}$ of G is the simple graph whose vertex set is V and whose edges are the pairs of nonadjacent vertices of G.

a) Express the degree sequence of $\overline{G}$ in terms of the degree sequence of G.
b) Show that if G is disconnected, then $\overline{G}$ is connected. Is the converse true?

1.1.18 Graphic Sequence
A sequence $\mathbf{d} = (d_1, d_2, \ldots, d_n)$ is *graphic* if there is a simple graph with degree sequence $\mathbf{d}$. Show that:

a) the sequences $(7, 6, 5, 4, 3, 3, 2)$ and $(6, 6, 5, 4, 3, 3, 1)$ are not graphic,
b) if $\mathbf{d} = (d_1, d_2, \ldots, d_n)$ is graphic and $d_1 \geq d_2 \geq \cdots \geq d_n$, then $\sum_{i=1}^{n} d_i$ is even and
$$\sum_{i=1}^{k} d_i \leq k(k-1) + \sum_{i=k+1}^{n} \min\{k, d_i\}, \quad 1 \leq k \leq n$$
(Erdős and Gallai (1960) showed that these necessary conditions for a sequence to be graphic are also sufficient.)

1.1.19 Let $\mathbf{d} = (d_1, d_2, \ldots, d_n)$ be a nonincreasing sequence of nonnegative integers. Set $\mathbf{d}' := (d_2 - 1, d_3 - 1, \ldots, d_{d_1+1} - 1, d_{d_1+2}, \ldots, d_n)$.

a) Show that $\mathbf{d}$ is graphic if and only if $\mathbf{d}'$ is graphic.
b) Using (a), describe an algorithm which accepts as input a nonincreasing sequence $\mathbf{d}$ of nonnegative integers, and returns either a simple graph with degree sequence $\mathbf{d}$, if such a graph exists, or else a proof that $\mathbf{d}$ is not graphic.

(V. Havel and S.L. Hakimi)

1.1.20 Let S be a set of n points in the plane, the distance between any two of which is at least one. Show that there are at most $3n$ pairs of points of S at distance exactly one.

1.1.21 Eigenvalues of a Graph
Recall that the eigenvalues of a square matrix $\mathbf{A}$ are the roots of its characteristic polynomial $\det(\mathbf{A} - x\mathbf{I})$. An *eigenvalue* of a graph is an eigenvalue of its adjacency matrix. Likewise, the *characteristic polynomial* of a graph is the characteristic polynomial of its adjacency matrix. Show that:

a) every eigenvalue of a graph is real,
b) every rational eigenvalue of a graph is integral.

1.1.22

a) Let G be a k-regular graph. Show that:
   i) $\mathbf{M}\mathbf{M}^t = \mathbf{A} + k\mathbf{I}$, where $\mathbf{I}$ is the $n \times n$ identity matrix,
   ii) k is an eigenvalue of G, with corresponding eigenvector $\mathbf{1}$, the n-vector in which each entry is 1.
b) Let G be a complete graph of order n. Denote by $\mathbf{J}$ the $n \times n$ matrix all of whose entries are 1. Show that:
   i) $\mathbf{A} = \mathbf{J} - \mathbf{I}$,
   ii) $\det(\mathbf{J} - (1 + \lambda)\mathbf{I}) = (1 + \lambda - n)(1 + \lambda)^{n-1}$.
c) Derive from (b) the eigenvalues of a complete graph and their multiplicities, and determine the corresponding eigenspaces.

1.1.23 Let G be a simple graph.

a) Show that $\overline{G}$ has adjacency matrix $\mathbf{J} - \mathbf{I} - \mathbf{A}$.
b) Suppose now that G is k-regular.
   i) Deduce from Exercise 1.1.22 that $n - k - 1$ is an eigenvalue of $\overline{G}$, with corresponding eigenvector $\mathbf{1}$.
   ii) Show that if λ is an eigenvalue of G different from k, then $-1 - \lambda$ is an eigenvalue of $\overline{G}$, with the same multiplicity. (Recall that eigenvectors corresponding to distinct eigenvalues of a real symmetric matrix are orthogonal.)

1.1.24 Show that:

a) no eigenvalue of a graph G has absolute value greater than Δ,
b) if G is a connected graph and Δ is an eigenvalue of G, then G is regular,
c) if G is a connected graph and $-\Delta$ is an eigenvalue of G, then G is both regular and bipartite.

1.1.25 Strongly Regular Graph
A simple graph G which is neither empty nor complete is said to be *strongly regular* with parameters (v, k, λ, μ) if:

▷ $v(G) = v$,
▷ G is k-regular,
▷ any two adjacent vertices of G have λ common neighbours,
▷ any two nonadjacent vertices of G have μ common neighbours.

Let G be a strongly regular graph with parameters (v, k, λ, μ). Show that:

a) $\overline{G}$ is strongly regular,
b) $k(k - \lambda - 1) = (v - k - 1)\mu$,
c) $\mathbf{A}^2 = k\,\mathbf{I} + \lambda\,\mathbf{A} + \mu\,(\mathbf{J} - \mathbf{I} - \mathbf{A})$.

1.2 Isomorphisms and Automorphisms

Isomorphisms

Two graphs G and H are *identical*, written $G = H$, if $V(G) = V(H)$, $E(G) = E(H)$, and $\psi_G = \psi_H$. If two graphs are identical, they can clearly be represented by identical diagrams. However, it is also possible for graphs that are not identical to have essentially the same diagram. For example, the graphs G and H in Figure 1.6 can be represented by diagrams which look exactly the same, as the second drawing of H shows; the sole difference lies in the labels of their vertices and edges. Although the graphs G and H are not identical, they do have identical structures, and are said to be isomorphic.

In general, two graphs G and H are *isomorphic*, written $G \cong H$, if there are bijections $\theta : V(G) \to V(H)$ and $\phi : E(G) \to E(H)$ such that $\psi_G(e) = uv$ if and only if $\psi_H(\phi(e)) = \theta(u)\theta(v)$; such a pair of mappings is called an *isomorphism* between G and H.

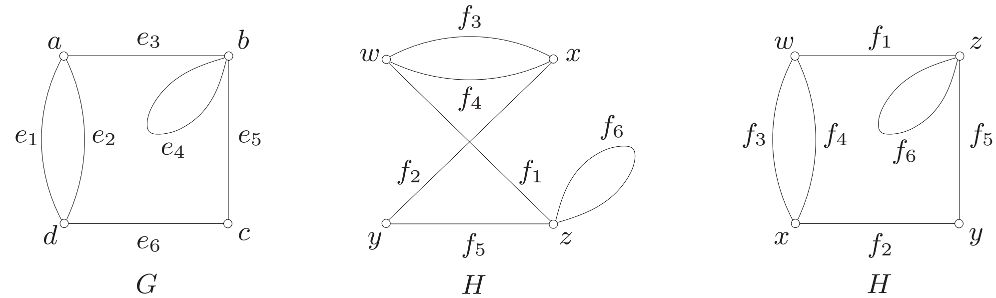

Fig. 1.6. Isomorphic graphs

In order to show that two graphs are isomorphic, one must indicate an isomorphism between them. The pair of mappings (θ, ϕ) defined by

$$\theta := \begin{pmatrix} a & b & c & d \\ w & z & y & x \end{pmatrix} \qquad \phi := \begin{pmatrix} e_1 & e_2 & e_3 & e_4 & e_5 & e_6 \\ f_3 & f_4 & f_1 & f_6 & f_5 & f_2 \end{pmatrix}$$

is an isomorphism between the graphs G and H in Figure 1.6.

In the case of simple graphs, the definition of isomorphism can be stated more concisely, because if (θ, ϕ) is an isomorphism between simple graphs G and H, the mapping ϕ is completely determined by θ; indeed, $\phi(e) = \theta(u)\theta(v)$ for any edge $e = uv$ of G. Thus one may define an isomorphism between two simple graphs G and H as a bijection $\theta : V(G) \to V(H)$ which preserves adjacency (that is, the vertices u and v are adjacent in G if and only if their images $\theta(u)$ and $\theta(v)$ are adjacent in H).

Consider, for example, the graphs G and H in Figure 1.7.

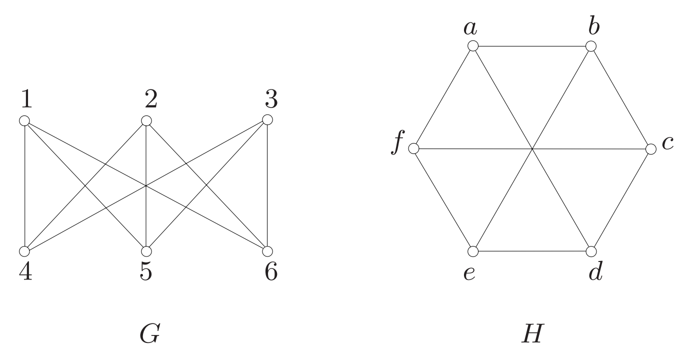

Fig. 1.7. Isomorphic simple graphs

The mapping

$$\theta := \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ b & d & f & c & e & a \end{pmatrix}$$

is an isomorphism between G and H, as is

$$\theta' := \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ a & c & e & d & f & b \end{pmatrix}$$

Isomorphic graphs clearly have the same numbers of vertices and edges. On the other hand, equality of these parameters does not guarantee isomorphism. For instance, the two graphs shown in Figure 1.8 both have eight vertices and twelve edges, but they are not isomorphic. To see this, observe that the graph G has four mutually nonadjacent vertices, v_1, v_3, v_6, and v_8. If there were an isomorphism θ between G and H, the vertices $\theta(v_1)$, $\theta(v_3)$, $\theta(v_6)$, and $\theta(v_8)$ of H would likewise be mutually nonadjacent. But it can readily be checked that no four vertices of H are mutually nonadjacent. We deduce that G and H are not isomorphic.

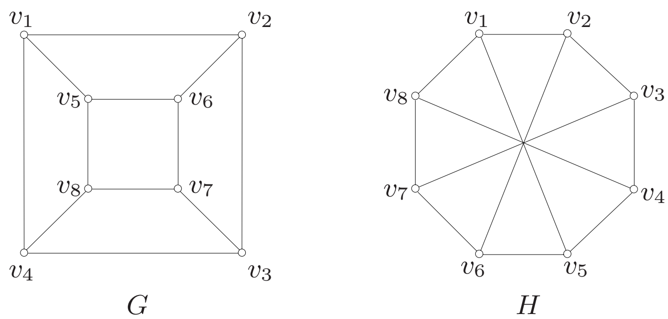

Fig. 1.8. Nonisomorphic graphs

It is clear from the foregoing discussion that if two graphs are isomorphic, then they are either identical or differ merely in the names of their vertices and edges, and thus have the same structure. Because it is primarily in structural properties that we are interested, we often omit labels when drawing graphs; formally, we may define an *unlabelled graph* as a representative of an equivalence class of isomorphic graphs. We assign labels to vertices and edges in a graph mainly for the purpose of referring to them (in proofs, for instance).

Up to isomorphism, there is just one complete graph on n vertices, denoted K_n. Similarly, given two positive integers m and n, there is a unique complete bipartite graph with parts of sizes m and n (again, up to isomorphism), denoted $K_{m,n}$. In this notation, the graphs in Figure 1.2 are K_5, $K_{3,3}$, and $K_{1,5}$, respectively. Likewise, for any positive integer n, there is a unique path on n vertices and a unique cycle on n vertices. These graphs are denoted P_n and C_n, respectively. The graphs depicted in Figure 1.3 are P_4 and C_5.

Testing for Isomorphism

Given two graphs on n vertices, it is certainly possible in principle to determine whether they are isomorphic. For instance, if G and H are simple, one could just consider each of the $n!$ bijections between $V(G)$ and $V(H)$ in turn, and check whether it is an isomorphism between the two graphs. If the graphs happen to be isomorphic, an isomorphism might (with luck) be found quickly. On the other hand, if they are not isomorphic, one would need to check all $n!$ bijections to discover this fact. Unfortunately, even for moderately small values of n (such as $n = 100$), the number $n!$ is unmanageably large (indeed, larger than the number of particles in the universe!), so this 'brute force' approach is not feasible. Of course, if the graphs are not regular, the number of bijections to be checked will be smaller, as an isomorphism must map each vertex to a vertex of the same degree (Exercise 1.2.1a). Nonetheless, except in particular cases, this restriction does not serve to reduce their number sufficiently. Indeed, no efficient generally applicable procedure for testing isomorphism is known. However, by employing powerful group-theoretic methods, Luks (1982) devised an efficient isomorphism-testing algorithm for cubic graphs and, more generally, for graphs of bounded maximum degree.

There is another important matter related to algorithmic questions such as graph isomorphism. Suppose that two simple graphs G and H are isomorphic. It might not be easy to find an isomorphism between them, but once such an isomorphism θ has been found, it is a simple matter to verify that θ is indeed an isomorphism: one need merely check that, for each of the $\binom{n}{2}$ pairs uv of vertices of G, $uv \in E(G)$ if and only if $\theta(u)\theta(v) \in E(H)$. On the other hand, if G and H happen not to be isomorphic, how can one verify this fact, short of checking all possible bijections between $V(G)$ and $V(H)$? In certain cases, one might be able to show that G and H are not isomorphic by isolating some structural property of G that is not shared by H, as we did for the graphs G and H of Figure 1.8. However, in general, verifying that two nonisomorphic graphs are indeed not isomorphic seems to be just as hard as determining in the first place whether they are isomorphic or not.

Automorphisms

An *automorphism* of a graph is an isomorphism of the graph to itself. In the case of a simple graph, an automorphism is just a permutation α of its vertex set which preserves adjacency: if uv is an edge then so is $\alpha(u)\alpha(v)$.

The automorphisms of a graph reflect its symmetries. For example, if u and v are two vertices of a simple graph, and if there is an automorphism α which maps u to v, then u and v are alike in the graph, and are referred to as *similar* vertices. Graphs in which all vertices are similar, such as the complete graph K_n, the complete bipartite graph $K_{n,n}$ and the n-cube Q_n, are called *vertex-transitive*. Graphs in which no two vertices are similar are called *asymmetric*; these are the graphs which have only the identity permutation as automorphism (see Exercise 1.2.14).

Particular drawings of a graph may often be used to display its symmetries. As an example, consider the three drawings shown in Figure 1.9 of the *Petersen graph*, a graph which turns out to have many special properties. (We leave it as an exercise (1.2.5) that they are indeed drawings of one and the same graph.) The first drawing shows that the five vertices of the outer pentagon are similar (under rotational symmetry), as are the five vertices of the inner pentagon. The third drawing exhibits six similar vertices (under reflective or rotational symmetry), namely the vertices of the outer hexagon. Combining these two observations, we conclude that all ten vertices of the Petersen graph are similar, and thus that the graph is vertex-transitive.

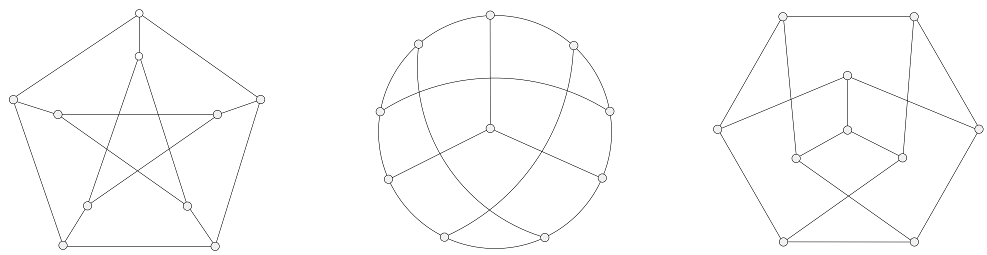

Fig. 1.9. Three drawings of the Petersen graph

We denote the set of all automorphisms of a graph G by $\mathrm{Aut}(G)$, and their number by $\mathrm{aut}(G)$. It can be verified that $\mathrm{Aut}(G)$ is a group under the operation of composition (Exercise 1.2.9). This group is called the *automorphism group* of G. The automorphism group of K_n is the symmetric group S_n, consisting of all permutations of its vertex set. In general, for any simple graph G on n vertices, $\mathrm{Aut}(G)$ is a subgroup of S_n. For instance, the automorphism group of C_n is D_n, the dihedral group on n elements (Exercise 1.2.10).

Labelled Graphs

As we have seen, the edge set E of a simple graph $G = (V, E)$ is usually considered to be a subset of $\binom{V}{2}$, the set of all 2-subsets of V; edge labels may then be omitted in drawings of such graphs. A simple graph whose vertices are labelled, but whose edges are not, is referred to as a *labelled simple graph*. If $|V| = n$, there are $2^{\binom{n}{2}}$ distinct subsets of $\binom{V}{2}$, so $2^{\binom{n}{2}}$ labelled simple graphs with vertex set V. We denote by $\mathcal{G}_n$ the set of labelled simple graphs with vertex set $V := \{v_1, v_2, \ldots, v_n\}$. The set $\mathcal{G}_3$ is shown in Figure 1.10.

A priori, there are $n!$ ways of assigning the labels $v_1, v_2, \ldots, v_n$ to the vertices of an unlabelled simple graph on n vertices. But two of these will yield the same labelled graph if there is an automorphism of the graph mapping one labelling to the other. For example, all six labellings of K_3 result in the same element of $\mathcal{G}_3$, whereas the six labellings of P_3 yield three distinct labelled graphs, as shown in Figure 1.10. The number of distinct labellings of a given unlabelled simple graph G on n vertices is, in fact, $n!/\mathrm{aut}(G)$ (Exercise 1.2.15). Consequently,

$$\sum_G \frac{n!}{\mathrm{aut}(G)} = 2^{\binom{n}{2}}$$

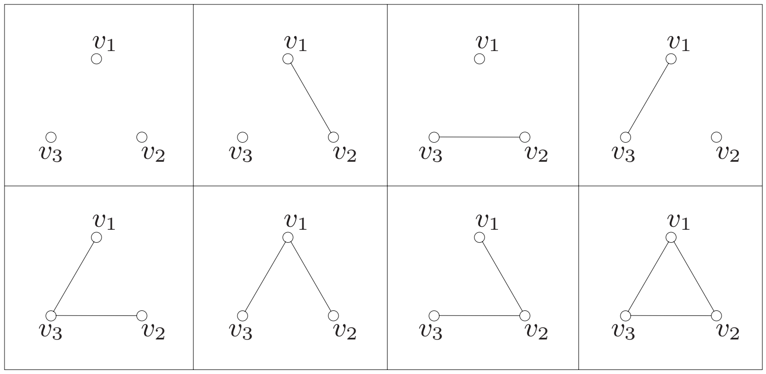

Fig. 1.10. The eight labelled graphs on three vertices

where the sum is over all unlabelled simple graphs on n vertices. In particular, the number of unlabelled simple graphs on n vertices is at least

$$\left\lceil \frac{2^{\binom{n}{2}}}{n!} \right\rceil \tag{1.2}$$

For small values of n, this bound is not particularly good. For example, there are four unlabelled simple graphs on three vertices, but the bound (1.2) is just two. Likewise, the number of unlabelled simple graphs on four vertices is eleven (Exercise 1.2.6), whereas the bound given by (1.2) is three. Nonetheless, when n is large, this bound turns out to be a good approximation to the actual number of unlabelled simple graphs on n vertices because the vast majority of graphs are asymmetric (see Exercise 1.2.15d).

Exercises

1.2.1

a) Show that any isomorphism between two graphs maps each vertex to a vertex of the same degree.
b) Deduce that isomorphic graphs necessarily have the same (nonincreasing) degree sequence.

1.2.2 Show that the graphs in Figure 1.11 are not isomorphic (even though they have the same degree sequence).

1.2.3 Let G be a connected graph. Show that every graph which is isomorphic to G is connected.

1.2.4 Determine:

a) the number of isomorphisms between the graphs G and H of Figure 1.7,

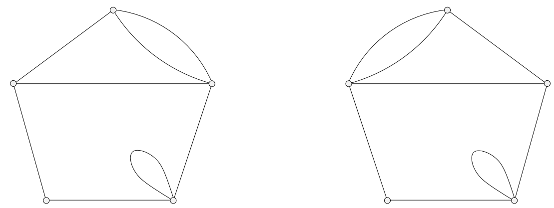

Fig. 1.11. Nonisomorphic graphs

b) the number of automorphisms of each of these graphs.

⋆**1.2.5** Show that the three graphs in Figure 1.9 are isomorphic.

1.2.6 Draw:

a) all the labelled simple graphs on four vertices,
b) all the unlabelled simple graphs on four vertices,
c) all the unlabelled simple cubic graphs on eight or fewer vertices.

1.2.7 Show that the n-cube Q_n and the boolean lattice BL_n (defined in Exercises 1.1.7 and 1.1.8) are isomorphic.

1.2.8 Show that two simple graphs G and H are isomorphic if and only if there exists a permutation matrix $\mathbf{P}$ such that $\mathbf{A}_H = \mathbf{P}\mathbf{A}_G\mathbf{P}^t$.

1.2.9 Show that $\mathrm{Aut}(G)$ is a group under the operation of composition.

1.2.10

a) Show that, for $n \geq 2$, $\mathrm{Aut}(P_n) \cong S_2$ and $\mathrm{Aut}(C_n) = D_n$, the dihedral group on n elements (where $\cong$ denotes isomorphism of groups; see, for example, Herstein (1996)).
b) Determine the automorphism group of the complete bipartite graph $K_{m,n}$.

1.2.11 Show that, for any simple graph G, $\mathrm{Aut}(G) = \mathrm{Aut}(\overline{G})$.

1.2.12 Consider the subgroup Γ of S_3 with elements (1)(2)(3), (123), and (132).

a) Show that there is no simple graph whose automorphism group is Γ.
b) Find a simple graph whose automorphism group is isomorphic to Γ.
(Frucht (1938) showed that every abstract group is isomorphic to the automorphism group of some simple graph.)

1.2.13 Orbits of a Graph

a) Show that similarity is an equivalence relation on the vertex set of a graph.
b) The equivalence classes with respect to similarity are called the *orbits* of the graph. Determine the orbits of the graphs in Figure 1.12.

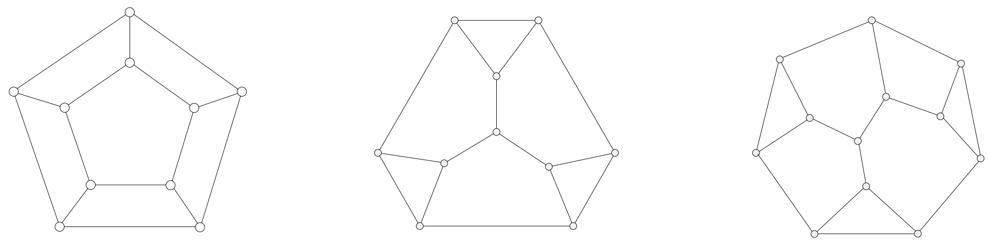

(a) (b) (c)

Fig. 1.12. Determine the orbits of these graphs (Exercise 1.2.13)

1.2.14

a) Show that there is no asymmetric simple graph on five or fewer vertices.
b) For each $n \geq 6$, find an asymmetric simple graph on n vertices.

1.2.15 Let G and H be isomorphic members of $\mathcal{G}_n$, let θ be an isomorphism between G and H, and let α be an automorphism of G.

a) Show that $\theta\alpha$ is an isomorphism between G and H.
b) Deduce that the set of all isomorphisms between G and H is the coset $\theta\mathrm{Aut}(G)$ of $\mathrm{Aut}(G)$.
c) Deduce that the number of labelled graphs isomorphic to G is equal to $n!/\mathrm{aut}(G)$.
d) Erdős and Rényi (1963) have shown that almost all simple graphs are asymmetric (that is, the proportion of simple graphs on n vertices that are asymmetric tends to one as n tends to infinity). Using this fact, deduce from (c) that the number of unlabelled graphs on n vertices is asymptotically equal to $2^{\binom{n}{2}}/n!$ (G. Pólya)

1.2.16 Self-Complementary Graph
A simple graph is *self-complementary* if it is isomorphic to its complement. Show that:

a) each of the graphs P_4 and C_5 (shown in Figure 1.3) is self-complementary,
b) every self-complementary graph is connected,
c) if G is self-complementary, then $n \equiv 0, 1 \pmod 4$,
d) every self-complementary graph on $4k+1$ vertices has a vertex of degree $2k$.

1.2.17 Edge-Transitive Graph
A simple graph is *edge-transitive* if, for any two edges uv and xy, there is an automorphism α such that $\alpha(u)\alpha(v) = xy$.

a) Find a graph which is vertex-transitive but not edge-transitive.
b) Show that any graph without isolated vertices which is edge-transitive but not vertex-transitive is bipartite. (E. Dauber)

1.2.18 The Folkman Graph

a) Show that the graph shown in Figure 1.13a is edge-transitive but not vertex-transitive.

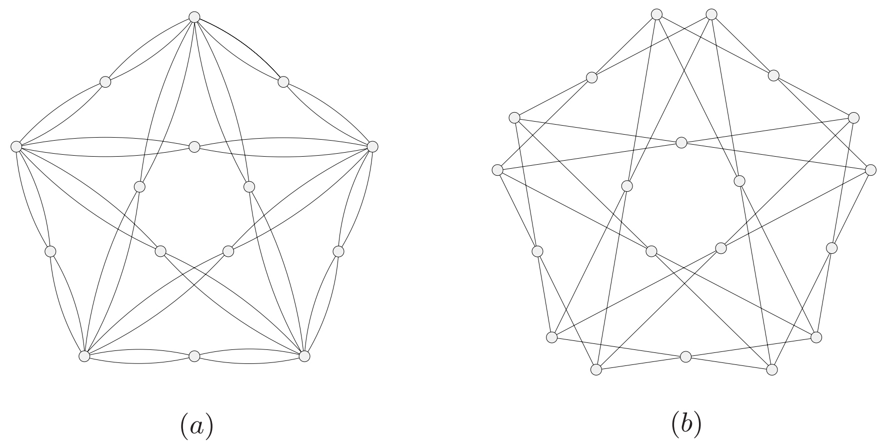

Fig. 1.13. Construction of the Folkman graph

b) The *Folkman graph*, depicted in Figure 1.13b, is the 4-regular graph obtained from the graph of Figure 1.13a by replacing each vertex v of degree eight by two vertices of degree four, both of which have the same four neighbours as v. Show that the Folkman graph is edge-transitive but not vertex-transitive. (J. Folkman)

1.2.19 Generalized Petersen Graph
Let k and n be positive integers, with $n > 2k$. The *generalized Petersen graph* $P_{k,n}$ is the simple graph with vertices $x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_n$, and edges $x_i x_{i+1}, y_i y_{i+k}, x_i y_i$, $1 \leq i \leq n$, indices being taken modulo n. (Note that $P_{2,5}$ is the Petersen graph.)

a) Draw the graphs $P_{2,7}$ and $P_{3,8}$.
b) Which of these two graphs are vertex-transitive, and which are edge-transitive?

1.2.20 Show that if G is simple and the eigenvalues of $\mathbf{A}$ are distinct, then every automorphism of G is of order one or two. (A. Mowshowitz)

1.3 Graphs Arising from Other Structures

As remarked earlier, interesting graphs can often be constructed from geometric and algebraic objects. Such constructions are often quite straightforward, but in some instances they rely on experience and insight.

Polyhedral Graphs

A *polyhedral graph* is the 1-skeleton of a polyhedron, that is, the graph whose vertices and edges are just the vertices and edges of the polyhedron, with the same incidence relation. In particular, the five platonic solids (the tetrahedron, the cube, the octahedron, the dodecahedron, and the icosahedron) give rise to the five *platonic graphs* shown in Figure 1.14. For classical polyhedra such as these, we give the graph the same name as the polyhedron from which it is derived.

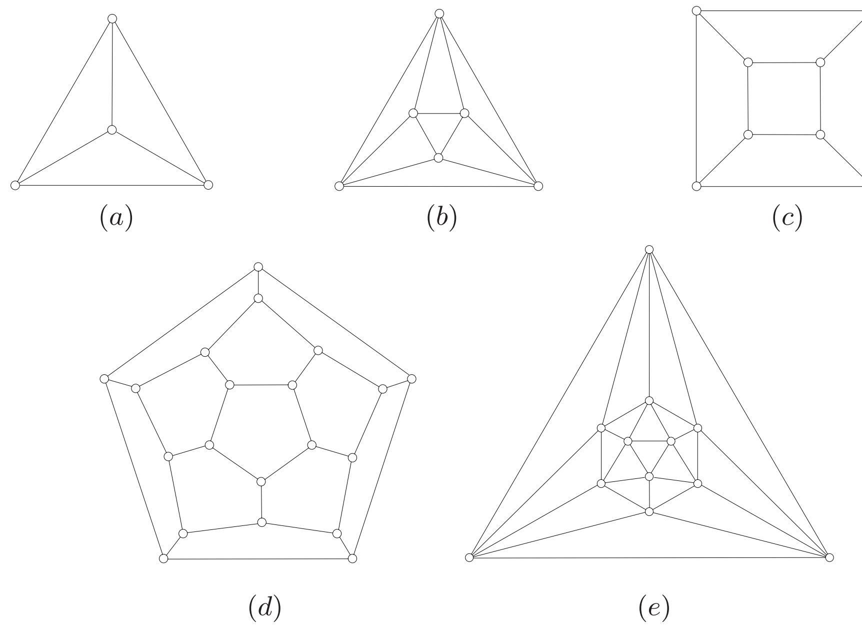

Fig. 1.14. The five platonic graphs: (a) the tetrahedron, (b) the octahedron, (c) the cube, (d) the dodecahedron, (e) the icosahedron

Set Systems and Hypergraphs

A *set system* is an ordered pair $(V, \mathcal{F})$, where V is a set of elements and $\mathcal{F}$ is a family of subsets of V. Note that when $\mathcal{F}$ consists of pairs of elements of V, the set system $(V, \mathcal{F})$ is a loopless graph. Thus set systems can be thought of as generalizations of graphs, and are usually referred to as *hypergraphs*, particularly when one seeks to extend properties of graphs to set systems (see Berge (1973)). The elements of V are then called the *vertices* of the hypergraph, and the elements of $\mathcal{F}$ its *edges* or *hyperedges*. A hypergraph is k*-uniform* if each edge is a k-set (a set of k elements). As we show below, set systems give rise to graphs in two principal ways: incidence graphs and intersection graphs.

Many interesting examples of hypergraphs are provided by geometric configurations. A *geometric configuration* $(P, \mathcal{L})$ consists of a finite set P of elements called *points*, and a finite family $\mathcal{L}$ of subsets of P called *lines*, with the property that at most one line contains any given pair of points. Two classical examples of geometric configurations are the *Fano plane* and the *Desargues configuration*. These two configurations are shown in Figure 1.15. In both cases, each line consists of three points. These configurations thus give rise to 3-uniform hypergraphs; the *Fano hypergraph* has seven vertices and seven edges, the *Desargues hypergraph* ten vertices and ten edges.

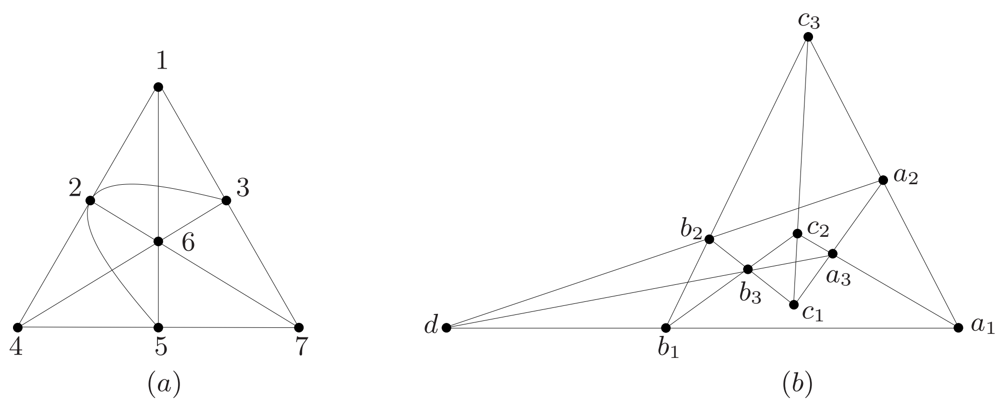

Fig. 1.15. (a) The Fano plane, and (b) the Desargues configuration

The Fano plane is the simplest of an important family of geometric configurations, the *projective planes* (see Exercise 1.3.13). The Desargues configuration arises from a well-known theorem in projective geometry. Other examples of interesting geometric configurations are described in Coxeter (1950) and Godsil and Royle (2001).

Incidence Graphs

A natural graph associated with a set system $H = (V, \mathcal{F})$ is the bipartite graph $G[V, \mathcal{F}]$, where $v \in V$ and $F \in \mathcal{F}$ are adjacent if $v \in F$. This bipartite graph G is called the *incidence graph* of the set system H, and the bipartite adjacency matrix of G the *incidence matrix* of H; these are simply alternative ways of representing a set system. Incidence graphs of geometric configurations often give rise to interesting bipartite graphs; in this context, the incidence graph is sometimes called the *Levi graph* of the configuration. The incidence graph of the Fano plane is shown in Figure 1.16. This graph is known as the *Heawood graph*.

Intersection Graphs

With each set system $(V, \mathcal{F})$ one may associate its *intersection graph*. This is the graph whose vertex set is $\mathcal{F}$, two sets in $\mathcal{F}$ being adjacent if their intersection is nonempty. For instance, when V is the vertex set of a simple graph G and $\mathcal{F} := E$,

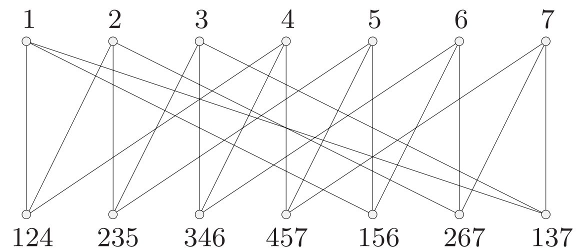

Fig. 1.16. The incidence graph of the Fano plane: the Heawood graph

the edge set of G, the intersection graph of $(V, \mathcal{F})$ has as vertices the edges of G, two edges being adjacent if they have an end in common. For historical reasons, this graph is known as the *line graph* of G and denoted $L(G)$. Figure 1.17 depicts a graph and its line graph.

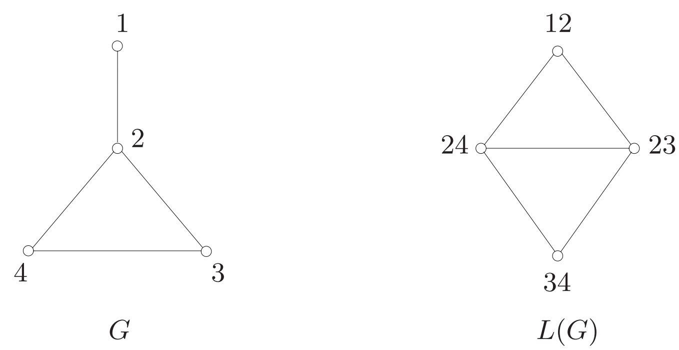

Fig. 1.17. A graph and its line graph

It can be shown that the intersection graph of the Desargues configuration is isomorphic to the line graph of K_5, which in turn is isomorphic to the complement of the Petersen graph (Exercise 1.3.2). As for the Fano plane, its intersection graph is isomorphic to K_7, because any two of its seven lines have a point in common.

The definition of the line graph $L(G)$ may be extended to all loopless graphs G as being the graph with vertex set E in which two vertices are joined by just as many edges as their number of common ends in G.

When $V = \mathbb{R}$ and $\mathcal{F}$ is a set of closed intervals of $\mathbb{R}$, the intersection graph of $(V, \mathcal{F})$ is called an *interval graph*. Examples of practical situations which give rise to interval graphs can be found in the book by Berge (1973). Berge even wrote a detective story whose resolution relies on the theory of interval graphs; see Berge (1995).

It should be evident from the above examples that graphs are implicit in a wide variety of structures. Many such graphs are not only interesting in their own right but also serve to provide insight into the structures from which they arise.

Exercises

1.3.1

a) Show that the graph in Figure 1.18 is isomorphic to the Heawood graph (Figure 1.16).

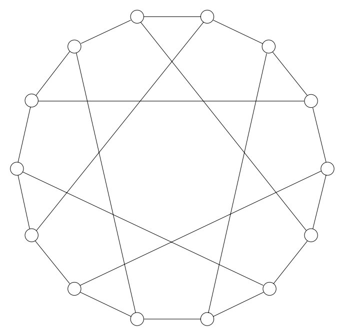

Fig. 1.18. Another drawing of the Heawood graph

b) Deduce that the Heawood graph is vertex-transitive.

1.3.2 Show that the following three graphs are isomorphic:

▷ the intersection graph of the Desargues configuration,
▷ the line graph of K_5,
▷ the complement of the Petersen graph.

1.3.3 Show that the line graph of $K_{3,3}$ is self-complementary.

1.3.4 Show that neither of the graphs displayed in Figure 1.19 is a line graph.

1.3.5 Let $H := (V, \mathcal{F})$ be a hypergraph. The number of edges incident with a vertex v of H is its *degree*, denoted $d(v)$. A *degree sequence* of H is a vector $\mathbf{d} := (d(v) : v \in V)$. Let $\mathbf{M}$ be the incidence matrix of H and $\mathbf{d}$ the corresponding degree sequence of H. Show that the sum of the columns of $\mathbf{M}$ is equal to $\mathbf{d}$.

1.3.6 Let $H := (V, \mathcal{F})$ be a hypergraph. For $v \in V$, let $\mathcal{F}_v$ denote the set of edges of H incident to v. The *dual* of H is the hypergraph H^* whose vertex set is $\mathcal{F}$ and whose edges are the sets $\mathcal{F}_v$, $v \in V$.

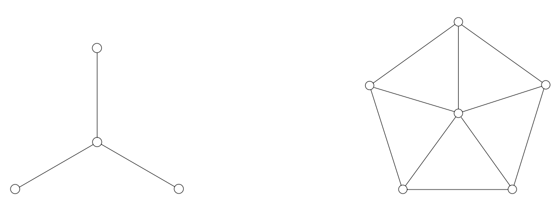

Fig. 1.19. Two graphs that are not line graphs

a) How are the incidence graphs of H and H^* related?
b) Show that the dual of H^* is isomorphic to H.
c) A hypergraph is *self-dual* if it is isomorphic to its dual. Show that the Fano and Desargues hypergraphs are self-dual.

1.3.7 HELLY PROPERTY
A family of sets has the *Helly Property* if the members of each pairwise intersecting subfamily have an element in common.

a) Show that the family of closed intervals on the real line has the Helly Property. (E. HELLY)
b) Deduce that the graph in Figure 1.20 is not an interval graph.

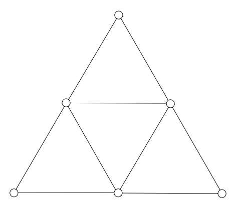

Fig. 1.20. A graph that is not an interval graph

1.3.8 KNESER GRAPH
Let m and n be positive integers, where $n > 2m$. The *Kneser graph* $KG_{m,n}$ is the graph whose vertices are the m-subsets of an n-set S, two such subsets being adjacent if and only if their intersection is empty. Show that:

a) $KG_{1,n} \cong K_n$, $n \geq 3$,
b) $KG_{2,n}$ is isomorphic to the complement of $L(K_n)$, $n \geq 5$.

1.3.9 Let G be a simple graph with incidence matrix $\mathbf{M}$.

a) Show that the adjacency matrix of its line graph $L(G)$ is $\mathbf{M}^t\mathbf{M} - 2\mathbf{I}$, where $\mathbf{I}$ is the $m \times m$ identity matrix.
b) Using the fact that $\mathbf{M}^t\mathbf{M}$ is positive-semidefinite, deduce that:
   i) each eigenvalue of $L(G)$ is at least -2,
   ii) if the rank of $\mathbf{M}$ is less than m, then -2 is an eigenvalue of $L(G)$.

1.3.10

a) Consider the following two matrices $\mathbf{B}$ and $\mathbf{C}$, where x is an indeterminate, $\mathbf{M}$ is an arbitrary $n \times m$ matrix, and $\mathbf{I}$ is an identity matrix of the appropriate dimension.

$$\mathbf{B} := \begin{bmatrix} \mathbf{I} & \mathbf{M} \\ \mathbf{M}^t & x\mathbf{I} \end{bmatrix} \qquad \mathbf{C} := \begin{bmatrix} x\mathbf{I} & -\mathbf{M} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}$$

By equating the determinants of $\mathbf{BC}$ and $\mathbf{CB}$, derive the identity

$$\det(x\mathbf{I} - \mathbf{M}^t\mathbf{M}) = x^{m-n}\det(x\mathbf{I} - \mathbf{M}\mathbf{M}^t)$$

b) Let G be a simple k-regular graph with $k \geq 2$. By appealing to Exercise 1.3.9 and using the above identity, establish the following relationship between the characteristic polynomials of $L(G)$ and G.

$$\det(\mathbf{A}_{L(G)} - x\mathbf{I}) = (-1)^{m-n}(x+2)^{m-n}\det(\mathbf{A}_G - (x+2-k)\mathbf{I})$$

c) Deduce that:
   i) to each eigenvalue $\lambda \neq -k$ of G, there corresponds an eigenvalue $\lambda + k - 2$ of $L(G)$, with the same multiplicity,
   ii) -2 is an eigenvalue of $L(G)$ with multiplicity $m - n + r$, where r is the multiplicity of the eigenvalue $-k$ of G. (If $-k$ is not an eigenvalue of G then $r = 0$.) (H. Sachs)

1.3.11

a) Using Exercises 1.1.22 and 1.3.10, show that the eigenvalues of $L(K_5)$ are

$$(6, 1, 1, 1, 1, -2, -2, -2, -2, -2)$$

b) Applying Exercise 1.1.23, deduce that the Petersen graph has eigenvalues

$$(3, 1, 1, 1, 1, 1, -2, -2, -2, -2)$$

1.3.12 Sperner's Lemma

Let T be a triangle in the plane. A subdivision of T into triangles is *simplicial* if any two of the triangles which intersect have either a vertex or an edge in common. Consider an arbitrary simplicial subdivision of T into triangles. Assign the colours red, blue, and green to the vertices of these triangles in such a way that each colour is missing from one side of T but appears on the other two sides. (Thus, in particular, the vertices of T are assigned the colours red, blue, and green in some order.)

a) Show that the number of triangles in the subdivision whose vertices receive all three colours is odd. (E. Sperner)
b) Deduce that there is always at least one such triangle.

(Sperner's Lemma, generalized to n-dimensional simplices, is the key ingredient in a proof of Brouwer's Fixed Point Theorem: *every continuous mapping of a closed n-disc to itself has a fixed point*; see Bondy and Murty (1976).)

1.3.13 Finite Projective Plane
A *finite projective plane* is a geometric configuration $(P, \mathcal{L})$ in which:

i) any two points lie on exactly one line,
ii) any two lines meet in exactly one point,

iii) there are four points no three of which lie on a line.

(Condition (iii) serves only to exclude two trivial configurations — the *pencil*, in which all points are collinear, and the *near-pencil*, in which all but one of the points are collinear.)

a) Let $(P, \mathcal{L})$ be a finite projective plane. Show that there is an integer $n \geq 2$ such that $|P| = |\mathcal{L}| = n^2 + n + 1$, each point lies on $n + 1$ lines, and each line contains $n + 1$ points (the instance $n = 2$ being the Fano plane). This integer n is called the *order* of the projective plane.
b) How many vertices has the incidence graph of a finite projective plane of order n, and what are their degrees?

1.3.14 Consider the nonzero vectors in $\mathbb{F}^3$, where $\mathbb{F} = GF(q)$ and q is a prime power. Define two of these vectors to be *equivalent* if one is a multiple of the other. One can form a finite projective plane $(P, \mathcal{L})$ of order q by taking as points and lines the $(q^3 - 1)/(q - 1) = q^2 + q + 1$ equivalence classes defined by this equivalence relation and defining a point (a, b, c) and line (x, y, z) to be incident if $ax + by + cz = 0$ (in $GF(q)$). This plane is denoted $PG_{2,q}$.

a) Show that $PG_{2,2}$ is isomorphic to the Fano plane.
b) Construct $PG_{2,3}$.

1.3.15 The de Bruijn–Erdős Theorem

a) Let $G[X, Y]$ be a bipartite graph, each vertex of which is joined to at least one, but not all, vertices in the other part. Suppose that $d(x) \geq d(y)$ for all $xy \notin E$. Show that $|Y| \geq |X|$, with equality if and only if $d(x) = d(y)$ for all $xy \notin E$ with $x \in X$ and $y \in Y$.
b) Deduce the following theorem.
Let $(P, \mathcal{L})$ be a geometric configuration in which any two points lie on exactly one line and not all points lie on a single line. Then $|\mathcal{L}| \geq |P|$. Furthermore, if $|\mathcal{L}| = |P|$, then $(P, \mathcal{L})$ is either a finite projective plane or a near-pencil.
(N.G. de Bruijn and P. Erdős)

1.3.16 Show that:

a) the line graphs $L(K_n)$, $n \geq 4$, and $L(K_{n,n})$, $n \geq 2$, are strongly regular,
b) the *Shrikhande graph*, displayed in Figure 1.21 (where vertices with the same label are to be identified), is strongly regular, with the same parameters as those of $L(K_{4,4})$, but is not isomorphic to $L(K_{4,4})$.

1.3.17

a) Show that:
 i) $\mathrm{Aut}(L(K_n)) \not\cong \mathrm{Aut}(K_n)$ for $n = 2$ and $n = 4$,
 ii) $\mathrm{Aut}(L(K_n)) \cong \mathrm{Aut}(K_n)$ for $n = 3$ and $n \geq 5$.
b) Appealing to Exercises 1.2.11 and 1.3.2, deduce that the automorphism group of the Petersen graph is isomorphic to the symmetric group S_5.

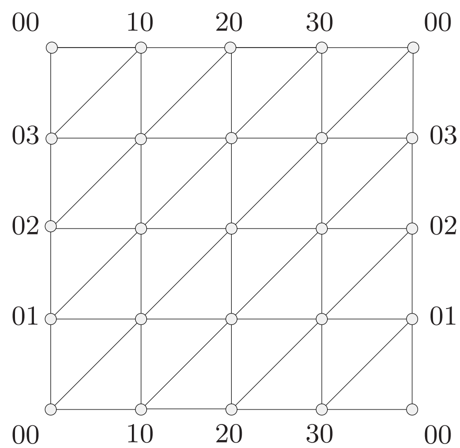

Fig. 1.21. An embedding of the Shrikhande graph on the torus

1.3.18 Cayley Graph
Let Γ be a group, and let S be a set of elements of Γ not including the identity element. Suppose, furthermore, that the inverse of every element of S also belongs to S. The *Cayley graph* of Γ with respect to S is the graph $\mathrm{CG}(\Gamma, S)$ with vertex set Γ in which two vertices x and y are adjacent if and only if $xy^{-1} \in S$. (Note that, because S is closed under taking inverses, if $xy^{-1} \in S$, then $yx^{-1} \in S$.)

a) Show that the n-cube is a Cayley graph.
b) Let G be a Cayley graph $\mathrm{CG}(\Gamma, S)$ and let x be an element of Γ.
   i) Show that the mapping α_x defined by the rule that $\alpha_x(y) := xy$ is an automorphism of G.
   ii) Deduce that every Cayley graph is vertex-transitive.
c) By considering the Petersen graph, show that not every vertex-transitive graph is a Cayley graph.

1.3.19 Circulant
A *circulant* is a Cayley graph $\mathrm{CG}(\mathbb{Z}_n, S)$, where $\mathbb{Z}_n$ is the additive group of integers modulo n. Let p be a prime, and let i and j be two nonzero elements of $\mathbb{Z}_p$.

a) Show that $\mathrm{CG}(\mathbb{Z}_p, \{i, -i\}) \cong \mathrm{CG}(\mathbb{Z}_p, \{j, -j\})$.
b) Determine when $\mathrm{CG}(\mathbb{Z}_p, \{1, -1, i, -i\}) \cong \mathrm{CG}(\mathbb{Z}_p, \{1, -1, j, -j\})$.

1.3.20 Paley Graph
Let q be a prime power, $q \equiv 1 \ (mod\ 4)$. The *Paley graph* PG_q is the graph whose vertex set is the set of elements of the field $GF(q)$, two vertices being adjacent if their difference is a nonzero square in $GF(q)$.

a) Draw PG_5, PG_9, and PG_{13}.
b) Show that these three graphs are self-complementary.
c) Let a be a nonsquare in $GF(q)$. By considering the mapping $\theta : GF(q) \to GF(q)$ defined by $\theta(x) := ax$, show that PG_q is self-complementary for all q.

1.4 Constructing Graphs from Other Graphs

We have already seen a couple of ways in which we may associate with each graph another graph: the complement (in the case of simple graphs) and the line graph. If we start with two graphs G and H rather than just one, a new graph may be defined in several ways. For notational simplicity, we assume that G and H are simple, so that each edge is an unordered pair of vertices; the concepts described here can be extended without difficulty to the general context.

Union and Intersection

Two graphs are *disjoint* if they have no vertex in common, and *edge-disjoint* if they have no edge in common. The most basic ways of combining graphs are by union and intersection. The *union* of simple graphs G and H is the graph $G \cup H$ with vertex set $V(G) \cup V(H)$ and edge set $E(G) \cup E(H)$. If G and H are disjoint, we refer to their union as a *disjoint union*, and generally denote it by $G + H$. These operations are associative and commutative, and may be extended to an arbitrary number of graphs. It can be seen that a graph is disconnected if and only if it is a disjoint union of two (nonnull) graphs. More generally, every graph G may be expressed uniquely (up to order) as a disjoint union of connected graphs (Exercise 1.4.1). These graphs are called the *connected components*, or simply the *components*, of G. The number of components of G is denoted $c(G)$. (The null graph has the anomalous property of being the only graph without components.)

The *intersection* $G \cap H$ of G and H is defined analogously. (Note that when G and H are disjoint, their intersection is the null graph.) Figure 1.22 illustrates these concepts. The graph $G \cup H$ shown in Figure 1.22 has just one component, whereas the graph $G \cap H$ has two components.

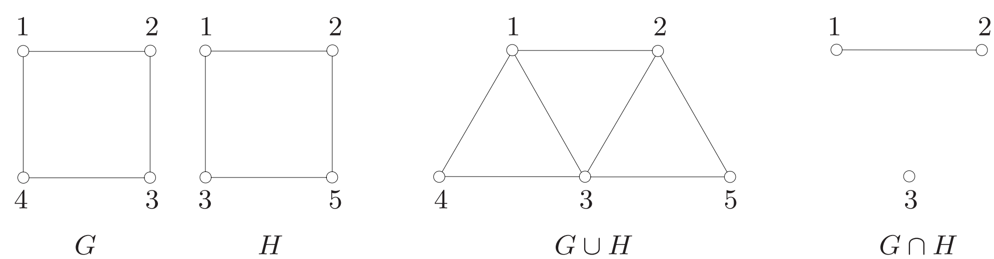

Fig. 1.22. The union and intersection of two graphs

Cartesian Product

There are also several ways of forming from two graphs a new graph whose vertex set is the cartesian product of their vertex sets. These constructions are consequently referred to as 'products'. We now describe one of them.

The *cartesian product* of simple graphs G and H is the graph $G \,\square\, H$ whose vertex set is $V(G) \times V(H)$ and whose edge set is the set of all pairs $(u_1, v_1)(u_2, v_2)$ such that either $u_1u_2 \in E(G)$ and $v_1 = v_2$, or $v_1v_2 \in E(H)$ and $u_1 = u_2$. Thus, for each edge u_1u_2 of G and each edge v_1v_2 of H, there are four edges in $G \,\square\, H$, namely $(u_1, v_1)(u_2, v_1)$, $(u_1, v_2)(u_2, v_2)$, $(u_1, v_1)(u_1, v_2)$, and $(u_2, v_1)(u_2, v_2)$ (see Figure 1.23a); the notation used for the cartesian product reflects this fact. More generally, the cartesian product $P_m \,\square\, P_n$ of two paths is the $(m \times n)$-*grid*. An example is shown in Figure 1.23b.

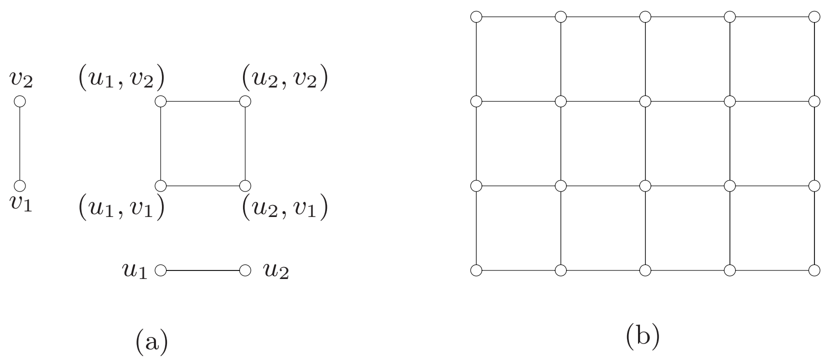

Fig. 1.23. (a) The cartesian product $K_2 \,\square\, K_2$, and (b) the (5×4)-grid

For $n \geq 3$, the cartesian product $C_n \,\square\, K_2$ is a polyhedral graph, the n-*prism*; the 3-prism, 4-prism, and 5-prism are commonly called the *triangular prism*, the *cube*, and the *pentagonal prism* (see Figure 1.24). The cartesian product is arguably the most basic of graph products. There exist a number of others, each arising naturally in various contexts. We encounter several of these in later chapters.

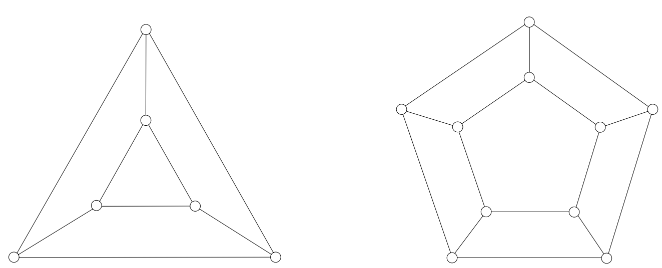

Fig. 1.24. The triangular and pentagonal prisms

Exercises

1.4.1 Show that every graph may be expressed uniquely (up to order) as a disjoint union of connected graphs.

1.4.2 Show that the rank over $GF(2)$ of the incidence matrix of a graph G is $n-c$.

1.4.3 Show that the cartesian product is both associative and commutative.

1.4.4 Find an embedding of the cartesian product $C_m \square C_n$ on the torus.

1.4.5

a) Show that the cartesian product of two vertex-transitive graphs is vertex-transitive.
b) Give an example to show that the cartesian product of two edge-transitive graphs need not be edge-transitive.

1.4.6

a) Let G be a self-complementary graph and let P be a path of length three disjoint from G. Form a new graph H from $G \cup P$ by joining the first and third vertices of P to each vertex of G. Show that H is self-complementary.
b) Deduce (by appealing to Exercise 1.2.16) that there exists a self-complementary graph on n vertices if and only if $n \equiv 0, 1 \pmod 4$.

1.5 Directed Graphs

Although many problems lend themselves to graph-theoretic formulation, the concept of a graph is sometimes not quite adequate. When dealing with problems of traffic flow, for example, it is necessary to know which roads in the network are one-way, and in which direction traffic is permitted. Clearly, a graph of the network is not of much use in such a situation. What we need is a graph in which each link has an assigned orientation, namely a directed graph.

Formally, a *directed graph* D is an ordered pair $(V(D), A(D))$ consisting of a set $V := V(D)$ of *vertices* and a set $A := A(D)$, disjoint from $V(D)$, of *arcs*, together with an *incidence function* ψ_D that associates with each arc of D an ordered pair of (not necessarily distinct) vertices of D. If a is an arc and $\psi_D(a) = (u, v)$, then a is said to *join* u *to* v; we also say that u *dominates* v. The vertex u is the *tail* of a, and the vertex v its *head*; they are the two *ends* of a. Occasionally, the orientation of an arc is irrelevant to the discussion. In such instances, we refer to the arc as an *edge* of the directed graph. The number of arcs in D is denoted by $a(D)$. The vertices which dominate a vertex v are its *in-neighbours*, those which are dominated by the vertex its *outneighbours*. These sets are denoted by $N_D^-(v)$ and $N_D^+(v)$, respectively.

For convenience, we abbreviate the term 'directed graph' to *digraph*. A *strict* digraph is one with no loops or parallel arcs (arcs with the same head and the same tail).

With any digraph D, we can associate a graph G on the same vertex set simply by replacing each arc by an edge with the same ends. This graph is the *underlying graph* of D, denoted $G(D)$. Conversely, any graph G can be regarded as a digraph, by replacing each of its edges by two oppositely oriented arcs with the same ends; this digraph is the *associated digraph* of G, denoted $D(G)$. One may also obtain a digraph from a graph G by replacing each edge by just one of the two possible arcs with the same ends. Such a digraph is called an *orientation* of G. We occasionally use the symbol $\overrightarrow{G}$ to specify an orientation of G (even though a graph generally has many orientations). An orientation of a simple graph is referred to as an *oriented graph*. One particularly interesting instance is an orientation of a complete graph. Such an oriented graph is called a *tournament*, because it can be viewed as representing the results of a round-robin tournament, one in which each team plays every other team (and there are no ties).

Digraphs, like graphs, have a simple pictorial representation. A digraph is represented by a diagram of its underlying graph together with arrows on its edges, each arrow pointing towards the head of the corresponding arc. The four unlabelled tournaments on four vertices are shown in Figure 1.25 (see Exercise 1.5.3a).

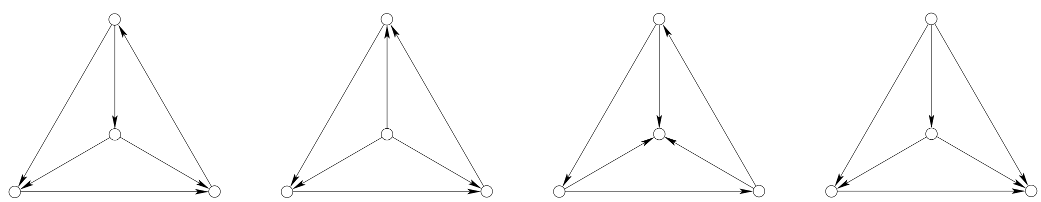

Fig. 1.25. The four unlabelled tournaments on four vertices

Every concept that is valid for graphs automatically applies to digraphs too. For example, the *degree* of a vertex v in a digraph D is simply the degree of v in $G(D)$, the underlying graph of D.[1] Likewise, a digraph is said to be *connected* if its underlying graph is connected.[2] But there are concepts in which orientations play an essential role. For instance, the *indegree* $d_D^-(v)$ of a vertex v in D is the number of arcs with head v, and the *outdegree* $d_D^+(v)$ of v is the number of arcs with tail v. The minimum indegree and outdegree of D are denoted by $\delta^-(D)$ and $\delta^+(D)$, respectively; likewise, the maximum indegree and outdegree of D are

[1] *In such cases, we employ the same notation as for graphs (with G replaced by D). Thus the degree of v in D is denoted by $d_D(v)$. These instances of identical notation are recorded only once in the glossaries, namely for graphs.*

[2] *The index includes only those definitions for digraphs which differ substantively from their analogues for graphs. Thus the term 'connected digraph' does not appear there, only 'connected graph'.*

denoted $\Delta^-(D)$ and $\Delta^+(D)$, respectively. A digraph is *k-diregular* if each indegree and each outdegree is equal to k. A vertex of indegree zero is called a *source*, one of outdegree zero a *sink*. A *directed path* or *directed cycle* is an orientation of a path or cycle in which each vertex dominates its successor in the sequence. There is also a notion of connectedness in digraphs which takes directions into account, as we shall see in Chapter 2.

Two special digraphs are shown in Figure 1.26. The first of these is a 2-diregular digraph, the second a 3-diregular digraph (see Bondy (1978)); we adopt here the convention of representing two oppositely oriented arcs by an edge. These digraphs can both be constructed from the Fano plane (Exercise 1.5.9). They also possess other unusual properties, to be described in Chapter 2.

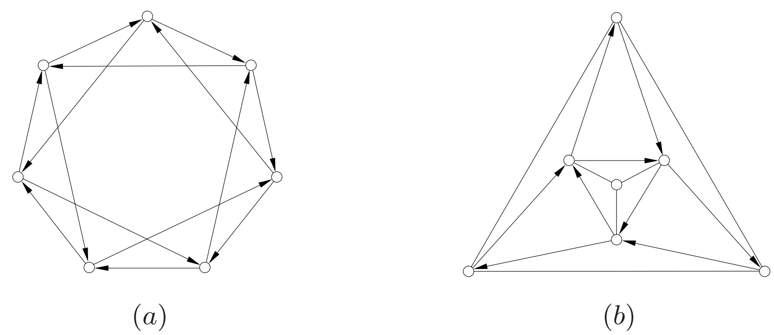

Fig. 1.26. (a) the Koh–Tindell digraph, and (b) a directed analogue of the Petersen graph

Further examples of interesting digraphs can be derived from other mathematical structures, such as groups. For example, there is a natural directed analogue of a Cayley graph. If Γ is a group, and S a subset of Γ not including the identity element, the *Cayley digraph* of Γ with respect to S is the digraph, denoted $\mathrm{CD}(\Gamma, S)$, whose vertex set is Γ and in which vertex x dominates vertex y if and only if $xy^{-1} \in S$. A *directed circulant* is a Cayley digraph $\mathrm{CD}(\mathbb{Z}_n, S)$, where $\mathbb{Z}_n$ is the group of integers modulo n. The *Koh–Tindell digraph* of Figure 1.26a is a directed circulant based on $\mathbb{Z}_7$.

With each digraph D, one may associate another digraph, $\overleftarrow{D}$, obtained by reversing each arc of D. The digraph $\overleftarrow{D}$ is called the *converse* of D. Because the converse of the converse is just the original digraph, the converse of a digraph can be thought of as its 'directional dual'. This point of view gives rise to a simple yet useful principle.

Principle of Directional Duality

Any statement about a digraph has an accompanying 'dual' statement, obtained by applying the statement to the converse of the digraph and reinterpreting it in terms of the original digraph.

For instance, the sum of the indegrees of the vertices of a digraph is equal to the total number of arcs (Exercise 1.5.2). Applying the Principle of Directional Duality, we immediately deduce that the sum of the outdegrees is also equal to the number of arcs.

Apart from the practical aspect mentioned earlier, assigning suitable orientations to the edges of a graph is a convenient way of exploring properties of the graph, as we shall see in Chapter 6.

Exercises

1.5.1 How many orientations are there of a labelled graph G?

$\star$**1.5.2** Let D be a digraph.

a) Show that $\sum_{v \in V} d^-(v) = m$.
b) Using the Principle of Directional Duality, deduce that $\sum_{v \in V} d^+(v) = m$.

1.5.3 Two digraphs D and D' are *isomorphic*, written $D \cong D'$, if there are bijections $\theta : V(D) \to V(D')$ and $\phi : A(D) \to A(D')$ such that $\psi_D(a) = (u, v)$ if and only if $\psi_{D'}(\phi(a)) = (\theta(u), \theta(v))$. Such a pair of mappings is called an *isomorphism* between D and D'.

a) Show that the four tournaments in Figure 1.25 are pairwise nonisomorphic, and that these are the only ones on four vertices, up to isomorphism.
b) How many tournaments are there on five vertices, up to isomorphism?

1.5.4

a) Define the notions of vertex-transitivity and arc-transitivity for digraphs.
b) Show that:
   i) every vertex-transitive digraph is diregular,
   ii) the Koh–Tindell digraph (Figure 1.26a) is vertex-transitive but not arc-transitive.

1.5.5 A digraph is *self-converse* if it is isomorphic to its converse. Show that both digraphs in Figure 1.26 are self-converse.

1.5.6 Incidence Matrix of a Digraph
Let D be a digraph with vertex set V and arc set A. The *incidence matrix* of D (with respect to given orderings of its vertices and arcs) is the $n \times m$ matrix $\mathbf{M}_D := (m_{va})$, where

$$m_{va} = \begin{cases} 1 & \text{if arc } a \text{ is a link and vertex } v \text{ is the tail of } a \\ -1 & \text{if arc } a \text{ is a link and vertex } v \text{ is the head of } a \\ 0 & \text{otherwise} \end{cases}$$

Let $\mathbf{M}$ be the incidence matrix of a connected digraph D. Show that the rank of $\mathbf{M}$ is $n - 1$.

⋆**1.5.7** Totally Unimodular Matrix
A matrix is *totally unimodular* if each of its square submatrices has determinant equal to 0, +1, or −1. Let $\mathbf{M}$ be the incidence matrix of a digraph.

a) Show that $\mathbf{M}$ is totally unimodular. (H. Poincaré)
b) Deduce that the matrix equation $\mathbf{Mx} = \mathbf{b}$ has a solution in integers provided that it is consistent and the vector $\mathbf{b}$ is integral.

——— ≀≀ ———

1.5.8 Balanced Digraph
A digraph D is *balanced* if $|d^+(v) - d^-(v)| \leq 1$, for all $v \in V$. Show that every graph has a balanced orientation.

1.5.9 Describe how the two digraphs in Figure 1.26 can be constructed from the Fano plane.

1.5.10 Paley Tournament
Let q be a prime power, $q \equiv 3 \pmod 4$. The *Paley tournament* PT_q is the tournament whose vertex set is the set of elements of the field $GF(q)$, vertex i dominating vertex j if and only if $j - i$ is a nonzero square in $GF(q)$.

a) Draw PT_3, PT_7, and PT_{11}.
b) Show that these three digraphs are self-converse.

1.5.11 Stockmeyer Tournament
For a nonzero integer k, let $\mathrm{pow}\,(k)$ denote the greatest integer p such that 2^p divides k, and set $\mathrm{odd}\,(k) := k/2^p$. (For example, $\mathrm{pow}\,(12) = 2$ and $\mathrm{odd}\,(12) = 3$, whereas $\mathrm{pow}\,(-1) = 0$ and $\mathrm{odd}\,(-1) = -1$.) The *Stockmeyer tournament* ST_n, where $n \geq 1$, is the tournament whose vertex set is $\{1, 2, 3, \ldots, 2^n\}$ in which vertex i dominates vertex j if $\mathrm{odd}\,(j - i) \equiv 1 \pmod 4$.

a) Draw ST_2 and ST_3.
b) Show that ST_n is both self-converse and asymmetric (that is, has no nontrivial automorphisms). (P.K. Stockmeyer)

1.5.12 Arc-transitive Graph
An undirected graph G is *arc-transitive* if its associated digraph $D(G)$ is arc-transitive. (Equivalently, G is arc-transitive if, given any two ordered pairs (x, y) and (u, v) of adjacent vertices, there exists an automorphism of G which maps (x, y) to (u, v).)

a) Show that any graph which is arc-transitive is both vertex-transitive and edge-transitive.
b) Let G be a k-regular graph which is both vertex-transitive and edge-transitive, but not arc-transitive. Show that k is even. (An example of such a graph with $k = 4$ may be found in Godsil and Royle (2001).)

1.5.13 Adjacency Matrix of a Digraph
The *adjacency matrix* of a digraph D is the $n \times n$ matrix $\mathbf{A}_D = (a_{uv})$, where a_{uv} is the number of arcs in D with tail u and head v. Let $\mathbf{A}$ be the adjacency matrix of a tournament on n vertices. Show that rank $\mathbf{A} = n - 1$ if n is odd and rank $\mathbf{A} = n$ if n is even.

1.6 Infinite Graphs

As already mentioned, the graphs studied in this book are assumed to be finite. There is, however, an extensive theory of graphs defined on infinite sets of vertices and/or edges. Such graphs are known as *infinite graphs*. An infinite graph is *countable* if both its vertex and edge sets are countable. Figure 1.27 depicts three well-known countable graphs, the *square lattice*, the *triangular lattice*, and the *hexagonal lattice*.

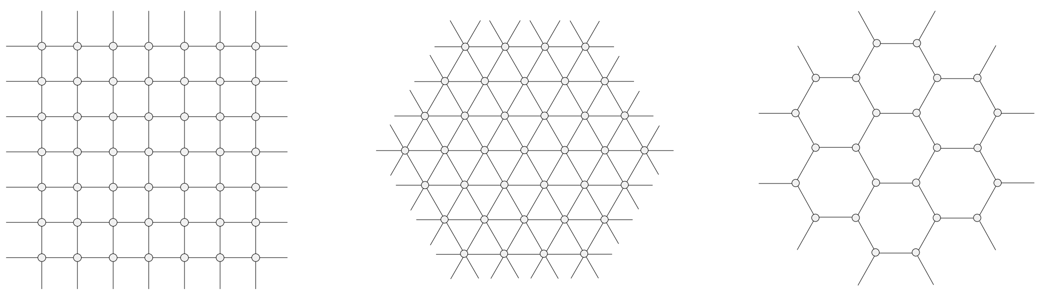

Fig. 1.27. The square, triangular and hexagonal lattices

Most notions that are valid for finite graphs are either directly applicable to infinite graphs or else require some simple modification. Whereas the definition of the degree of a vertex is essentially the same as for finite graphs (with 'number' replaced by 'cardinality'), there are two types of infinite path, one having an initial but no terminal vertex (called a *one-way infinite path*), and one having neither initial nor terminal vertices (called a *two-way infinite path*); the square lattice is the cartesian product of two two-way infinite paths. However, certain concepts for finite graphs have no natural 'infinite' analogue, the cycle for instance (although, in some circumstances, a two-way infinite path may be regarded as an infinite cycle).

While the focus of this book is on finite graphs, we include occasional remarks and exercises on infinite graphs, mainly to illustrate the differences between finite and infinite graphs. Readers interested in pursuing the topic are referred to the survey article by Thomassen (1983a) or the book by Diestel (2005), which includes a chapter on infinite graphs.

Exercises

1.6.1 Locally Finite Graph
An infinite graph is *locally finite* if every vertex is of finite degree. Give an example of a locally finite graph in which no two vertices have the same degree.

1.6.2 For each positive integer d, describe a simple infinite planar graph with minimum degree d. (We shall see, in Chapter 10, that every simple finite planar graph has a vertex of degree at most five.)

1.6.3 Give an example of a self-complementary infinite graph.

1.6.4 Unit Distance Graph
The *unit distance graph* on a subset V of $\mathbb{R}^2$ is the graph with vertex set V in which two vertices (x_1, y_1) and (x_2, y_2) are adjacent if their euclidean distance is equal to 1, that is, if $(x_1 - x_2)^2 + (y_1 - y_2)^2 = 1$. When $V = \mathbb{Q}^2$, this graph is called the *rational* unit distance graph, and when $V = \mathbb{R}^2$, the *real* unit distance graph. (Note that these are both infinite graphs.)

a) Let V be a finite subset of the vertex set of the infinite 2-dimensional integer lattice (see Figure 1.27), and let d be an odd positive integer. Denote by G the graph with vertex set V in which two vertices (x_1, y_1) and (x_2, y_2) are adjacent if their euclidean distance is equal to d. Show that G is bipartite.
b) Deduce that the rational unit distance graph is bipartite.
c) Show, on the other hand, that the real unit distance graph is not bipartite.

1.7 Related Reading

History of Graph Theory

An attractive account of the history of graph theory up to 1936, complete with annotated extracts from pivotal papers, can be found in Biggs et al. (1986). The first book on graph theory was published by König (1936). It led to the development of a strong school of graph theorists in Hungary which included P. Erdős and T. Gallai. Also in the thirties, H. Whitney published a series of influential articles (see Whitney (1992)).

As with every branch of mathematics, graph theory is best learnt by doing. The book *Combinatorial Problems and Exercises* by Lovász (1993) is highly recommended as a source of stimulating problems and proof techniques. A general guide to solving problems in mathematics is the very readable classic *How to Solve It* by Pólya (2004). The delightful *Proofs from the Book* by Aigner and Ziegler (2004) is a compilation of beautiful proofs in mathematics, many of which treat combinatorial questions.