3

Connected Graphs

Contents

3.1 Walks and Connection

Walks

In Section 1.1, the notion of connectedness was defined in terms of edge cuts. Here, we give an alternative definition based on the notion of a walk in a graph.

A *walk* in a graph G is a sequence $W := v_0 e_1 v_1 \ldots v_{\ell-1} e_\ell v_\ell$, whose terms are alternately vertices and edges of G (not necessarily distinct), such that v_{i-1} and v_i are the ends of e_i, $1 \leq i \leq \ell$. (We regard loops as giving rise to distinct walks with the same sequence, because they may be traversed in either sense. Thus if e is a loop incident with a vertex v, we count the walk vev not just once, but twice.) If $v_0 = x$ and $v_\ell = y$, we say that W *connects* x to y and refer to W as an *xy-walk*. The vertices x and y are called the *ends* of the walk, x being its *initial vertex* and

y its *terminal vertex*; the vertices $v_1, \ldots, v_{\ell-1}$ are its *internal vertices*. The integer ℓ (the number of edge terms) is the *length* of W. An *x-walk* is a walk with initial vertex x. If u and v are two vertices of a walk W, where u precedes v on W, the subsequence of W starting with u and ending with v is denoted by uWv and called the *segment* of W from u to v. The notation uWv is also used simply to signify a uv-walk W.

In a simple graph, a walk $v_0e_1v_1 \ldots v_{\ell-1}e_\ell v_\ell$ is determined, and is commonly specified, by the sequence $v_0v_1 \ldots v_\ell$ of its vertices. Indeed, even if a graph is not simple, we frequently refer to a sequence of vertices in which consecutive terms are adjacent vertices as a 'walk'. In such cases, it should be understood that the discussion is valid for any walk with that vertex sequence. This convention is especially useful in discussing paths, which may be viewed as walks whose vertices (and edges) are distinct.

A walk in a graph is *closed* if its initial and terminal vertices are identical, and is a *trail* if all its edge terms are distinct. A closed trail of positive length whose initial and internal vertices are distinct is simply the sequence of vertices and edges of a cycle. Reciprocally, with any cycle one may associate a closed trail whose terms are just the vertices and edges of the cycle. Even though this correspondence is not one-to-one (the trail may start and end at any vertex of the cycle, and traverse it in either sense), we often specify a cycle by describing an associated closed trail and refer to that trail as the cycle itself.

Connection

Connectedness of pairs of vertices in a graph G is an equivalence relation on V. Clearly, each vertex x is connected to itself by the trivial walk $W := x$; also, if x is connected to y by a walk W, then y is connected to x by the walk obtained on reversing the sequence W; finally, for any three vertices, x, y, and z of G, if xWy and $yW'z$ are walks, the sequence $xWyW'z$, obtained by concatenating W and W' at y, is a walk; thus, if x is connected to y and y is connected to z, then x is connected to z. The equivalence classes determined by this relation of connectedness are simply the vertex sets of the components of G (Exercise 3.1.3).

If there is an xy-walk in a graph G, then there is also an xy-path (Exercise 3.1.1). The length of a shortest such path is called the *distance* between x and y and denoted $d_G(x,y)$. If there is no path connecting x and y (that is, if x and y lie in distinct components of G), we set $d_G(x,y) := \infty$.

We may extend the notion of an xy-path to paths connecting subsets X and Y of V. An *(X,Y)-path* is a path which starts at a vertex of X, ends at a vertex of Y, and whose internal vertices belong to neither X nor Y; if F_1 and F_2 are subgraphs of a graph G, we write *(F_1,F_2)-path* instead of $(V(F_1),V(F_2))$-path. A useful property of connected graphs is that any two nonempty sets of vertices (or subgraphs) are connected by such a path (Exercise 3.1.4).

PROOF TECHNIQUE: EIGENVALUES

We saw in Chapter 2 how certain problems can be solved by making use of arguments involving linear independence. Another powerful linear algebraic tool involves the computation of eigenvalues of appropriate matrices. Although this technique is suitable only for certain rather special problems, it is remarkably effective when applicable. Here is an illustration.

A *friendship graph* is a simple graph in which any two vertices have exactly one common neighbour. By using a clever mixture of graph-theoretical and eigenvalue arguments, Erdős et al. (1966) proved that all friendship graphs have a very simple structure.

Theorem 3.1 THE FRIENDSHIP THEOREM
Let G be a simple graph in which any two vertices (people) have exactly one common neighbour (friend). Then G has a vertex of degree $n-1$ (a politician, everyone's friend).

Proof Suppose the theorem false, and let G be a friendship graph with $\Delta < n-1$. Let us show first of all that G is regular. Consider two nonadjacent vertices x and y, where, without loss of generality, $d(x) \geq d(y)$. By assumption, x and y have exactly one common neighbour, z. For each neighbour v of x other than z, denote by $f(v)$ the common neighbour of v and y. Then f is a one-to-one mapping from $N(x) \setminus \{z\}$ to $N(y) \setminus \{z\}$. Because $|N(x)| = d(x) \geq d(y) = |N(y)|$, we conclude that f is a bijection and hence that $d(x) = d(y)$. Thus any two nonadjacent vertices of G have the same degree; equivalently, any two adjacent vertices of $\overline{G}$ have the same degree.

In order to prove that G is regular, it therefore suffices to show that $\overline{G}$ is connected. But $\overline{G}$ has no singleton component, because $\delta(\overline{G}) = n-1-\Delta(G) > 0$, and cannot have two components of order two or more, because G would then contain a 4-cycle, thus two vertices with two common neighbours. Therefore G is k-regular for some positive integer k. Moreover, by counting the number of 2-paths in G in two ways, we have $n\binom{k}{2} = \binom{n}{2}$; that is, $n = k^2 - k + 1$.

Let $\mathbf{A}$ be the adjacency matrix of G. Then (Exercise 3.1.2) $\mathbf{A}^2 = \mathbf{J} + (k-1)\mathbf{I}$, where $\mathbf{J}$ is the $n \times n$ matrix all of whose entries are 1, and $\mathbf{I}$ is the $n \times n$ identity matrix. Because the eigenvalues of $\mathbf{J}$ are 0, with multiplicity $n-1$, and n, with multiplicity 1, the eigenvalues of $\mathbf{A}^2$ are $k-1$, with multiplicity $n-1$, and $n+k-1 = k^2$, with multiplicity 1. The graph G therefore has eigenvalues $\pm\sqrt{k-1}$, with total multiplicity $n-1$, and k, with multiplicity 1 (see Exercise 1.1.22a).

Because G is simple, the sum of its eigenvalues, the trace of $\mathbf{A}$, is zero. Thus $t\sqrt{k-1} = k$ for some integer t. But this implies that $k = 2$ and $n = 3$, contradicting the assumption that $\Delta < n-1$. $\square$

Further applications of eigenvalues are outlined in Exercises 3.1.11 and 3.1.12.

The above notions apply equally to digraphs. If $W := v_0 a_1 v_1 \ldots v_{\ell-1} a_\ell v_\ell$ is a walk in a digraph, an arc a_i of W is a *forward arc* if v_{i-1} is the tail of a_i and v_i is its head, and a *reverse arc* if v_i is the tail of a_i and v_{i-1} its head. The sets of forward and reverse arcs of W are denoted by W^+ and W^-, respectively. Walks in which all arcs are forward arcs, called directed walks, are discussed in Section 3.4.

Connectedness plays an essential role in applications of graph theory. For example, the graph representing a communications network needs to be connected for communication to be possible between all vertices. Connectedness also plays a basic role in theoretical considerations. For instance, in developing an algorithm to determine whether a given graph is planar, we may restrict our attention to connected graphs, because a graph is planar if and only if each of its components is planar.

Exercises

$\star$**3.1.1** If there is an xy-walk in a graph G, show that there is also an xy-path in G.

3.1.2 Let G be a graph with vertex set V and adjacency matrix $\mathbf{A} = (a_{uv})$. Show that the number of uv-walks of length k in G is the (u, v) entry of $\mathbf{A}^k$.

$\star$**3.1.3** Show that the equivalence classes determined by the relation of connectedness between vertices are precisely the vertex sets of the components of the graph.

$\star$**3.1.4** Show that a graph G is connected if and only if there is an (X, Y)-path in G for any two nonempty subsets X and Y of V.

3.1.5 Show that, in any graph G, the distance function satisfies the *triangle inequality*: for any three vertices x, y, and z, $d(x, y) + d(y, z) \geq d(x, z)$.

3.1.6 Power of a Graph
The *kth power* of a simple graph $G = (V, E)$ is the graph G^k whose vertex set is V, two distinct vertices being adjacent in G^k if and only if their distance in G is at most k. The graph G^2 is referred to as the *square* of G, the graph G^3 as the *cube* of G. Consider P_n^k, the kth power of a path on n vertices, where $n > k^2 + k$. Show that:

a) $d(P_n^k) > 2k - 1$,
b) $\delta(F) \leq k$ for every induced subgraph F of P_n^k.

3.1.7 Diameter
The *diameter* of a graph G is the greatest distance between two vertices of G.

a) Let G be a simple graph of diameter greater than three. Show that $\overline{G}$ has diameter less than three.
b) Deduce that every self-complementary graph has diameter at most three.

c) For $k = 0, 1, 2, 3$, give an example of a self-complementary graph of diameter k, if there is one.

3.1.8 Show that if G is a simple graph of diameter two with $\Delta = n - 2$, then $m \geq 2n - 4$.

3.1.9 Show that the incidence graph of a finite projective plane has diameter three.

3.1.10 If the girth of a graph is at least $2k$, show that its diameter is at least k.

3.1.11

a) Let G_1 and G_2 be edge-disjoint copies of the Petersen graph on the same vertex set. Show that 2 is an eigenvalue of $G_1 \cup G_2$ by proceeding as follows.
   i) Observe that $\mathbf{1}$ is an eigenvector of both G_1 and G_2 corresponding to the eigenvalue 3.
   ii) Let S_1 and S_2 denote the eigenspaces of G_1 and G_2, respectively, corresponding to the eigenvalue 1. (Since 1 is an eigenvalue of the Petersen graph with multiplicity five, S_1 and S_2 are 5-dimensional subspaces of $\mathbb{R}^{10}$.) Using the fact that $\mathbf{1}$ is orthogonal to both S_1 and S_2, show that the dimension of $S_1 \cap S_2$ is at least one.
   iii) Noting that $\mathbf{A}_{G_1 \cup G_2} = \mathbf{A}_{G_1} + \mathbf{A}_{G_2}$, show that any nonzero vector in $S_1 \cap S_2$ is an eigenvector of $G_1 \cup G_2$ corresponding to the eigenvalue 2.

b) Appealing now to Exercises 1.3.2 and 1.3.11, conclude that K_{10} cannot be decomposed into three copies of the Petersen graph. (A.J. Schwenk)

3.1.12 Moore Graph

A *Moore graph of diameter* d is a regular graph of diameter d and girth $2d + 1$. Consider a k-regular Moore graph G of diameter two.

a) Show that $n = k^2 + 1$.

b) Let $\mathbf{A}$ be the adjacency matrix of G and $\operatorname{tr}(\mathbf{A})$ its trace.
   i) Show that $\operatorname{tr}(\mathbf{A}) = 0$.
   ii) Evaluate the matrix $\mathbf{A}^2 + \mathbf{A}$, determine its eigenvalues and their multiplicities, and deduce the possible eigenvalues of $\mathbf{A}$ (but not their multiplicities).
   iii) Expressing $\operatorname{tr}(\mathbf{A})$ in terms of the eigenvalues of $\mathbf{A}$ and their multiplicities, and noting that these multiplicities are necessarily integers, conclude that such a graph G can exist only if $k = 2, 3, 7$, or 57.

(A.J. Hoffman and R.R. Singleton)

c) Find such a graph G for $k = 2$ and $k = 3$.

(A 7-regular example, the *Hoffman–Singleton graph*, discovered by Hoffman and Singleton (1960), is depicted in Figure 3.1; vertex i of P_j is joined to vertex $i + jk \pmod 5$ of Q_k. A 57-regular example would have 3250 vertices. No such graph is known.)

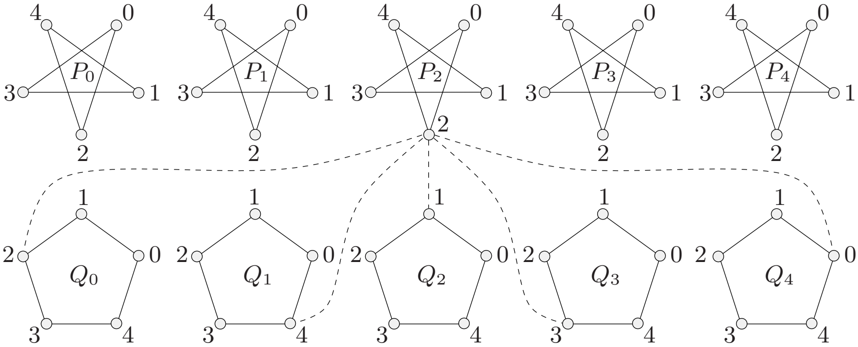

Fig. 3.1. The Hoffman–Singleton graph

3.1.13 Cage

A k-regular graph of girth g with the least possible number of vertices is called a *(k, g)-cage*. A $(3, g)$-cage is often simply referred to as a *g-cage*. Let $f(k, g)$ denote the number of vertices in a (k, g)-cage. Observe that $f(2, g) = g$.

a) For $k \geq 3$, show that:
   i) $f(k, 2r) \geq (2(k-1)^r - 2)/(k-2)$,
   ii) $f(k, 2r+1) \geq (k(k-1)^r - 2)/(k-2)$.
b) Determine all g-cages, $g = 3, 4, 5, 6$.
c) Show that the incidence graph of a projective plane of order $k-1$ is a $(k, 6)$-cage.

(Singleton (1966) showed, conversely, that any $(k, 6)$-cage of order $2(k^2 - k + 1)$ is necessarily the incidence graph of a projective plane of order $k - 1$.)

3.1.14 The Tutte–Coxeter Graph

A highly symmetric cubic graph, known as *the Tutte–Coxeter graph*, is shown in Figure 3.2. Show that:

a) the Tutte–Coxeter graph is isomorphic to the bipartite graph $G[X, Y]$ derived from K_6 in the following manner. The vertices of X are the fifteen edges of K_6 and the vertices of Y are the fifteen 1-factors of K_6, an element e of X being adjacent to an element F of Y whenever e is an edge of the 1-factor F.

(H.S.M. Coxeter)

b) the Tutte–Coxeter graph is an 8-cage.
   (Tutte (1947b) showed that this graph is, in fact, the unique 8-cage.)

3.1.15 t-Arc-Transitive Graph

A walk $(v_0, v_1, \ldots, v_t)$ in a graph such that $v_{i-1} \neq v_{i+1}$, for $1 \leq i \leq t-1$, is called a *t-arc*. A simple connected graph G is *t-arc-transitive* if, given any two t-arcs $(v_0, v_1, \ldots, v_t)$ and $(w_0, w_1, \ldots, w_t)$, there is an automorphism of G which maps v_i to w_i, for $0 \leq i \leq t$. (Thus a 1-arc-transitive graph is the same as an arc-transitive graph, defined in Exercise 1.5.12.) Show that:

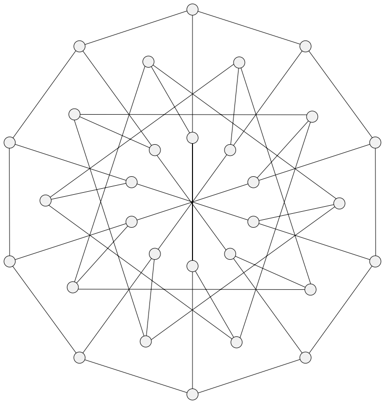

Fig. 3.2. The Tutte–Coxeter graph: the 8-cage

a) $K_{3,3}$ is 2-arc-transitive,
b) the Petersen graph is 3-arc-transitive,
c) the Heawood graph is 4-arc-transitive,
d) the Tutte–Coxeter graph is 5-arc-transitive.

(Tutte (1947b) showed that there are no t-arc-transitive cubic graphs when $t > 5$.)

3.2 Cut Edges

For any edge e of a graph G, it is easy to see that either $c(G \setminus e) = c(G)$ or $c(G \setminus e) = c(G) + 1$ (Exercise 3.2.1). If $c(G \setminus e) = c(G) + 1$, the edge e is called a *cut edge* of G. Thus a cut edge of a connected graph is one whose deletion results in a disconnected graph. More generally, the cut edges of a graph correspond to its bonds of size one (Exercise 3.2.2).

The graph in Figure 3.3 has three cut edges.

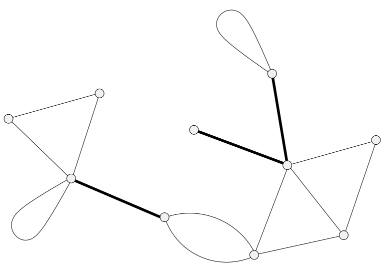

Fig. 3.3. The cut edges of a graph

If e is a cut edge of a graph G, its ends x and y belong to different components of $G \setminus e$, and so are not connected by a path in $G \setminus e$; equivalently, e lies in no cycle of G. Conversely, if $e = xy$ is not a cut edge of G, the vertices x and y belong to the same component of $G \setminus e$, so there is an xy-path P in $G \setminus e$, and $P + e$ is a cycle in G through e. Hence we have the following characterization of cut edges.

Proposition 3.2 *An edge e of a graph G is a cut edge if and only if e belongs to no cycle of G.* □

Exercises

$\star$**3.2.1** Show that if $e \in E$, then either $c(G \setminus e) = c(G)$ or $c(G \setminus e) = c(G) + 1$.

$\star$**3.2.2** Show that an edge e is a cut edge of a graph G if and only if $\{e\}$ is a bond of G.

3.2.3 Let G be a connected even graph. Show that:

a) G has no cut edge,
b) for any vertex $v \in V$, $c(G - v) \leq \frac{1}{2}d(v)$.

3.2.4 Let G be a k-regular bipartite graph with $k \geq 2$. Show that G has no cut edge.

3.3 Euler Tours

A trail that traverses every edge of a graph is called an *Euler trail*, because Euler (1736) was the first to investigate the existence of such trails. In the earliest known paper on graph theory, he showed that it was impossible to cross each of the seven bridges of Königsberg once and only once during a walk through the town. A plan of Königsberg and the river Pregel is shown in Figure 3.4a. As can be seen, proving that such a walk is impossible amounts to showing that the graph in Figure 3.4b has no Euler trail.

A *tour* of a connected graph G is a closed walk that traverses each edge of G at least once, and an *Euler tour* one that traverses each edge exactly once (in other words, a closed Euler trail). A graph is *eulerian* if it admits an Euler tour.

Fleury's Algorithm

Let G be an eulerian graph, and let W be an Euler tour of G with initial and terminal vertex u. Each time a vertex v occurs as an internal vertex of W, two edges incident with v are accounted for. Since an Euler tour traverses each edge

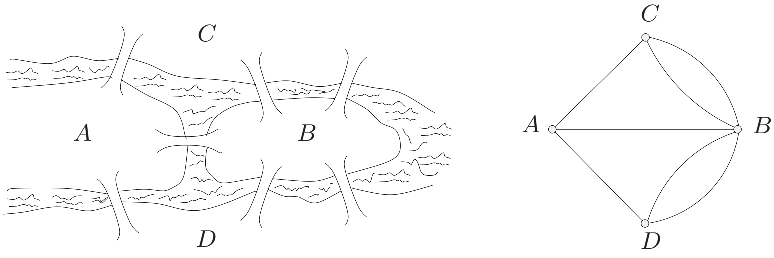

(a) (b)

Fig. 3.4. The bridges of Königsberg and their graph

exactly once, $d(v)$ is even for all $v \neq u$. Similarly, $d(u)$ is even, because W both starts and ends at u. Thus an eulerian graph is necessarily even.

The above necessary condition for the existence of an Euler tour in a connected graph also turns out to be sufficient. Moreover, there is a simple algorithm, due to Fleury (1883), which finds an Euler tour in an arbitrary connected even graph G (see also Lucas (1894)). Fleury's Algorithm constructs such a tour of G by tracing out a trail subject to the condition that, at any stage, a cut edge of the untraced subgraph F is taken only if there is no alternative.

Algorithm 3.3 FLEURY'S ALGORITHM

INPUT: a connected even graph G and a specified vertex u of G
OUTPUT: an Euler tour W of G starting (and ending) at u

1: *set* $W := u$, $x := u$, $F := G$
2: ***while*** $\partial_F(x) \neq \emptyset$ ***do***
3: *choose an edge* $e := xy \in \partial_F(x)$, *where* e *is not a cut edge of* F *unless there is no alternative*
4: *replace* uWx *by* $uWxey$, x *by* y, *and* F *by* $F \setminus e$
5: ***end while***
6: *return* W

Theorem 3.4 *If G is a connected even graph, the walk W returned by Fleury's Algorithm is an Euler tour of G.*

Proof The sequence W is initially a trail, and remains one throughout the procedure, because Fleury's Algorithm always selects an edge of F (that is, an as yet unchosen edge) which is incident to the terminal vertex x of W. Moreover, the algorithm terminates when $\partial_F(x) = \emptyset$, that is, when all the edges incident to the terminal vertex x of W have already been selected. Because G is even, we deduce that $x = u$; in other words, the trail W returned by the algorithm is a closed trail of G.

Suppose that W is not an Euler tour of G. Denote by X the set of vertices of positive degree in F when the algorithm terminates. Then $X \neq \emptyset$, and $F[X]$ is an even subgraph of G. Likewise $V \setminus X \neq \emptyset$, because $u \in V \setminus X$. Since G is connected, $\partial_G(X) \neq \emptyset$. On the other hand, $\partial_F(X) = \emptyset$. The last edge of $\partial_G(X)$ selected for inclusion in W was therefore a cut edge $e = xy$ of F at the time it was chosen, with $x \in X$ and $y \in V \setminus X$ (see Figure 3.5). But this violates the rule for choosing the next edge of the trail W, because the edges in $\partial_F(x)$, which were also candidates for selection at the time, were not cut edges of F, by Theorem 2.10. □

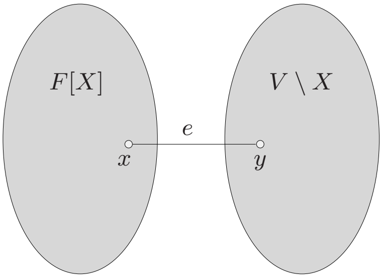

Fig. 3.5. Choosing a cut edge in Fleury's Algorithm

The validity of Fleury's Algorithm provides the following characterization of eulerian graphs.

Theorem 3.5 *A connected graph is eulerian if and only if it is even.* □

Let now x and y be two distinct vertices of a graph G. Suppose that we wish to find an Euler xy-trail of G, if one exists. We may do so by adding a new edge e joining x and y. The graph G has an Euler trail connecting x and y if and only if $G+e$ has an Euler tour (Exercise 3.3.3). Thus Fleury's Algorithm may be adapted easily to find an Euler xy-trail in G, if one exists.

We remark that Fleury's Algorithm is an efficient algorithm, in a sense to be made precise in Chapter 8. When an edge is considered for inclusion in the current trail W, it must be examined to determine whether or not it is a cut edge of the remaining subgraph F. If it is not, it is appended to W right away. On the other hand, if it is found to be a cut edge of F, it remains a cut edge of F until it is eventually selected for inclusion in W; therefore, each edge needs to be examined only once. In Chapter 7, we present an efficient algorithm for determining whether or not an edge is a cut edge of a graph.

A comprehensive treatment of eulerian graphs and related topics can be found in Fleischner (1990, 1991).

Exercises

3.3.1 Which of the pictures in Figure 3.6 can be drawn without lifting one's pen from the paper and without tracing a line more than once?

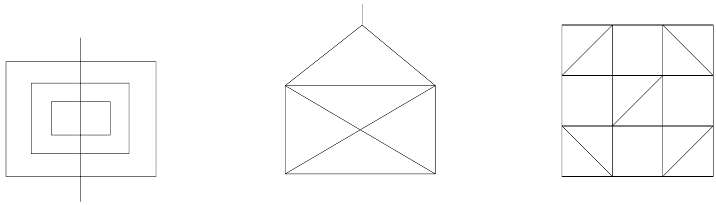

Fig. 3.6. Tracing pictures

3.3.2 If possible, give an example of an eulerian graph G with n even and m odd. Otherwise, explain why there is no such graph.

$\star$**3.3.3** Let G be a graph with two distinct specified vertices x and y, and let $G+e$ be the graph obtained from G by the addition of a new edge e joining x and y.

a) Show that G has an Euler trail connecting x and y if and only if $G+e$ has an Euler tour.
b) Deduce that G has an Euler trail connecting x and y if and only if $d(x)$ and $d(y)$ are odd and $d(v)$ is even for all $v \in V \setminus \{x, y\}$.

3.3.4 Let G be a connected graph, and let X be the set of vertices of G of odd degree. Suppose that $|X| = 2k$, where $k \geq 1$.

a) Show that there are k edge-disjoint trails $Q_1, Q_2, \ldots, Q_k$ in G such that $E(G) = E(Q_1) \cup E(Q_2) \cup \ldots \cup E(Q_k)$.
b) Deduce that G contains k edge-disjoint paths connecting the vertices of X in pairs.

3.3.5 Let G be a nontrivial eulerian graph, and let $v \in V$. Show that each v-trail in G can be extended to an Euler tour of G if and only if $G - v$ is acyclic.

(O. Ore)

3.3.6 Dominating Subgraph
A subgraph F of a graph G is *dominating* if every edge of G has at least one end in F. Let G be a graph with at least three edges. Show that $L(G)$ is hamiltonian if and only if G has a dominating eulerian subgraph.

(F. Harary and C.St.J.A. Nash-Williams)

3.3.7 A cycle decomposition of a loopless eulerian graph G induces a family of pairs of edges of G, namely the consecutive pairs of edges in the cycles comprising the decomposition. Each edge thus appears in two pairs, and each trivial edge cut $\partial(v)$, $v \in V$, is partitioned into pairs. An Euler tour of G likewise induces a family of pairs of edges with these same two properties. A cycle decomposition and Euler tour are said to be *compatible* if, for all vertices v, the resulting partitions of

$\partial(v)$ have no pairs in common. Show that every cycle decomposition of a loopless eulerian graph of minimum degree at least four is compatible with some Euler tour. (A. Kotzig)
(G. Sabidussi has conjectured that, conversely, every Euler tour of a loopless eulerian graph of minimum degree at least four is compatible with some cycle decomposition; see Appendix A.)

3.4 Connection in Digraphs

As we saw earlier, in Section 3.1, the property of connection in graphs may be expressed not only in terms of edge cuts but also in terms of walks. By the same token, the property of strong connection, defined in terms of outcuts in Section 2.5, may be expressed alternatively in terms of directed walks. This is an immediate consequence of Theorem 3.6 below.

A *directed walk* in a digraph D is an alternating sequence of vertices and arcs

$$W := (v_0, a_1, v_1, \ldots, v_{\ell-1}, a_\ell, v_\ell)$$

such that v_{i-1} and v_i are the tail and head of a_i, respectively, $1 \leq i \leq \ell$.[1] If x and y are the initial and terminal vertices of W, we refer to W as a *directed* (x, y)*-walk*. Directed trails, tours, paths, and cycles in digraphs are defined analogously. As for undirected graphs, the (u, v)*-segment* of a directed walk W, where u and v are two vertices of W, u preceding v, is the subsequence of W starting with u and ending with v, and is denoted uWv (the same notation as for undirected graphs).

We say that a vertex y is *reachable* from a vertex x if there is a directed (x, y)-path. The property of reachability can be expressed in terms of outcuts, as follows.

Theorem 3.6 *Let x and y be two vertices of a digraph D. Then y is reachable from x in D if and only if $\partial^+(X) \neq \emptyset$ for every subset X of V which contains x but not y.*

Proof Suppose, first, that y is reachable from x by a directed path P. Consider any subset X of V which contains x but not y. Let u be the last vertex of P which belongs to X and let v be its successor on P. Then $(u, v) \in \partial^+(X)$, so $\partial^+(X) \neq \emptyset$.

Conversely, suppose that y is not reachable from x, and let X be the set of vertices which are reachable from x. Then $x \in X$ and $y \notin X$. Furthermore, because no vertex of $V \setminus X$ is reachable from x, the outcut $\partial^+(X)$ is empty. □

In a digraph D, two vertices x and y are *strongly connected* if there is a directed (x, y)-walk and also a directed (y, x)-walk (that is, if each of x and y is reachable from the other). Just as connection is an equivalence relation on the vertex set of a graph, strong connection is an equivalence relation on the vertex set of a digraph (Exercise 3.4.1). The subdigraphs of D induced by the equivalence classes

[1] Thus a walk in a graph corresponds to a directed walk in its associated digraph. This is consistent with our convention regarding the traversal of loops in walks.

with respect to this relation are called the *strong components* of D. The strong components of the digraph shown in Figure 3.7a are indicated in Figure 3.7b. We leave it to the reader to verify that a digraph is strong if and only if it has exactly one strong component (Exercise 3.4.2).

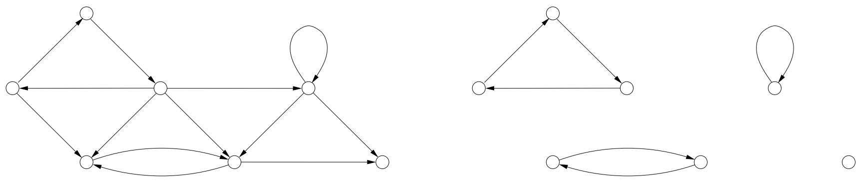

(a) (b)

Fig. 3.7. (a) A digraph and (b) its strong components

A *directed Euler trail* is a directed trail which traverses each arc of the digraph exactly once, and a *directed Euler tour* is a directed tour with this same property. A digraph is *eulerian* if it admits a directed Euler tour. There is a directed version of Theorem 3.5, whose proof we leave as an exercise (3.4.8).

Theorem 3.7 *A connected digraph is eulerian if and only if it is even.* □

Exercises

$\star$**3.4.1** Show that strong connection is an equivalence relation on the vertex set of a digraph.

$\star$**3.4.2** Show that a digraph is strong if and only if it has exactly one strong component.

$\star$**3.4.3** Let C be a strong component of a digraph D, and let P be a directed path in D connecting two vertices of C. Show that P is contained in C.

3.4.4 Let D be a digraph with adjacency matrix $\mathbf{A} = (a_{uv})$. Show that the number of directed (u, v)-walks of length k in D is the (u, v) entry of $\mathbf{A}^k$.

3.4.5 Show that every tournament is either strong or can be transformed into a strong tournament by the reorientation of just one arc.

$\star$**3.4.6** Condensation of a Digraph

a) Show that all the arcs linking two strong components of a digraph have their tails in one strong component (and their heads in the other).

b) The *condensation* $C(D)$ of a digraph D is the digraph whose vertices correspond to the strong components of D, two vertices of $C(D)$ being linked by an arc if and only if there is an arc in D linking the corresponding strong components, and with the same orientation. Draw the condensations of:
   i) the digraph of Figure 3.7a,
   ii) the four tournaments of Figure 1.25.
c) Show that the condensation of any digraph is acyclic.
d) Deduce that:
   i) every digraph has a *minimal* strong component, namely one that dominates no other strong component,
   ii) the condensation of any tournament is a transitive tournament.

3.4.7 A digraph is *unilateral* if any two vertices x and y are connected either by a directed (x, y)-path or by a directed (y, x)-path, or both. Show that a digraph is unilateral if and only if its condensation has a directed Hamilton path.

$\star$**3.4.8** Prove Theorem 3.7.

3.4.9 de Bruijn–Good Digraph
The *de Bruijn–Good digraph* BG_n has as vertex set the set of all binary sequences of length n, vertex $a_1a_2 \ldots a_n$ being joined to vertex $b_1b_2 \ldots b_n$ if and only if $a_{i+1} = b_i$ for $1 \le i \le n-1$. Show that BG_n is an eulerian digraph of order 2^n and directed diameter n.

3.4.10 de Bruijn–Good Sequence
A circular sequence $s_1s_2 \ldots s_{2^n}$ of zeros and ones is called a *de Bruijn–Good sequence* of order n if the 2^n subsequences $s_is_{i+1} \ldots s_{i+n-1}$, $1 \le i \le 2^n$ (where subscripts are taken modulo 2^n) are distinct, and so constitute all possible binary sequences of length n. For example, the sequence 00011101 is a de Bruijn–Good sequence of order three. Show how to derive such a sequence of any order n by considering a directed Euler tour in the de Bruijn–Good digraph BG_{n-1}. (N.G. de Bruijn; I.J. Good)
(An application of de Bruijn–Good sequences can be found in Chapter 10 of Bondy and Murty (1976).)

$\star$**3.4.11**

a) Show that a digraph which has a closed directed walk of odd length contains a directed odd cycle.
b) Deduce that a strong digraph which contains an odd cycle contains a directed odd cycle.

$\star$**3.4.12** Show that:

a) every nontrivial strong tournament has a directed Hamilton cycle,
(P. Camion)

b) each vertex of a nontrivial strong tournament D is contained in a directed cycle of every length l, $3 \leq l \leq n$, (J.W. Moon)
c) each arc of an even tournament D is contained in a directed cycle of every length l, $3 \leq l \leq n$. (B. Alspach)

3.5 Cycle Double Covers

In this section, we discuss a beautiful conjecture concerning cycle coverings of graphs. In order for a graph to admit a cycle covering, each of its edges must certainly lie in some cycle. On the other hand, once this requirement is fulfilled, the set of all cycles of the graph clearly constitutes a covering. Thus, by Proposition 3.2, a graph admits a cycle covering if and only if it has no cut edge. We are interested here in cycle coverings which cover no edge too many times.

Recall that a *decomposition* is a covering in which each edge is covered exactly once. According to Veblen's Theorem (2.7), the only graphs which admit such cycle coverings are the even graphs. Thus, if a graph has vertices of odd degree, some edges will necessarily be covered more than once in a cycle covering. One is led to ask whether every graph without cut edges admits a cycle covering in which no edge is covered more than twice.

All the known evidence suggests that this is indeed so. For example, each of the platonic graphs (shown in Figure 1.14) has such a cycle covering consisting of its *facial cycles*, those which bound its regions, or *faces*, as in Figure 3.8. More generally, the same is true of all polyhedral graphs, and indeed of all planar graphs without cut edges, as we show in Chapter 10.

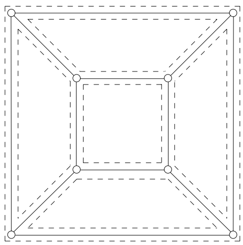

Fig. 3.8. A double covering of the cube by its facial cycles

In the example of Figure 3.8, observe that any five of the six facial cycles already constitute a cycle covering. Indeed, the covering shown, consisting of all six facial cycles, covers each edge exactly twice. Such a covering is called a *cycle double cover* of the graph. It turns out that cycle coverings and cycle double covers are closely related.

Proposition 3.8 *If a graph has a cycle covering in which each edge is covered at most twice, then it has a cycle double cover.*

Proof Let $\mathcal{C}$ be a cycle covering of a graph G in which each edge is covered at most twice. The symmetric difference $\triangle\{E(C) | C \in \mathcal{C}\}$ of the edge sets of the cycles in $\mathcal{C}$ is then the set of edges of G which are covered just once by $\mathcal{C}$. Moreover, by Corollary 2.16, this set of edges is an even subgraph C' of G. By Veblen's Theorem (2.7), C' has a cycle decomposition $\mathcal{C}'$. It is now easily checked that $\mathcal{C} \cup \mathcal{C}'$ is a cycle double cover of G. □

Motivated by quite different considerations, Szekeres (1973) and Seymour (1979b) each put forward the conjecture that every graph without cut edges admits a cycle double cover.

The Cycle Double Cover Conjecture

Conjecture 3.9 *Every graph without cut edges has a cycle double cover.*

A graph has a cycle double cover if and only if each of its components has one. Thus, in order to prove the Cycle Double Cover Conjecture, it is enough to prove it for nontrivial connected graphs. Indeed, one may restrict one's attention even further, to *nonseparable graphs*. Roughly speaking, these are the connected graphs which cannot be obtained by piecing together two smaller connected graphs at a single vertex. (Nonseparable graphs are defined and discussed in Chapter 5.) In the case of planar graphs, the boundaries of the faces in any planar embedding are then cycles, as we show in Chapter 10, and these facial cycles constitute a cycle double cover of the graph. This suggests one natural approach to the Cycle Double Cover Conjecture: find a suitable embedding of the graph on some surface, an embedding in which each face is bounded by a cycle; the facial cycles then form a cycle double cover.

Consider, for example, the toroidal embeddings of the complete graph K_7 and the Petersen graph shown in Figure 3.9. The torus is represented here by a rectangle whose opposite sides are identified; identifying one pair of sides yields a cylinder, and identifying the two open ends of the cylinder results in a torus. In the embedding of K_7, there are fourteen faces, each bounded by a triangle; these triangles form a cycle double cover of K_7. In the embedding of the Petersen graph, there are five faces; three are bounded by cycles of length five (faces A, B, C), one by a cycle of length six (face D), and one by a cycle of length nine (face E). These five cycles constitute a cycle double cover of the Petersen graph.

The above approach to the Cycle Double Cover Conjecture, via surface embeddings, is supported by the following conjecture, which asserts that every loopless nonseparable graph can indeed be embedded in some surface in an appropriate fashion.

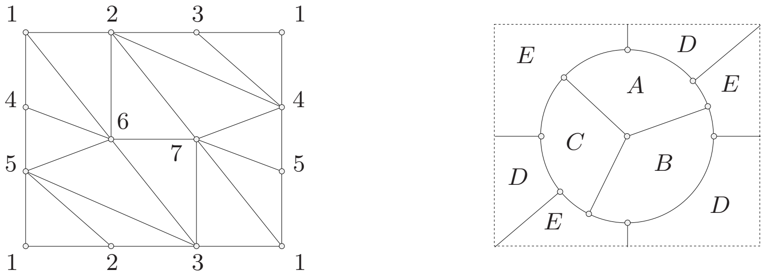

Fig. 3.9. Toroidal embeddings of (a) the complete graph K_7, and (b) the Petersen graph

THE CIRCULAR EMBEDDING CONJECTURE

Conjecture 3.10 *Every loopless nonseparable graph can be embedded in some surface in such a way that each face in the embedding is bounded by a cycle.*

The origins of Conjecture 3.10 are uncertain. It was mentioned by W.T. Tutte (unpublished) in the mid-1960s, but was apparently already known at the time to several other graph-theorists, according to Robertson (2007). We discuss surface embeddings of graphs in greater detail in Chapter 10, and describe there a stronger conjecture on embeddings of graphs.

Apart from its intrinsic beauty, due to the simplicity of its statement and the fact that it applies to essentially all graphs, the Cycle Double Cover Conjecture is of interest because it is closely related to a number of other basic problems in graph theory, including the Circular Embedding Conjecture. We encounter several more in future chapters.

DOUBLE COVERS BY EVEN SUBGRAPHS

There is another attractive formulation of the Cycle Double Cover Conjecture, in terms of even subgraphs; here, by an even subgraph we mean the edge set of such a subgraph.

If a graph has a cycle covering, then it has a covering by even subgraphs because cycles are even subgraphs. Conversely, by virtue of Theorem 2.17, any covering by even subgraphs can be converted into a cycle covering by simply decomposing each even subgraph into cycles. It follows that a graph has a cycle double cover if and only if it has a double cover by even subgraphs. Coverings by even subgraphs therefore provide an alternative approach to the Cycle Double Cover Conjecture. If every graph without cut edges had a covering by at most two even subgraphs, such a covering would yield a cycle covering in which each edge was covered at most twice, thereby establishing the Cycle Double Cover Conjecture by virtue of Proposition 3.8. Unfortunately, this is not the case. Although many graphs do indeed admit such coverings, many do not. The Petersen graph, for instance, cannot be covered by two even subgraphs (Exercise 3.5.3a). On the other hand, it may be shown that every graph without cut edges admits a covering by three even subgraphs (Theorem 21.21).

Suppose, now, that every graph without cut edges does indeed have a cycle double cover. It is then natural to ask how few cycles there can be in such a covering; a covering with few cycles may be thought of as an efficient covering, in some sense. Let $\mathcal{C}$ be a cycle double cover of a graph G. As each edge of G is covered exactly twice,

$$\sum_{C \in \mathcal{C}} e(C) = 2m$$

Because $e(C) \leq n$ for all $C \in \mathcal{C}$, we deduce that $|\mathcal{C}| \geq 2m/n$, the average degree of G. In particular, if G is a complete graph K_n, the number of cycles in a cycle double cover of G must be at least $n - 1$. A cycle double cover consisting of no more than this number of cycles is called a *small cycle double cover*. Bondy (1990) conjectures that every simple graph G without cut edges admits such a covering.

Conjecture 3.11 The Small Cycle Double Cover Conjecture
Every simple graph without cut edges has a small cycle double cover.

Several other strengthenings of the Cycle Double Cover Conjecture have been proposed. One of these is a conjecture put forward by Jaeger (1988).

Conjecture 3.12 The Oriented Cycle Double Cover Conjecture
Let G be a graph without cut edges. Then the associated digraph $D(G)$ of G admits a decomposition into directed cycles of length at least three.

Further information on these and a number of related conjectures can be found in the book by Zhang (1997).

Exercises

3.5.1 Show that every loopless graph has a double covering by bonds.

3.5.2 Let $\{C_1, C_2, C_3\}$ be a covering of a graph G by three even subgraphs such that $C_1 \cap C_2 \cap C_3 = \emptyset$. Show that $\{C_1 \triangle C_2,\ C_1 \triangle C_3\}$ is a covering of G by two even subgraphs.

$\star$**3.5.3**

a) Show that the Petersen graph has no covering by two even subgraphs.
b) Deduce, using Exercise 3.5.2, that this graph has no double cover by four even subgraphs.
c) Find a covering of the Petersen graph by three even subgraphs, and a double cover by five even subgraphs.

3.5.4

a) i) Let $\{C_1, C_2\}$ be a covering of a graph G by two even subgraphs. Show that $\{C_1, C_2, C_1 \triangle C_2\}$ is a double cover of G by three even subgraphs.
   ii) Deduce that a graph has a covering by two even subgraphs if and only if it has a double cover by three even subgraphs.

.

b) Let $\{C_1, C_2, C_3\}$ be a covering of a graph G by three even subgraphs. Show that G has a quadruple cover (a covering in which each edge is covered exactly four times) by seven even subgraphs.
(We show in Theorem 21.25 that every graph without cut edges has a covering by three even subgraphs, and hence a quadruple cover by seven even subgraphs.)

3.5.5 Find a small cycle double cover of K_6.

3.5.6 Find a decomposition of $D(K_6)$ into directed cycles of length at least three.

3.5.7 Show that every graph without cut edges has a uniform cycle covering.

3.5.8 Let G be a graph, and let $\mathcal{C}$ be the set of all cycles of G. For $C \in \mathcal{C}$, denote by $\mathbf{f}_C$ the incidence vector of C, and set $\mathbf{F}_{\mathcal{C}} := \{\mathbf{f}_C : C \in \mathcal{C}\}$.

a) Let $\mathbf{x} \in \mathbb{R}^E$. Show that:
   i) the vector $\mathbf{x}$ lies in the vector space generated by $\mathbf{F}_{\mathcal{C}}$ if and only if the following two conditions hold:
      ▷ $x(e) = 0$ for every cut edge e,
      ▷ $x(e) = x(f)$ for every edge cut $\{e, f\}$ of cardinality two,
   ii) if $\mathbf{x}$ is a nonnegative linear combination of vectors in $\mathbf{F}_{\mathcal{C}}$, then for any bond B of G and any edge e of B:

$$x(e) \leq \sum_{f \in B \setminus \{e\}} x(f) \tag{3.1}$$

   (Seymour (1979b) showed that this necessary condition is also sufficient for a nonnegative vector $\mathbf{x}$ to be a nonnegative linear combination of vectors in $\mathbf{F}_{\mathcal{C}}$.)
   iii) if $\mathbf{x}$ is a nonnegative integer linear combination of vectors in $\mathbf{F}_{\mathcal{C}}$, then for any bond B, in addition to (3.1), $\mathbf{x}$ must satisfy the condition:

$$\sum_{e \in B} x(e) \equiv 0 \pmod 2 \tag{3.2}$$

b) With the aid of Exercise 2.4.6, give an example showing that conditions (3.1) and (3.2) are not sufficient for a nonnegative integer vector $\mathbf{x}$ in $\mathbb{R}^E$ to be a nonnegative integer linear combination of vectors in $\mathbf{F}_{\mathcal{C}}$.

(Seymour (1979b) showed, however, that these two conditions are sufficient when G is a planar graph. Furthermore, he conjectured that they are sufficient in any graph if each component of $\mathbf{x}$ is an even integer. This conjecture clearly implies the Cycle Double Cover Conjecture. For related work, see Alspach et al. (1994).)

3.6 Related Reading

Cages

Cages were introduced in Exercise 3.1.13. There are many interesting examples of such graphs, the Petersen graph and the Heawood graph being but two. Numerous others are described in the survey by Wong (1982). Two particularly interesting infinite families of examples are those constructed from projective geometries by Benson (1966), namely the $(k, 8)$- and $(k, 12)$-cages, where $k-1$ is a prime power. For $\ell = 3, 5$, the Benson cages furnish examples of dense graphs (graphs with many edges) containing no 2ℓ-cycles. For $\ell = 2$, examples are provided by polarity graphs of projective planes (see Exercises 12.2.12, 12.2.13, and 12.2.14.) The question as to how many edges a graph on n vertices can have without containing a 2ℓ-cycle is unsolved for other values of ℓ, and in particular for $\ell = 4$; see Appendix A.

The study of *directed cages*, smallest k-diregular digraphs with specified directed girth g, was initiated by Behzad et al. (1970). They conjectured that the directed circulants on $k(g-1)+1$ vertices in which each vertex dominates the k vertices succeeding it are directed cages. This conjecture remains open; see Appendix A.