Chapter 1. Graphs

Section 1.1. Graphs and Their Representations

Note. In this section, we lay down some basic definitions and notation. We largely follow a definition/theorem/proof approach. Bondy and Murty present definitions in the "conversational" parts of the text.

Definition. A *graph* G is an ordered pair $(V(G), E(G))$ where $V(G)$ is a set whose elements are called *vertices* ($V(G)$ is called the *vertex set*) and $E(G)$ is a set disjoint from $V(G)$ whose elements are called *edges* ($E(G)$ is called the *edge set*), together with an *incidence function* ψ_G that associates with each edge of G an unordered pair of (not necessarily different) vertices u and v, denoted $\{u, v\} = \{v, u\}$. If $e \in E(G)$ and $\psi_G(e) = \{u, v\}$ then edge e is said to *join* u and v, and the vertices u and v are called the *ends* of e. The *order* of graph G, denoted $v(G)$, is the cardinality of the vertex set $|V(G)|$ (most often, taken to be finite). The *size* of graph G, denoted $e(G)$, is the cardinality of the edge set $|E(G)|$ (also usually taken to be finite).

Note. Unless stated otherwise, **we consider graphs to be of finite size and finite order!** We adopt the standard notation $n = v(G)$ and $m = e(G)$. We usually denote edge $\{u, v\}$ simply as uv. Notice that an edge where both ends are vertex v is denoted $\{v, v\} = vv$, even though in a set elements are not repeated (so "$\{u, v\}$" does not technically denote, in general, the set consisting of elements u and v unless u and v are distinct).

Example 1.1.1. Let $G = (V(G), E(G))$ where $V(G) = \{u, v, w, x, y\}$, $E(G) = \{a, b, c, d, e, f, g, h\}$, and ψ_G is defined by

$$\psi_G(a) = uv, \quad \psi_G(b) = uu, \quad \psi_G(c) = vw, \quad \psi_G(d) = wx,$$
$$\psi_G(e) = vx, \quad \psi_G(f) = wx, \quad \psi_G(g) = ux, \quad \psi_G(h) = xy.$$

Then the order of G is $v(G) = |V(G)| = 5$ and the size of G is $E(G) = |E(G)| = 8$.

Note. We graphically represent a graph (Bondy and Murty state that this is why graphs are called "graphs"—see page 2) by representing vertices as little circles or disks and representing edges by joining the ends of the edge with some curve. Figure 1.1(a) gives such a graphical representation (or "drawing") of the graph of Example 1.1.1.

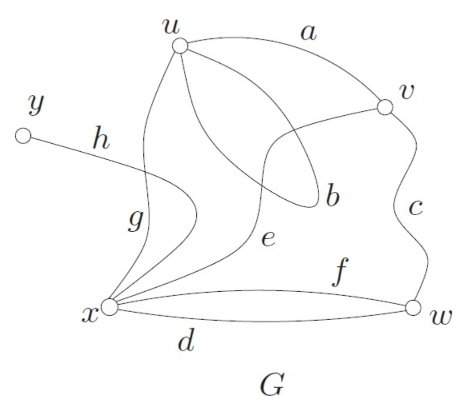

Figure 1.1(a)

Since the vertices have no particular location and the curves representing the edges have no particular shape, then there are multiple ways to represent a graph with a drawing (we do follow the convention that no curve representing an edge contains any vertices other than its ends, and no edge intersects itself except that a loop intersects itself at its ends).

Definition. In graph G with $e \in E(G)$ where $\psi_G(e) = \{u, v\}$ we say that vertices u and v are *incident* with edge e and edge e is *incident* with vertices u and v. Two vertices which are incident with a common edge are *adjacent* and two adjacent vertices are *neighbors*. The set of neighbors of vertex v in graph G is denoted $N_G(v)$. Also, two edges which are incident with a common vertex are *adjacent*. An edge with identical ends is called a *loop* and an edge with distinct ends is called a *link*. Two or more links with the same pair of ends are *parallel edges*. A graph is *finite* if both its vertex set and edge set are finite. A graph is *simple* if it has no loops and no parallel edges.

Note. The graph in Example 1.1.1 and Figure 1.1(a) has edge b as a loop, and edges d and f are parallel edges. So this graph is not simple! We are primarily concerned with finite simple graphs.

Definition. The graph with no vertices and no edges is the *null graph*. A graph with one vertex is a *trivial graph*. Graphs other than the null graph and the trivial graphs are *nontrivial graphs*.

Note 1.1.A. We'll often abbreviate a simple graph as (V, E) where V is the vertex set and E is a set of two-element subsets of V. In this way we can gloss over the incidence function ψ. In drawings, we'll label the vertices and not the edges but will refer to an edge with ends u and v simply as uv.

Definition. A *complete graph* is a simple graph in which every pair of distinct vertices are adjacent. An *empty graph* is a graph with no edges. 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; the pair of sets (X, Y) is the *bipartition* of the bipartite graph and X and Y are its *parts*. A bipartite graph G with bipartition (X, Y) is denoted $G[X, Y]$. If $G[X, Y]$ is simple and every vertex in X is joined to every vertex in Y, then G is a *complete bipartite graph*. A *star* is a complete bipartite graph $G[X, Y]$ with $|X| = 1$ or $|Y| = 1$.

Definition. 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.

Note. In Theorem 4.7, it is shown that a graph is bipartite if and only if it contains no odd cycle.

Note. Drawings of a complete graph, a complete bipartite graph, and a star are given in Figure 1.2, parts (a), (b), and (c), respectively.

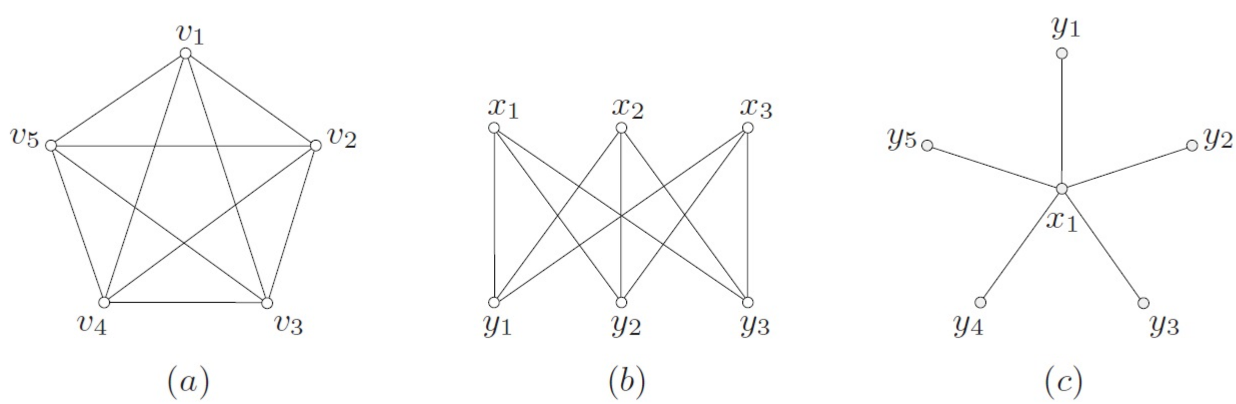

Figure 1.2

Definition. A *path* is a simple graph whose vertices can be arranged in a linear sequence, say $v_1, v_2, \ldots, v_n$, where v_i and v_j are adjacent if and only if either $i = j+1$ or $j = i+1$ (this means that the vertices and edges of a path are distinct). A *cycle on three or more vertices* is a simple graph whose vertices can be arranged as $v_0, v_1, \ldots, v_{n-1}$, where v_i and v_j are adjacent if and only if either $i \equiv j+1 \pmod{n}$ or $j \equiv i+1 \pmod{n}$; 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 edges, and a path or cycle of length k is called a *k-path* or a *k-cycle*, respectively. The parity (*even* or *odd*) of a k-path or k-cycle is the parity of k.

Note. In keeping with geometric intuition, a 3-cycle is called a triangle, a 4-cycle a square, a 5-cycle a pentagon, and so forth.

Note 1.1.B. Recall that, in the setting of $\mathbb{R}$, a set A of real numbers has a *separation* if there are open sets U and V with $U \cap V = \varnothing$, $U \cap A \neq \varnothing$, $V \cap A \neq \varnothing$, and $(U \cap A) \cup (V \cap A) = A$. If there is no separation of A then set A is *connected* (see my online Analysis 1 [MATH 4217/5217] on 3.1. Topology of the Real Numbers). A similar idea holds in topological spaces (see my online notes for Introduction to Topology [MATH 4357/5357] on 3.23. Connected Spaces). We use this same concept in defining a connected graph.

Definition. 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*.

Note. If a graph is disconnected because there exists a partition of the vertex set into sets X and Y for which there is no edge with one end in set X and the other end in a set Y, then the sets X and Y form a "separation" of the graph, similar to the topological concept (we should be careful with the use of the term "separation" here, since in Section 5.2, "Separations and Blocks," we will define a separation of a connected graph, which is a very different concept). Connectivity (and the "degree" of connectivity) is the topic of Chapter 9. In Exercise 3.1.4 it is to be shown that a graph G is connected if and only if for any two nonempty subsets X and Y of $V(G)$, there is a path $v_1, v_2, \ldots, v_n$ in G such that $v_1 \in X$ and $v_n \in Y$. In particular, for any two vertices of a connected graph, there is a path in the graph between them.

Note. We have not formally defined a "drawing" of a graph, making it impossible to formally define a planar graph. So we intuitively define a *planar graph* as a graph that can be drawn in the Cartesian plane in such a way that edges meet only at points corresponding to their common ends. Such a drawing is called a *planar embedding* of the graph. For example, the graph in Figure 1.1(a) is planar—we can "move" loop b and edge h so that they do not intersect any other edges, thus producing a planar embedding of that graph. The graphs in Figure 1.2(a) and 1.2(b) are not planar (as shown in Theorem 10.2 and Exercise 10.1.1(b), and again in Corollaries 10.23 and 10.24; see also Kuratowski's Theorem, Theorem 10.30). Graphs can also be embedded on surfaces other then the Cartesian plane. This idea is explained some in Chapter 2 and in more detail in Chapter 10.

Note. A drawing of a graph is an good way to visualize the graph (if it isn't too big). Another way to present a graph is to tabulate the connectivity in an array.

Definition. Let G be a graph with vertex set V and edge set E. The *incidence matrix* of G is the $n \times m$ matrix (where $n = v(G)$ and $m = e(G)$) $\mathbf{M}_G = (m_{ve})$, where m_{ve} is the number of times (0, 1, or 2) that vertex v and edge e are incident. The *adjacency matrix* of G is the $n \times n$ matrix (where $n = v(G)$) $\mathbf{A}_G = (a_{uv})$, where a_{uv} is the number of edges joining vertices u and v, each loop counting as two edges.

Note. The graph of Figure 1.1(a) is given again in Figure 1.5, along with the incidence and adjacency matrix. Notice that the adjacency matrix is symmetric.

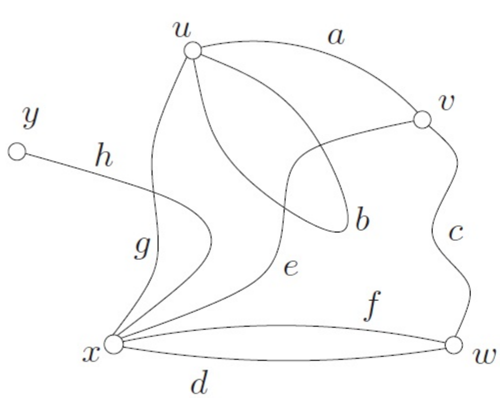

G

	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

$\mathbf{A}$

Figure 1.5

Definition. Let G be a simple finite graph. A list $(N(v) \mid v \in V) = N(v_1), N(v_2), \ldots, N(v_n)$ of the neighbors of the vertices is an *adjacency list* of G. Suppose $G[X, Y]$ is a bipartite graph, where $X = \{x_1, x_2, \ldots, x_r\}$ and $Y = \{y_1, y_2, \ldots, y_s\}$. The *bipartite adjacency matrix* of G is the $r \times s$ matrix $\mathbf{B}_G = (b_{ij})$, where b_{ij} is the number of edges joining x_i and y_j.

Note. We now turn our attention to the "degree" of vertices in a graph and present our first theorems.

Definition. The *degree* of a vertex v in a finite graph G, denoted $d_G(v)$, is the number of edges of G incident with v, each loop counting as two edges. A vertex of degree zero is called an *isolated vertex*. The minimum of the degrees of vertices in graph G is denoted $\delta(G)$, and the maximum of the degrees of vertices in (finite) graph G is denoted $\Delta(G)$. The *average of the degrees* is $d(G) = \frac{1}{n}\sum_{v\in V} d(v)$ where $n = |V| < \infty$.

Note. The earliest known paper on graph theory is by Leonard Euler: "Solutio problematis ad geometriam situs pertinentis" (The Solution of a Problem Relating to the Geometry of Position), *Commentarii Academiae Scientiarum Imperialis Petropolitanae* **8** (1736), 128-140. A translation of Euler's foundational paper is in N.L. Biggs, E.K. Lloyd, and R.J. Wilson's *Graph Theory: 1736–1936*, 2nd edition (NY: Clarendon Press, 1986) in Chapter 1, "Paths," on pages 3 to 8. In it he considered the "Bridges of Königsberg Problem" (see also Section 3.4 of these online notes) and proved the following two results.

Theorem 1.1. For any graph G, $\sum_{v\in V} d(v) = 2m$ where $m = |E|$.

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

Note. Theorem 1.1 is sometimes called "The Handshaking Lemma" since it implies that in a group of people, an even number of people must have shaken the hands of an odd number of other people (this observation is more explicitly Corollary 1.2).

Definition. 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.

Note. You will characterize k-regular graphs for $k = 0, 1, 2$ in Exercise 1.1.5. The 3-regular graphs are called *cubic graphs*.

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 $xy \in E$.

Revised: 2/20/2023