Notes for Recitation 6

1 Graph Basics

Let $G = (V, E)$ be a graph. Here is a picture of a graph.

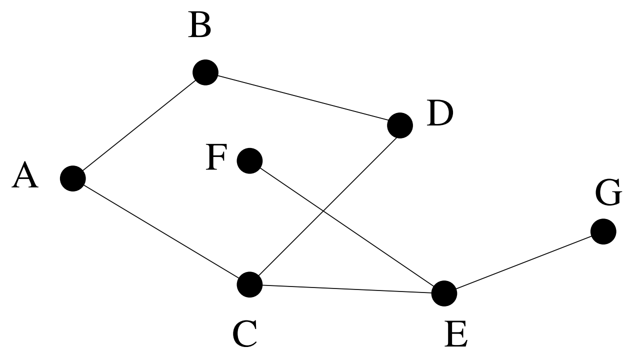

Recall that the elements of V are called vertices, and those of E are called edges. In this example the vertices are $\{A, B, C, D, E, F, G\}$ and the edges are

$$\{A—B, B—D, C—D, A—C, E—F, C—E, E—G\}.$$

Deleting some vertices or edges from a graph leaves a *subgraph*. Formally, a subgraph of $G = (V, E)$ is a graph $G' = (V', E')$ where V' is a nonempty subset of V and E' is a subset of E. Since a subgraph is itself a graph, the endpoints of every edge in E' must be vertices in V'. For example, $V' = \{A, B, C, D\}$ and $E' = \{A—B, B—D, C—D, A—C\}$ forms a subgraph of G.

In the special case where we only remove edges incident to removed nodes, we say that G' is the *subgraph induced on* V' if $E' = \{(x—y | x, y \in V'$ and $x—y \in E\}$. In other words, we keep all edges unless they are incident to a node not in V'. For instance, for a new set of vertices $V' = \{A, B, C, D\}$, the induced subgraph G' has the set of edges $E' = \{A—B, B—D, C—D, A—C\}$.

Definition 1. *Let $G = (V, E)$ be a graph. A* path *in G is a sequence of vertices*

$$v_0, \ldots, v_k$$

with $k \geq 0$ such that $v_i—v_{i+1}$ is an edge in E for all $i \geq 0$ such that $i < k$, and all the v_i's, except possibly v_0 and v_k, are different. That is, if $0 \leq i < j \leq k$, then $v_i = v_j$ only if both

$i = 0$ and $j = k$. The path is said to start *at v_0, to* end *at v_k, and* length *of the path is defined to be k. If it happens that $v_0 = v_k$, then we say the path is* closed. *In this case we call the path a* cycle.

For example, the graph in the figure above has a length 5 path A,B,D,C,E,G. There is a cycle of length 4 – namely A,B,D,C,A.

The *distance* between two vertices is just the length of the shortest path between them. So in the figure above, the distance $d(A, D)$ between A and D is 2.

Definition 2. *A* walk *in a graph G is an alternating sequence of vertices and edges of the form:*

$$v_0, v_0 — v_1, v_1, v_1 — v_2, v_2, \ldots, v_{n-1}, v_{n-1} — v_n, v_n$$

If $v_0 = v_n$, then the walk is closed. In this case we say the walk is a circuit.

Walks are similar to paths. However, a walk can cross itself, traverse the same edge multiple times, etc. There is a walk between two vertices if and only if there is a path between the vertices. The length of a walk is defined in the same way as the length for a path. It is just the number of edges on the walk. So, in the above graph A, C, D, C is a walk, but not a path. Moreover, A, C, D, C, A is a circuit.

We then say that a graph $G = (V, E)$ is *connected* if for all pairs of nodes $v_i, v_j \in V$, there is a path from v_i to v_j in G.

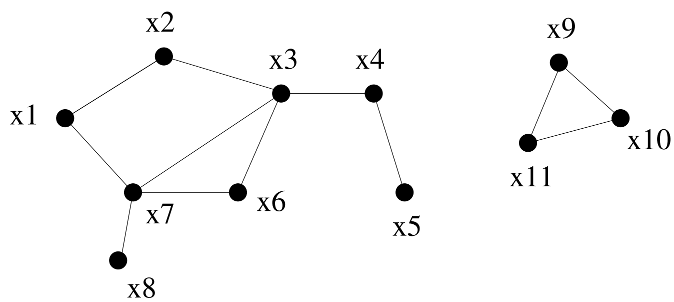

The graph shown here is not connected since, for example, there is no path from x_5 to x_9. But it does have *two* connected components – that is, it consists of two subgraphs that are themselves connected.

A *connected component* of G is a maximal connected subgraph of G. Maximal means that you can't add any nodes or edges to the subgraph without making it be disconnected. For example, x_9, x_{10}, x_{11} is a connected component above. But, $x_1, \ldots, x_7$ is not, since you can add x_8 and still stay connected.

2 Bipartite Graphs

Now we're going to talk about *coloring* a graph, which means assigning a color to each of its vertices. We will look at the very special case when we only have two colors to work with, say, black and white. A graph G is *bipartite* if each vertex can be assigned one of the two colors so that the endpoints of each edge in the graph are colored differently.

Then we can put all the black vertices in a clump on the left and all the white vertices in a clump on the right. Since every edge joins differently-colored vertices, every edge must run between the two clumps. Therefore, every bipartite graph looks something like this:

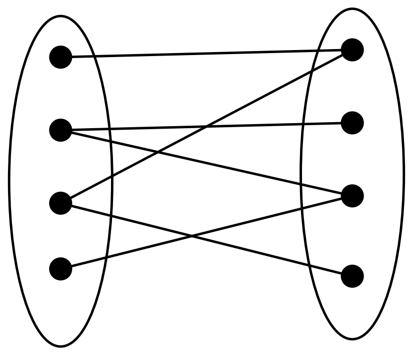

Bipartite graphs are both useful and common. In this problem, you will prove:

Theorem. *A simple graph G is 2-colorable iff it contains no odd length cycle.*

As usual with "iff" assertions, the proof splits into two parts: part (1) asks you to prove that the left side of the "iff" implies the right side. The other problem parts (2–6) prove that the right side implies the left.

1. Assume the left side and prove the right side. Three to five sentences should suffice.

   Solution. First, we assume that G is 2-colorable and prove that G contains no odd length cycle.

   Select a 2-coloring of G. Consider an arbitrary cycle with successive vertices $v_1, v_2, \ldots,$ v_k, v_1. Then the vertices v_i must be one color for all even i and the other color for all odd i. (We could confirm this claim with a proof by induction, but it seems obvious enough to accept without further proof.) Since v_1 and v_k must be colored differently, k must be even. Thus, the cycle has even length. We can make the same argument for any cycle in G, so every cycle has even length. ■

2. Now we will prove that if G contains no odd length cycle, then it is 2-colorable. As a first step toward proving this, explain why we can focus on a single connected component H within G.

   Solution. If we can 2-color every connected component of G, then we can 2-color all of G since vertices in different connected components have no edges in between them. Thus, it suffices to show that an arbitrary connected component H of G is 2-colorable. ■

3. Take an arbitrary connected component H of G. Let v be an arbitrary vertex in H. Suppose we color all vertices at an even distance from v the same color, and all vertices at an odd distance from v another color. Explain why either H is bipartite, or we can find two vertices a and b and an edge (a, b) in H, such that $d(v, a) \equiv d(v, b) \bmod 2$.

   Solution. If H is bipartite, we are done. Otherwise there is no valid coloring of the vertices of H, and so the coloring that we've chosen must have an edge between two vertices of the same color. This means that we can find an edge (a, b) with $d(v, a) \equiv d(v, b) \bmod 2$. ■

4. Now we will prove that H is 2-colorable by proving the contrapositive. Namely, if H is not 2-colorable, then it contains an odd-length cycle. Here we show in the second case from part 3, that we can find an odd cycle. Assume we've found a and b as above. Let $P(v, a)$ and $P(v, b)$ be *shortest* paths from v to a and from v to b, respectively. Explain why the walk $P(v, a)$ reverse$(P(v, b))$ is a circuit of odd length, where reverse$(v_1, \ldots, v_k)$ is defined to be $v_k, \ldots, v_1$.

   Solution. The walk starts and ends at the vertex v, so by definition it is a circuit. Its length is the length of $P(v, a)$ plus that of $P(v, b)$, plus 1 to account for the edge (a, b). Since $P(v, a)$ and $P(v, b)$ are shortest paths, their lengths are $d(v, a)$ and $d(v, b)$, respectively. It follows that modulo 2, the length of the walk is $d(v, a) + d(v, b) + 1 \equiv 1 \bmod 2$, since $d(v, a) \equiv d(v, b) \bmod 2$. ■

5. Show that any odd-length circuit contains an odd-length cycle.

   Solution. The proof is by strong induction on the number of edges in the circuit. Note that any odd length circuit must have length at least 3, since there is no circuit of length 1. Let $P(r)$ be the proposition that any circuit of length r contains an odd-length cycle, for any odd integer r such that $r \geq 3$.

   Base Case ($r = 3$): If the circuit has length 3, then it has the form a, b, c, a for some distinct vertices a, b and c. This follows from the fact that the graph is simple (no multiple edges or self loops). But in this case the circuit is a cycle of length 3.

   Inductive Step: Assume $P(k), 3 \leq k \leq r, k$ odd, show $P(r+2)$. Suppose, inductively, that all odd-length circuits of size at most r contain an odd-length cycle, for some odd integer $r \geq 3$. We will prove that any circuit of length $r + 2$ contains an odd-length cycle. Consider a circuit of the form $v_0, \ldots, v_{r+2}$, which has length $r + 2$. If it is in fact a cycle, we are done. Otherwise there is some vertex w which occurs more than once in the circuit. Suppose it occurs at v_i and v_j for $i < j$. Then the walks $v_0, \ldots, v_i, v_{j+1}, \ldots, v_{r+2}$ and $v_i, v_{i+1}, \ldots, v_j$ are both circuits. The sum of their lengths is the length of the original circuit $v_0, \ldots, v_{r+2}$ (i.e., $r + 2$), which is odd. Therefore, one of the walks' lengths must also be odd and strictly less than $r + 2$. By the inductive hypothesis this circuit contains an odd cycle, and thus so does $v_0, \ldots, v_{r+2}$. ■

6. We see that in this case, H contains an odd length cycle. We proved that if H is not bipartite, then it contains an odd-length cycle. The contrapositive of this is that if H contains no odd-length cycle, then it is bipartite.

3 Euler Circuits

Now we'll consider a special kind of circuit called an *Euler circuit*, named after the famous mathematician Leonhard Euler (the same Euler from Euler's Theorem).

Euler considered circuits of graphs of a very special form.

Definition 3. *An Euler circuit of a graph G is a circuit of G which visits every edge exactly once.*

Does the graph in the figure of Section 1 contain an Euler circuit? Well, if it did, the edge (E, F) would need to be included. If the walk does not start at F then at some point it traverses edge (E, F), and now it is stuck at F since F has no other edges incident to it and an Euler circuit can't traverse (E, F) twice. But then the walk could not be a circuit. On the other hand, if the walk starts at F, it must then go to E along (E, F), but now it cannot return to F. It again cannot be a circuit. This argument generalizes to show that if a graph has a vertex of degree 1, it cannot contain an Euler circuit.

It is easy to see that any cycle contains an Euler circuit. You can just start at any vertex and walk around back to it.

Naturally, this leads us to the question of which graphs contain Euler circuits. At first glance this may seem like a daunting problem. Nevertheless, you will now completely solve it. Recall that the *degree* of a vertex is the number of edges adjacent to it.

1. Show that if a graph G has an Euler circuit, then the degree of each of its vertices is even.

   Solution. Consider any Euler circuit $C = v_1, v_2, \ldots, v_r, v_1$ of G. Consider any vertex v. Then every time v occurs in C, there is a vertex a which comes immediately before v and a vertex b which comes immediately after v. Note that this holds for $v = v_1$ as well since C is a circuit. Moreover, (a, v) and (v, b) must be distinct edges of G since C is an Euler circuit. It follows that if v occurs s times in C, then it has degree $2s$ since every edge incident to v occurs in C exactly once. Thus, v has even degree. ■

2. We will now show that if the degree of each of the vertices of a connected graph G is even, then G has an Euler circuit. To do this, consider a longest walk $W = v_0, v_1, \ldots, v_r$ in G using no edge more than once. Show that W is a circuit.

   Solution. First, observe that all of the edges incident to v_r must already occur in W, as otherwise we could extend W to a walk with one more edge, while preserving the property that each edge in G occurs at most once in W. It follows that if v_r is different from v_0, it must have odd degree. To see why, remember that each of its edges occurs exactly once in W, and that therefore, if v_r were to appear k times before its last occurrence, its degree would be $2k + 1$. But v_r having odd degree contradicts the hypothesis that the degree of each of the vertices is even. It follows that $v_r = v_0$, and therefore W is a circuit. ■

3. Suppose that W is not an Euler circuit. Show that this implies it cannot be a longest walk in which each edge is visited at most once. Note that this yields a contradiction to our choice of W.

   Solution. By construction, W uses no edge of G more than once. If W is not an Euler circuit, by the definition of an Euler circuit, there is some edge in G which does not occur in W. Since G is connected, let this edge be $e = (v, v_i)$, incident to some vertex v_i in W. (If there were no such edge, then G would be disconnected.) Consider the walk

   $$W' = v, v_i, v_{i-1}, \ldots, v_0, v_{r-1}, v_{r-2}, \ldots, v_{i+1}, v_i$$

   This walk uses each edge at most once and is of length longer than W, which contradicts that W is a longest walk in G using each edge at most once. ■