Chapter I

Graphs and Subgraphs

I.1. DEFINITIONS

In this section we give a formal definition of a graph and introduce some of the basic terminology of graph theory. We give examples of graphs and some sample theorems.

A *graph* G is defined by a set $V(G)$ of elements called *vertices*, a set $E(G)$ of elements called *edges*, and a relation of *incidence*, which associates with each edge either one or two vertices called its *ends*.

The present work is concerned only with *finite* graphs, those in which the sets $V(G)$ and $E(G)$ are both finite. Much interesting work has been done on the other graphs, the *infinite* ones. But even the theory of finite graphs is too big to be adequately covered in one volume. Let us therefore make the rule that from here on the word "graph" is to mean a finite graph unless the contrary is stated explicitly.

The terminology of graph theory is not yet standardized. Some authors prefer to use the terms "point" and "line" rather than "vertex" and "edge." This usage may be found inconvenient in problems involving both graphs and geometrical or topological structures. In some of the older papers we may find "branch" used for "edge," and "node" for "vertex."

An edge is called a *link* or a *loop* according as the number of its ends is two or one. We shall however get into the habit of saying that each edge has two ends, with the explanation that in the case of a loop the two ends are coincident. The two ends of an edge are said to be *joined* by that edge, and to be *adjacent*. Accordingly we say that a vertex is joined to itself, or is adjacent to itself, if and only if it is incident with a loop. Two or more links with the same pair of ends are said to constitute a *multiple join*. A graph without loops or multiple joins is called a *strict* graph.

There are many problems of graph theory in which only strict graphs are of interest. Accordingly some authors restrict the term "graph" to mean what we have called a strict graph. When they have occasion to add loops or multiple joins to their structures they speak of "multigraphs."

Examples of graphs are not difficult to find. For one, the edges and vertices of a convex polyhedron are the edges and vertices, respectively, of a graph G. The ends in G of an edge are its ends in the geometric sense. We call G the graph of the polyhedron.

A roadmap can be interpreted as a graph. The vertices are the junctions, and an edge is the stretch of road from one junction to the next, or from a junction back to itself. Similarly an electrical circuit may give us a graph in which the vertices are terminals and the edges wires.

It is not difficult to see graphs in genealogical tables and computer programs. All through mathematics they are visible to the graph-theoretical eye of faith, for much of mathematics can be described in terms of binary relations, and what is a binary relation but a graph?

It is customary to represent a graph G by a drawing on paper. The vertices are drawn as dots. A link with ends x and y is represented by a straight or curved line joining the dots of x and y, and not meeting any other vertex-dot. A loop with end x is drawn as a curve leaving the dot of x and returning to it again, without meeting any other vertex-dot on the way.

In such a drawing it may happen that two edge-curves intersect at some point away from all the vertex-dots. Normally such edge-crossings are ignored as not representing anything in the structure of G. But when we ask what is the least possible number of edge-crossings in a drawing of G, deep and difficult problems arise. (See [8], 122–3.)

Some examples follow. Fig. I.1.1 is a drawing of the graph of a cube, and Fig. I.1.2 shows a graph with loops and multiple joins.

Some graphs with simple structures are thought to deserve special names. For some purposes it is convenient to recognize a *null* graph, having no edges and no vertices. A *vertex-graph* is an edgeless graph having exactly one vertex (Fig. I.1.3(i)). A *loop-graph* consists of a single loop with its one end (Fig. I.1.3(ii)), and a *link-graph* consists of a single link with its two ends (Fig. I.1.3(iii)).

Let n be a nonnegative integer. Then an *n-clique* is defined as a loopless graph with exactly n vertices and $\frac{1}{2}n(n-1)$ edges, each pair of vertices being joined by a single link. Thus the n-cliques are strict graphs. Evidently the null graphs are the 0-cliques, the vertex-graphs are the 1-cliques, and the link-graphs are the 2-cliques. Figure I.1.4 shows a

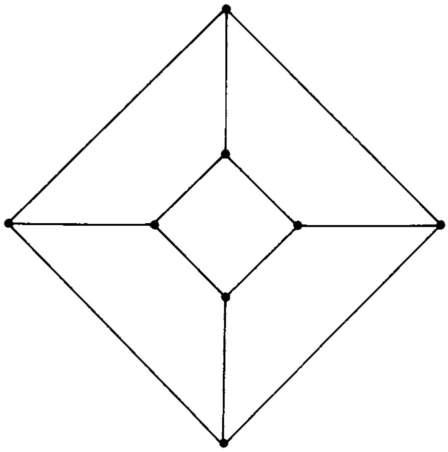

FIGURE I.1.1

3-clique, a 4-clique and a 5-clique, labeled respectively (i), (ii) and (iii). Note the edge-crossing in the diagram of the 5-clique.

Now let n be a positive integer. We define an *n-arc* as a graph G with n edges and $(n+1)$ vertices, having the following property: The edges can be enumerated as $A_1, A_2, \ldots, A_n$, and the vertices as $a_0, a_1, a_2, \ldots, a_n$, in such a way that the ends of A_j are a_{j-1} and a_j, for each relevant suffix j.

The 1-arcs are the link-graphs. Fig. I.1.5 shows a 2-arc a 3-arc and a 4-arc, labelled (i), (ii) and (iii), respectively.

For any positive integer n we define an *n-circuit* as a graph G with n vertices and n edges, having the following property. The vertices can be enumerated as $a_1, a_2, \ldots, a_n$, and the edges as $A_1, A_2, \ldots, A_n$, in such a way

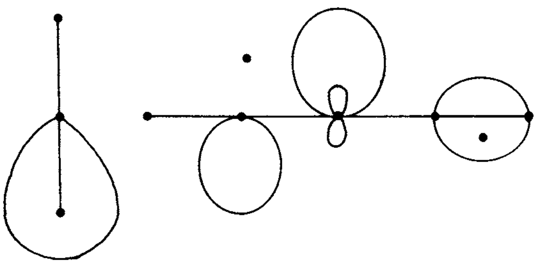

FIGURE I.1.2

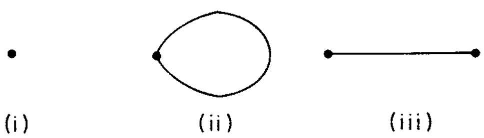

FIGURE I.1.3

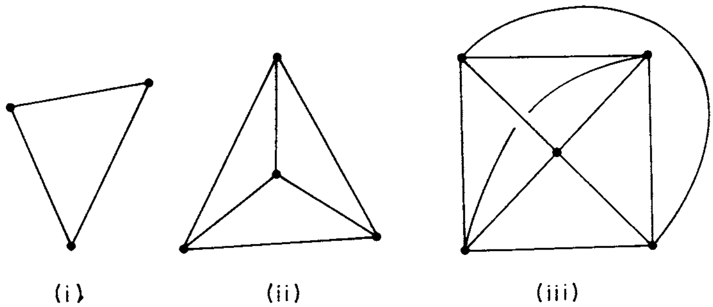

FIGURE I.1.4

that the ends of A_n are a_{j-1} and a_j (where $a_0 = a_n$) for each edge A_j. The 1-circuits are the loop-graphs and the 3-circuits are the 3-cliques. Figure I.1.6 shows a 2-circuit, a 3-circuit, and a 4-circuit, labelled (i), (ii) and (iii), respectively.

When it is unnecessary to assert or reassert the value of the integer n, an n-clique, n-arc, or n-circuit may be called simply a *clique*, *arc*, or *circuit*, respectively. In the two latter cases n is called the *length* of the arc or circuit concerned.

Let us now define the *valency* val(G, x) of a vertex x in a graph G. It is the number of edges incident with x, loops being counted twice. It is clear that

$$\sum_{x \in V(G)} \text{val}(G, x) = 2|E(G)|. \tag{I.1.1}$$

Here and in the following pages we denote the cardinality of a finite set S by $|S|$. From Eq. (I.1.1) we can deduce the following theorem.

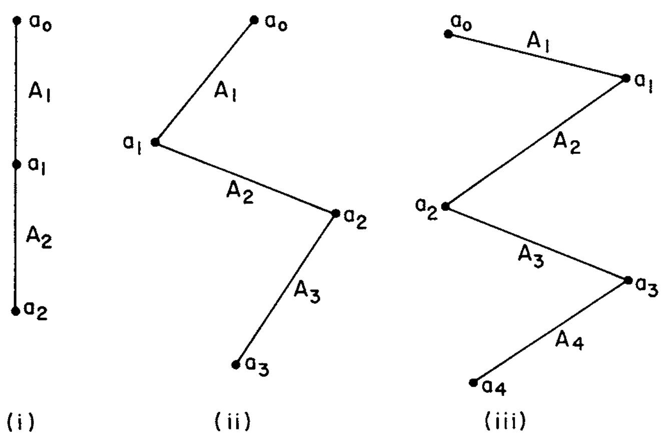

FIGURE I.1.5

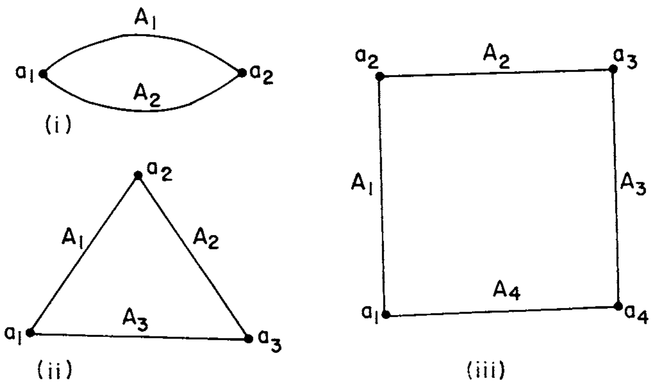

FIGURE I.1.6

Theorem I.1. *In any graph the number of vertices of odd valency is even*. (This is one of the oldest and best-known results in graph theory.)

The term "valency" or "valence" is suggested by a chemical analogy. But many workers in graph theory say "degree" instead.

A vertex of zero valency is said to be *isolated*. For example, the single vertex of a vertex-graph is isolated. It is convenient to describe a vertex of valency 1 as *monovalent*, one of valency 2 as *divalent*, and so on. Thus the single vertex of a loop-graph is divalent and the two vertices of the link-graph are monovalent. Each vertex of a circuit is divalent, and each vertex of an n-clique has valency $(n-1)$.

Consider an n-arc. Let its vertices and edges be enumerated in the manner specified in its definition. Then it is clear that a_0 and a_n are monovalent, and that every other vertex is divalent. It is usual to describe the two monovalent vertices of an arc as its *ends*. The remaining vertices, all divalent, are called the *internal vertices* of the arc.

A graph in which each vertex has the same valency n is called a *regular* graph of that valency. Thus every circuit is a regular graph of valency 2. If the integer m is positive, then each m-clique is a regular graph of valency $(m-1)$. Regular graphs of valency 3 are of special interest in graph theory. They are called the *cubic* graphs.

I.2. ISOMORPHISM

Given two graphs, such as the graphs of a large and a small cube, we may decide that they are not essentially different. We may say that they have the same structure, or that each is a copy of the other. What is meant is that the two graphs can be represented by the same diagram. To make our meaning precise we may appeal to the following notion of *isomorphism*.

Let G and H be graphs. Let f be a 1–1 mapping of $V(G)$ onto $V(H)$, and g a 1–1 mapping of $E(G)$ onto $E(H)$. Let θ denote the ordered pair (f, g). We say that θ is an *isomorphism* of G onto H if the following condition holds: The vertex x is incident with the edge A in G if and only if the vertex fx is incident with the edge gA in H.

If such an isomorphism θ exists, we say that the graphs G and H are *isomorphic*. Clearly we then have

$$|V(G)| = |V(H)| \qquad \text{and} \qquad |E(G)| = |E(H)|.$$

We may think of θ as an operation transforming G into H, and we accordingly write $\theta G = H$. It is also convenient to write $\theta v = fv$ and $\theta A = gA$ for each vertex v and edge A of G. Clearly G and H can be represented by the same diagram. The representative of an edge or vertex x of G can be reinterpreted as the representative of θx in H.

It is possible for G and H to be the same graph. An isomorphism of G onto itself is called an *automorphism* of G. Any graph G has the *identical* or *trivial* automorphism I such that $Ix = x$ for each edge or vertex x of G.

The relation of isomorphism between graphs is reflexive, because of the trivial automorphisms. It is symmetrical; if $\theta = (f, g)$ is an isomorphism of G onto H, then there is an inverse isomorphism $\theta^{-1} = (f^{-1}, g^{-1})$ of H onto G. Finally, it is transitive; if $\theta = (f, g)$ is an isomorphism of G onto H and $\phi = (f_1, g_1)$ is an isomorphism of H onto K, then there is an isomorphism $\phi\theta = (f_1 f, g_1 g)$ of G onto K. Here $f_1 f$ is the mapping obtained by applying first f and then f_1. Similarly $\phi\theta$ is the isomorphism obtained by applying first θ and then ϕ. It is easy to verify that the multiplication of isomorphisms thus defined is associative.

We have now verified that isomorphism is an equivalence relation. It therefore partitions the class of all graphs into disjoint nonnull subclasses, called *isomorphism classes*, such that two graphs belong to the same isomorphism class if and only if they are isomorphic.

Pure graph theory is concerned with those properties of graphs that are invariant under isomorphism, for example the number of vertices, the number of loops, the number of links, and the number of vertices of a given valency. It is therefore natural for a graph theorist to identify two graphs that are isomorphic. For example, all link-graphs are isomorphic, and therefore he speaks of the "link-graph" as though there were only one. Similarly one hears of "the null graph," "the vertex-graph," and "the graph of the cube." When this language is used, it is really an isomorphism class (also called an *abstract graph*) that is under discussion.

It is a convention rather than a theorem that all null graphs are isomorphic. We can justify it by postulating that if S and T are null sets, there is a unique 1–1 mapping of S onto T. The author remembers being taught that there is only one null set. Those who adhere to this mystical doctrine must hold that there is only one null graph, the unique member of its isomorphism class. The null graph often fits rather oddly into combinatorial definitions.

Theorem I.2. *Let G and H be strict graphs. Let f be a 1–1 mapping of $V(G)$ onto $V(H)$ having the following property: Two distinct vertices x and y of G are adjacent in G if and only if the corresponding vertices fx and fy of H are adjacent in H. Then there is a uniquely determined 1–1 mapping g of $E(G)$ onto $E(H)$ such that (f, g) is an isomorphism of G onto H.*

Proof. Let A be any edge of G. It has distinct ends x and y, since G is strict. By hypothesis there is a uniquely determined edge A' of H whose ends are fx and fy. We define a 1–1 mapping g by the rule that $gA = A'$ for each edge A of G. It is then clear that (f, g) is an isomorphism of G onto H. Conversely, if there is a g such that (f, g) is an isomorphism of G onto H, then g must satisfy the above rule. □

In a theory of strict graphs an isomorphism of G onto H is often defined as a 1–1 mapping of $V(G)$ onto $V(H)$ that preserves adjacency. This specialization can be regarded as an application of Theorem I.2.

Let n be a non-negative integer, and let G and H be n-cliques. They are strict graphs. There exists a 1–1 mapping f of $V(G)$ onto $V(H)$, and any such mapping preserves adjacency by the definition of an n-clique. Accordingly G and H are isomorphic, by Theorem I.2.

Now let n be a positive integer, and let G and H be n-arcs. In their defining enumerations let the vertices of G be enumerated as $a_0, a_1, \dots, a_n$ and those of H as $b_0, b_1, \dots, b_n$. Let f be the 1–1 mapping of $V(G)$ onto $V(H)$ such that $fa_j = b_j$ for each suffix j. Then f has the property specified in the enunciation of Theorem I.2, and therefore G and H are isomorphic.

A similar argument shows that if n is any integer exceeding 2, then any two n-circuits are isomorphic. It is a trivial matter to extend this result to the cases $n = 1$ and $n = 2$. But the 1-circuit and 2-circuit are not strict graphs, and so Theorem I.2 does not apply. If G and H are 2-circuits, as shown in Fig. I.1.6(i), then any 1–1 mapping of $V(G)$ onto $V(H)$ can be combined with any 1–1 mapping of $E(G)$ onto $E(H)$ to yield an isomorphism of G onto H.

The automorphisms of a graph G are the elements of a group $A(G)$ with respect to the associative multiplication of isomorphisms defined above. The unit element of $A(G)$ is the trivial automorphism. The inverse of an automorphism θ in $A(G)$ is the inverse isomorphism θ^{-1} defined above. A product $\phi\theta$ of isomorphisms is defined only when θ is a mapping onto, and ϕ is a mapping from, the same graph G. But in the case of automorphisms this condition is always satisfied. We call $A(G)$ the *automorphism group* of G.

Mathematical structures other than graphs have their theories of isomorphism. For example, two groups P and Q are said to be isomorphic if there is a 1–1 mapping f of P onto Q such that f and f^{-1} preserve products. We should note the following theorem.

Theorem I.3. *Isomorphic graphs have isomorphic automorphism groups.*

Proof. Let θ be an isomorphism of a graph G onto a graph H. Then if ξ is in $A(H)$, the isomorphism $\theta^{-1}\xi\theta$ is in $A(G)$. Similarly, if η is in $A(G)$, then $\theta\eta\theta^{-1}$ is in $A(H)$. So we have a 1–1 correspondence $\eta \to \theta\eta\theta^{-1}$ of $A(G)$ onto $A(H)$, and this clearly preserves products. □

We can now assert that the automorphism group of a graph, regarded as an abstract group, is invariant under graph isomorphism. It is thus a legitimate concern of the pure graph theorist.

Let us investigate the automorphism groups of the graphs so far defined. For the null graph, the vertex-graph, and the loop-graph, there is only the trivial automorphism. Hence the automorphism group of each of these graphs is "trivial"; that is, it consists of a unit element only.

Now consider an n-clique G, with $n > 0$. Referring to Theorem I.2 we see that any permutation of the n vertices defines a unique corresponding automorphism of G. We deduce that $A(G)$ is isomorphic to the group of permutations of n objects.

Consider next an n-arc g with ends x and y. We have a defining enumeration in which the vertices are enumerated as $a_0, a_1, \ldots, a_n$ and the edges as $A_1, A_2, \ldots, A_n$. We may suppose that $a_0 = x$ and $a_n = y$. We can get a second defining enumeration, called the *reverse* of the first, by reversing the order of edges or vertices in each of the above sequences. Adjusting the notation we can say that in the reverse enumeration $a_0 = y$ and $a_n = x$. However, a defining sequence is uniquely determined by valency considerations as soon as a_0 is given: A_1 must be the single edge incident with a_0, a_1 must be the other end of A_1, A_2 must be the other edge incident with a_1 (if $n > 1$), and so on. Evidently an automorphism of G is determined by an ordered pair of defining enumerations: It maps the first member onto the second. We deduce that $A(G)$ has exactly two elements, say I and θ. The nontrivial automorphism θ interchanges the two ends of G; it changes each defining enumeration into its reverse.

Consider next an n-circuit G, with $n > 1$. We have a defining enumeration in which the vertices are enumerated as $a_1, a_2, \ldots, a_n$ and the edges as $A_1, A_2, \ldots, A_n$. Evidently there is an automorphism θ that increases each suffix by 1. (Addition and subtraction in the suffixes is to be modulo n.) There is also an automorphism ϕ that reverses the original sequence of vertices, and that reverses the original cyclic sequence of edges. Using the powers of θ and their combinations with ϕ, we find that it is possible to obtain a defining enumeration of the circuit in which a_1 is an arbitrary vertex and A_1 is any edge incident with a_1. But it follows from valency considerations that a defining sequence is fixed uniquely when a_1 and A_1 are given. As with the arcs, an automorphism is uniquely determined by an ordered pair of defining enumerations. We deduce that $A(G)$ has exactly $2n$ elements, which we can write as

$$I, \theta, \theta^2, \ldots, \theta^{n-1}, \phi, \theta\phi, \theta^2\phi, \ldots, \theta^{n-1}\phi.$$

Indeed $A(G)$ is a dihedral group of order $2n$; its generators θ and ϕ satisfy the generating relations $\theta^n = I$, $\phi^2 = I$, and $\theta\phi = \phi\theta^{-1}$.

I.3. SUBGRAPHS

In constructing a theory of graphs, we can start with the observation that some graphs contain other graphs. A graph H contained in a graph G is called a *subgraph* of G. We shall use the following precise definition.

A graph H is said to be a *subgraph* of a graph G if

$$V(H) \subseteq V(G), \qquad E(H) \subseteq E(G),$$

and each edge of H has the same ends in H as in G. Under these conditions we say also that H is *contained* in G.

According to the above definition, G is itself a subgraph of G. The other subgraphs of G are called its *proper* subgraphs. They are said to be *properly* contained in G. We write $H \subseteq G$ to indicate that H is a subgraph of G, and $H \subset G$ to indicate that H is a proper subgraph of G. We use a corresponding notation for finite sets. Thus $S \subseteq T$ means that S is a subset of T, and $S \subset T$ that S is a proper subset of T.

The above formal definition leads to the following rule for constructing subgraphs.

Rule I.4. *Let G be a graph, P a subset of $V(G)$ and Q a subset of $E(G)$. Then the necessary and sufficient condition for the existence of a subgraph H of G such that $V(H) = P$ and $E(H) = Q$ is that P shall contain both ends of each member of Q.*

Condition I.4 is always satisfied when $P = V(G)$. We then call H a *spanning* subgraph of G. We shall denote the spanning subgraph of G with edge-set S by $G:S$. Those spanning subgraphs of G that are circuits are called its *Hamiltonian circuits*; those that are regular graphs of valency 1 are called its 1-*factors*. Both Hamiltonian circuits and 1-factors are prominent in the graph-theoretical literature.

If S is any set of edges of G, let $J(S)$ denote the set of all vertices v of G such that v is incident with a member of S. By Rule I.4 there is a subgraph H of G such that $V(H) = J(S)$ and $E(H) = S$. We call H the *reduction* of G to S, and we denote it by $G \cdot S$. We note that a reduction can have no isolated vertex. In particular $G \cdot E(G)$ is the graph obtained from G by deleting all its isolated vertices.

If U is any set of vertices of G, let $J(U)$ denote the set of all edges of G having both ends in U. By Rule I.4 there is a subgraph H of G such that

$$V(H) = U \quad \text{and} \quad E(H) = J(U).$$

We call such a subgraph an *induced* subgraph of G. In particular, H is the subgraph of G induced by U; we denote it by $G[U]$.

Each graph G has a *null subgraph*, which we denote by $\varnothing$. This null subgraph is both a reduction and an induced subgraph of G, but it is not a spanning subgraph of G unless G itself is null.

Let H and K be subgraphs of G, not necessarily distinct. By I.4 there is a subgraph L of G such that

$$V(L) = V(H) \cup V(K) \quad \text{and} \quad E(L) = E(H) \cup E(K).$$

We call L the *union* of H and K, and we denote it by $H \cup K$. By I.4 there is also a subgraph M of G such that

$$V(M) = V(H) \cap V(K) \quad \text{and} \quad E(M) = E(H) \cap E(K).$$

We call M the *intersection* of H and K, and we denote it by $H \cap K$. For obvious reasons, unions and intersections of subgraphs obey algebraic rules like those for unions and intersections of subsets of a finite set. For example,

$$(H \cup K) \cap L = (H \cap L) \cup (K \cap L). \tag{I.3.1}$$

Two subgraphs H and K of G are said to be *disjoint* if they have no common edge and no common vertex. One way of asserting the disjointness is to write

$$H \cap K = \varnothing.$$

There is a similar device for asserting that a graph G is nonnull: We write $\varnothing \subset G$.

The following theorem is an easy consequence of the definition of a subgraph.

Theorem I.5. *Any subgraph of a subgraph of G is a subgraph of G.*

We can go further. Any spanning subgraph of a spanning subgraph of G is a spanning subgraph of G, any reduction of a reduction of G is a reduction of G, and any induced subgraph of an induced subgraph of G is an induced subgraph of G.

We conclude the section with a notation for relating the subgraphs of two isomorphic graphs G and H. Let $\theta = (f, g)$ be an isomorphism of G onto H.

If S is any subset of $V(G)$, then f maps S onto a subset fS of $V(H)$. This is the set of images under f of the members of S. We can denote it also by θS. The *restriction $f \cdot S$ of f to S* is the 1–1 mapping of S onto fS that agrees with f for every member of S. Similarly, if T is any subset of $E(G)$, it