1 Bipartite graphs

One interesting class of graphs rather akin to trees and acyclic graphs is the *bipartite graph*:

Definition 1. A graph G is *bipartite* if the vertex-set of G can be partitioned into two sets A and B such that if u and v are in the same set, u and v are non-adjacent.

We've seen one good example of these already: the complete bipartite graph $K_{a,b}$ is a bipartite graph in which every possible edge *between* the two sets exists. We've in fact also seen several other bipartite graphs.

Proposition 1. *Every tree (or forest) is a bipartite graph.*

Proof. We may use the structural theorem on trees to prove this result fairly easily by induction. Clearly when $|T| = 1$, T is bipartite (although one of the parts will be empty).

For the inductive step, let $|T| = n$, and let v be a leaf; $T - v$ is a tree by the structural theorem, and by the inductive hypothesis, $T - v$ is bipartite. Thus, all vertices of T except for v may easily be classified into partitions satisfying the non-adjacency criterion. Since $d_T(v) = 1$, v has one neighbor. Put v in the partition not occupied by its neighbor to satisfy the non-adjacency condition.

To extend this result to forests, partition each component tree of the forest individually, and let the parts of the overall graph just be the union of the parts from each component. □

It's pretty easy to show that even cycles are also bipartite, but that odd cycles are not. In fact, this observiation, slightly generalized, forms the entire criterion for a graph to be bipartite.

Theorem 1. *A graph G is bipartite if and only if it contains no odd cycles.*

Proof. First, let us show that if a graph contains an odd cycle it is not bipartite. Let $v_1 \sim v_2 \sim \cdots \sim v_{2n-1} \sim v_1$ be the vertices of an odd cycle in G. If G were bipartite, then v_1 would be in some part; without loss of generality we may say $v_1 \in A$, so $v_2 \in B$ since it is adjacent to v_1, and $v_3 \in A$ since it is adjacent to v_2, and so forth up to showing that $v_{2n-1} \in A$. But then v_1 and v_{2n-1} would be adjacent vertices in a single part, violating the partition conditions, so such a partition cannot be valid, so G is non-bipartite.

Now, to show the converse, we will, for simplicity, look at a connected graph G which lacks odd cycles. If G is not connected, we could demonstrate bipartiteness on each component, and, as in the previous proof, build an overall partition as a union of the partitions on each component.

Now, we will make use of the structural theorem on connected graphs: we order the vertices of G with $v_1, v_2, \ldots, v_n$ such that each $G[v_1, v_2, \ldots, v_k]$ is connected. Then, we will step through the vertices, assigning them in order to the parts A and B, following the procedure below:

- Arbitrarily assign v_1 to part A.
- For each v_k from 2 through n, assign them in sequence. By the structural theorem, since $G[v_1, \ldots, v_k]$ is connected, v_k has at least one neighbor in $v_1, \ldots, v_{k-1}$; these vertices have already been assigned to parts. If all of them are in part A, assign v_k to part B, and vice versa. These two possibilities guarantee that the partition satisfies the non-adjacency criterion on $G[v_1, \ldots, v_k]$. The one situation not dealt with is when v_k has neighobrs in both parts, which we shall show below is impossible.

If the above procedure can be followed, it is guaranteed to assign all vertices of G to parts with no adjacencies among vertices inside a single part. The only danger is the assumption that all of v_k's neighbors lie in a single part.

Let us suppose, to the contrary, that for some $i, j < k$, $v_i \in A$ and $v_j \in B$, and v_k is adjacent to both v_i and v_j. By connectivity of $G[v_1, \dots, v_{k-1}]$, there is a path from v_i to v_j in $G[v_1, \dots, v_{k-1}]$. Since all the vertices $v_1, \dots, v_{k-1}$ have been assigned to a part, the aforementioned path $v_i = u_1 \sim u_2 \sim u_3 \sim \cdots \sim u_r = v_k$ consists of vertices already assigned to A or B. We know $u_1 = v_i \in A$; thus by adjacency, $u_2 \in B$, and $u_3 \in A$, and so forth, with odd indices in A and even indices in B. Since $u_r = v_k \in B$, it thus follows that r is even. But then $v_k \sim u_1 \sim u_2 \sim \cdots \sim u_r \sim v_k$ is an odd cycle, which we were guaranteed not to find in G. □

Bipartite groups will become significant when we look at several other problems later, including matchings and colorings.

2 Path-determined properties

2.1 Distance

The concepts of paths, connectedness, etc. have several ancillary properties which deserve development on their own. One significant property is the "shortest path" between two points.

Definition 2. The *length* of a path is the number of edges appearing in the path (or, alternatively, one less than the nubmer of vertices). If G is a connected graph, then the *distance* between two vertices u and v of G, denoted $d_G(u, v)$, is the minimum length among all paths from u to v.

There is an occasionally-seen, but by no means universal, convention that extends the concept of distance to disconnected graphs, stating that if u and v are in different components, then $d_G(u, v) = \infty$.

There are a couple of easy-to-show facts hardly worth mentioning, and not worth proving in detail:

- $d_G(u, v) = 0$ if and only if $u = v$.
- $d_G(u, v) = 1$ if and only if u is adjacent to v.
- $d_G(u, v) = d_G(v, u)$.

However, there are also several more fairly intuitive facts sufficiently nontrivial to deserve proof.

Proposition 2. *For any vertices u, v, and w in a graph G, $d(u, v) + d(v, w) \geq d(u, w)$.*

Proof. Definitionally, there is a path from u to v of length $d(u, v)$ and a path from v to w of length $d(v, w)$. Splicing these end-to-end gives a walk of length $d(u, v) + d(v, w)$ from u to w. As seen last semester, a path can be produced by excising cycles and backtracks from a walk; doing so will only reduce its length. Thus there is a path from u to w of length less than $d(u, v) + d(v, w)$. The minimum path length is clearly no more than the length of this particular path, so $d(u, w) \leq d(u, v) + d(v, w)$. □

Proposition 3. *If a connected graph H is a subgraph of some connected G, then for any $u, v \in V(H)$, $d_G(u, v) \leq d_H(u, v)$.*

Proof. Since every vertex and edge in H is also in G, any path in H is also a path in G. By the definition of distance, H contains a path of length $d_H(u, v)$ between u and v; this path is also in G. Since G contains a path of length $d_H(u, v)$ between u and v, it follows that the shortest path in G between them is no longer than $d_H(u, v)$; thus $d_G(u, v) \leq d_H(u, v)$.

If we adopt the infinite-distance convention mentioned above, we can relax the connectedness condition in the proof, so that G and H need not be connected. □

The shortest distance between two points in a graph is easy to find; what is more interesting, as a measure of a graph's size, is the distance between very distant points.

Definition 3. The *diameter* of a connected graph G is $d(G) = \max_{u,v\in V(G)} d_G(u, v)$. The *radius* of G is $r(G) = \min_{u\in V(G)} \max_{v\in V(G)} d_G(u, v)$.

The diameter is simply the distance between the two furthest apart points in the graph. The radius is a bit more complicated: it's the distance from a point chosen so as to be as central as possible to the furthest away point (such a point is sometimes called the *center*) of the graph.

The diameter is a popular measure of graph size; the radius less so. The diameter alone is guaranteed to give a pretty good idea of graph magnitude, since the radius is closely linked to it.

Proposition 4. *For any connected graph G, $r(G) \leq d(G) \leq 2r(G)$.*

Proof. It is trivial to see that $r(G) \leq d(G)$, since

$$r(G) = \min_{u\in V(G)} \max_{v\in V(G)} d(u, v) \leq \max_{u\in V(G)} \max_{v\in V(G)} d(u, v) = d(G)$$

To show the other inequality, let u and v be points such that $d(u, v) = d(G)$; let w be a point such that $\max_{x\in V(G)} d(w, x) = r(G)$. These points are guaranteed to exist by the definitions of radius and diameter. Since $r(G) = \max_{x\in V(G)} d(w, x)$, we know specifically that $d(w, u)$ and $d(w, v)$ are no more than $r(G)$. Thus, $d(G) = d(u, v) \leq d(u, w) + d(w, v) = d(w, u) + d(w, v) \leq r(G) + r(G) = 2r(G)$. □

Note that both extremes in the above proposition are achievable. In a graph where each vertex is essentially identical, like C_n or K_n, $d(G) = r(G)$; on the other hand, the path P_{2n+1} has diameter $2n$, since that is the distance between the two ends of the path, but radius n, since the point in the middle of the path has distance n from the ends.

2.2 Path-contraction substructures

There are several ways one graph can be a substructure of another graph beyond subgraphs and induced subgraphs. These are path-contraction structures, where edges and paths are reduced.

Definition 4. If x and y are adjacent vertices of a graph G, then the *edge contraction* G/xy is a graph on the vertex set $V(G) \setminus \{x, y\} \cup \{xy\}$ such that edges in G among vertices other than x and y are replicated in G/xy, and the edge $\{u, xy\}$ is in G/xy if and only if u is adjacent to x or y in G.

As an example, here we see the result of an edge-contraction on the labeled vertices x and y:

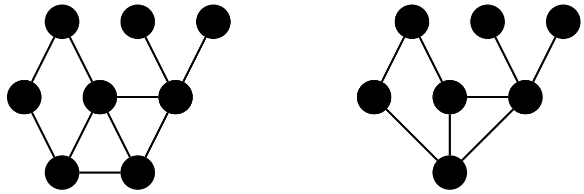

Definition 5. If H is a subgraph of the graph that results from some sequence of edge contractions of G, then H is said to be a *minor* of H.

It's not necessarily clear why this characterization is useful, but it turns out to be a rich way of identifying graph types. The following extremely simple proposition is a very simple case of a large and unsolved problem in graph theory (Hadwiger's Conjecture, for those who wish to investigate further):

Proposition 5. *If G is not bipartite, then G has a K_3 minor.*

Proof. A non-bipartite graph contains an odd cycle $v_1 \sim v_2 \sim \cdots \sim v_{2n+1} \sim v_1$. By successively contracting the edges $\{v_3, v_4\}, \{v_3v_4, v_5\}, \{v_3v_4v_5, v_6\}, \ldots, \{v_3v_4v_5 \ldots v_{2n}, v_{2n+1}\}$, we may contract this cycle down to the K_3 on the vertices v_1, v_2, and $v_3v_4v_5 \cdots v_{2n+1}$. Thus, this contraction of G has a K_3 subgraph, so K_3 is a minor of G. □

Note that the converse is not true: C_4 has a K_3 minor, but is bipartite.

Another concept of contraction is treating paths as if they were edges, which gives what is frequently called "topological structure" of the graph.

Definition 6. A graph H is a *topological minor* of a graph G if the vertices of H can be mapped injectively onto the vertices of G such that there are nonintersecting paths $P_1, P_2, \ldots, P_m$ in G corresponding respectively to each edge $e_1, \ldots, e_m$ in H such that the endpoints of each path are the images under the vertex map of the endpoints of the associated edge.

Both of the above are clearly generalizations of the concept of a subgraph: in fact, every subgraph of G is trivially both a minor and a topological minor, since a subgraph could be considered a minor by opting to do no contractions, or a topological minor by considering the simple length-one paths which are identical to the edges in the subgraph.

Furthermore, one can see that every topological minor is a minor. If H is a topological minor of a graph G, then H can be found as a subgraph of a contraction of G, since by contracting all edges of a path, it can be reduced to an edge; doing so on all paths used in the topological minor will yield H itself. The converse is not true: note, for instance, that a cycle together with a vertex adjacent to all vertices of the cycle has a K_4 minor but no topological K_4 minor.

In general, the forms of substructures one can find in a graph can be ranked as follows in descending order of strictness:

$$\{\text{induced subgraphs of } G\} \subseteq \{\text{subgraphs of } G\} \subseteq \{\text{topological minors of } G\} \subseteq \{\text{minors of } G\}$$

2.3 Eulerian and Hamiltonian traversals

Another question arising from investigations into paths (and other, more general walks) is their ability to traverse the entire graph with minimum self-intersection. There are two major questions, one asked about paths (and cycles) and the other about trails (and tours):

Definition 7. A path or cycle in a graph G which contains every vertex of the graph is called *hamiltonian*; a trail or tour which contains every edge of the graph is called *eulerian*. A graph G is *hamiltonian* or *eulerian* if it has a hamiltonian cycle or an eulerian tour respectively.

Note that hamiltonicity could also be defined in terms of subgraphs: for $|G| = n$, G is hamiltonian if it has a C_n subgraph. Hamiltonicity is a complicated subject with a great deal of open-ended questions associated with its determination; by way of contrast, the question of finding eulerian tours on graphs has been definitively answered for nearly four centuries.

A historical aside: the first work definitively identified as graph-theoretical is Euler's 1736 *Soluio problematis ad geometriam situs pertinentis*, which described its field, rather imaginatively, as, "the geometry of position". The question under consideration was whether a tour through the city of Königsberg could make use of every bridge exactly once: in fact, considering the shores and islands as vertices and the bridges as edges, it can be observed that the city is not eulerian and no such tour is possible. Euler's main line of argument is generalized in the proof below. His actual argument is slightly more complicated, as it applies to *multigraphs*, structures in which the same pair of vertices may be connected by multiple edges.

Theorem 2. *A simple finite graph is eulerian if and only if every vertex has even degree and all of the edges are in a single component.*

Proof. It is clearly necessary that all edges lie in a single component in order for a trail to traverse all of them (note that connectivity of G is not strictly necessary; isolated vertices are not a problem, from an eulerian-tour perspective).

Let us start by supposing G has an eulerian tour: we shall show that all vertices of G have even degree. Let us denote our eulerian tour by $v_1 \sim v_2 \sim v_3 \sim v_4 \sim \cdots \sim v_m \sim v_1$, and let $e_i = \{v_i, v_{i+1}\}$, with the special case $e_m = \{v_m, v_1\}$. Then, by the definition of an eulerian tour, all the e_i are distinct and $E(G) = \{e_1, e_2, \ldots, e_m\}$. Now, consider a vertex u in G. Some number k (which might be zero) of the v_i are equal to u: let us denote $u = v_{i_1} = v_{i_2} = \cdots = v_{i_k}$. Then, for any j, note that $v_{i_j \pm 1}$ (using the wraparound convention that $v_0 = v_k$ and $v_{k+1} = v_1$ where applicable) must *not* be identical to u, since $v_{i_j \pm 1}$ is adjacent to u and G is a simple graph. Thus, since none of the i_j are consecutive, the indices $\{i_1 - 1, i_i, i_2 - 1, i_2, \ldots, i_k - 1, i_k\}$ are all distinct. Since the edge e_i is incident on u if and only if either $i = i_j$ or $i + 1 = i_j$, then the above set of indices is exactly the indices of the edges incident on u. Since every element of $E(G)$ is associated with exactly one index, the set of edges incident on u is equinumerous with this index set, so $d(u) = 2k$. Since u was arbitrarily chosen, every vertex of G must have even degree.

To prove that even degree and connectedness among edges are sufficient, we dispense with the trivial case where G consists only of isolated vertices and prove the nontrivial graphs are eulerian by contradiction. Consider a nontrivial graph G in which every vertex has even degree and all edges lie in a single component. For simplicity we may assume G is connected, since isolated vertices have no bearing on eulerian tours. Since G is connected and nontrivial, $\delta(G) \geq 1$; since every vertex has even degree, $\delta(G) \geq 2$. As seen when investigating trees, this property guarantees that G contains a cycle. Let C be a maximal tour in G. If C is non-eulerian, then there is at least one edge e in $G - C$. Either e is incident on a vertex of C, or there is a vertex in v not in C, and thus there is a path from some vertex of C to v; in either case we are guaranteed that some vertex u in C is incident on an edge in $G-C$. Let H be the component of $G-C$ containing u. Since u is incident on an edge, H is nontrivial and connected; furthermore, since a tour, as seen in the necessity argument, utilizes an even number of edges incident to each vertex, the degrees in H have the same parity as degrees in G; that is, they are all even. Thus, $\delta(H) \geq 2$, so H contains some cycle D. Since D and C share noe dges and share at least one vertex u, they can be spliced together to get a tour whose length is the sum of the lengths of C and D, contradicting C's maximality. □

Producing an eulerian tour is pretty easy, by following a procedure suggested by the above necessity proof: choose a tour by any traversal mechanism (randomly following unused edges until you return to your start point will work); if it's not eulerian, find a vertex on you tour incident to untraveled edges, and build a new tour starting with one of its unused edges. Splice the new tour into your original tour; repeat this procuedure until all edges are used.

Eulerian trails, which are occiasionally useful, are also easy to produce and test for existence:

Corollary 1. *A graph G has an eulerian trail if and only if it has either zero or two vertices of odd degree and all edges lie in a single component.*

Proof. If G has no vertices of odd degree, then G is eulerian and contains an eulerian tour; this tour is also a trail.

Since $\sum_{v \in V(G)} d(v) = 2\|G\|$, the sum of all degrees must be even, so G cannot have exactly one vertex of odd degree.

If G has two vertices u and v of odd degree, then adding the edge $e = \{u, v\}$ to G results in a graph G' in which every vertex has even degree. G' is eulerian, so there is an eulerian tour of G'. Removal of a single edge from a tour produces a non-closed trail; thus removal of e from the eulerian tour of G' yields and eulerian trail in G.

Conversely, suppose that G has an eulerian trail from some u to v. If $u = v$, G must be eulerian and thus has zero vertices of odd degree; if $u \neq v$, then the trail can be modified to give an eulerian tour in $G' = G + \{u, v\}$, by adding the edge $\{u, v\}$ to the trail. Since G' is eulerian and all degrees except two are identical to those in G, G must have only two vertices of odd degree. □

3 Random Selecta: Matrices and Subgraph Parameters

3.1 Matrices

This is a collection of odds-and-ends that don't fit elsewhere. They're not much of a cohesive story, so they get to be presented in text on our off-day instead of trying to get a good lecture out of them.

Definition 8. The *adjacency matrix* of a graph G, with order $v_1, \ldots, v_n$ imposed on the vertices, is an $n \times n$ matrix A such that $A_{ij} = 1$ if v_i and v_j are adjacent, and $A_{ij} = 0$ if v_i and v_j are non-adjacent.

Obviously, a graph could yield several different matrices if the vertices are arranged in different patterns, but there are a few common configurations worth mentioning. A bipartite graph is conventionally represented with each part ordered consecutively, so that the adjacency matrix of a bipartite graph with parts S_1 and S_2 is the block matrix $\begin{bmatrix} \mathbf{0}_{|S_1|} & \mathbf{A} \\ \mathbf{A}^T & \mathbf{0}_{|S_2|} \end{bmatrix}$, where $\mathbf{0}_n$ is an $n \times n$ matrix of zeroes, and A is an $|S_1| \times |S_2|$ matrix indicating adjacency between elements of S_1 and S_2. Likewise, a disconnected graph is conventionally ordered with vertices in components grouped together, so that if G consists of components $C_1, C_2, \dots, C_r$, then G has adjacency matrix given by the block structure

$$A(G) = \begin{bmatrix} A(C_1) & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & A(C_2) & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & A(C_3) & \dots & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \dots & A(C_r) \end{bmatrix}$$

There are several tricks one can do with adjacency matrices: for instance, if A is the adjacency matrix of a graph G, it is easy to see that $A \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} = \begin{bmatrix} d_G(v_1) \\ d_G(v_2) \\ \vdots \\ d_G(v_n) \end{bmatrix}$. As a simple corollary to this fact, we see that A has $(1, 1, 1, \dots, 1)$ as an eigenvector if and only if G is regular, and it is specifically k-regular where k is the eigenvalue associated with this eigenvector.

One useful purpose to which an adjacency matrix can be put is investigations into walks.

Proposition 6. *The ij-th entry of A^k is equal to the number of walks from v_i to v_j of length k.*

Proof. We prove this by induction on k. For $k = 0$, note $A^0 = I$; v_i and v_j are connected by a walk of length zero iff $i = j$.

Considering the inductive step: a walk of length k from v_i to v_j consists of a walk of length $k-1$ from v_i to some v_t, and then a walk of length 1 — that is, a single edge — from v_t to v_j. Thus, the number of walks of length k from v_i to v_j is equal to the sum of the number of walks of length $k-1$ from v_i to some neighbor v_t of v_j. Using the inductive hypothesis, this quantity is:

$$\sum_{v_t \sim v_j} (A^{k-1})_{it} = \sum_{v_t \in G} (A^{k-1})_{it} A_{tj} = \sum_{t=1}^{n} (A^{k-1})_{it} A_{tj} = (A^{k-1}A)_{ij} = (A^k)_{ij}$$

□

A somewhat less commonly used matrix representation of a graph is the *incidence matrix*:

Definition 9. The *incidence matrix* of a graph G, with order $v_1, \dots, v_n$ imposed on the vertices and $e_1, \dots, e_m$ on the edges, is an $n \times m$ matrix B such that $B_{ij} = 1$ if $v_i \in e_j$.

One simple fact can relate the incidence matrix to the adjacency matrix:

Proposition 7. *If B is the incidence matrix of a graph G, and A is the adjacency matrix with the same vertex order, and D is a diagonal matrix whose ith diagonal entry is $d_G(v_i)$, then $BB^T = A + D$.*

Proof. Let us look at the ij-th entry of BB^T. By the definition of matrix multiplication:

$$(BB^T)_{ij} = \sum_{k=1}^{m} b_{ik} b_{jk}$$

Note that $b_{ik}b_{jk}$ is a product of indicator functions which are 1 if e_k is incident on v_i and v_j respectively, and zero otherwise. Their product is thus 1 if e_k is incident on both v_i and v_j. So $(BB^T)_{ij}$ is the sum over all edges of the indicator function testing if that edge is incident to both v_i and v_j; that is, it is the number of edges incident to both v_i and v_j. When $i = j$, this is the number of edges incident to v_i, so $(BB^T)_{ii} = d_G(i) = D_{ii} + A_{ii}$, since the diagonal entries of the adjacency matrix are zero; when $i \neq j$, since v_i and v_j are distinct, the only edge which could be incident on both of them is $\{v_i, v_j\}$; this edge exists only if v_i is adjacent to v_j, so $(BB^T)_{ij} = A_{ij} = D_{ij} + A_{ij}$, since off-diagonal elements of D are zero. □

3.2 Subgraph properties

These were briefly mentioned last semester, but they're orphaned concepts that don't really have a place, so they get to go on the end of this notes section.

Definition 10. A *clique* of a graph G is a set of mutually adjacent vertices; an *independent set* of G is a set of mutually non-adjacent numbers. The *clique number* $\omega(G)$ is the size of the largest clique in G; the *independence number* $\alpha(G)$ is the size of the largest independent set.

Note that cliques are just complete subgraphs, so the clique nubmer can be reinterpreted as the largest r such that $K_r \subseteq G$.

Also, note that independent sets of G are cliques in G^c and vice versa; consequently $\alpha(G) = \omega(G^c)$ and $\omega(G) = \alpha(G^c)$.

A simple example of relating these properties to other graph properties:

Proposition 8. $\omega(G) \leq \frac{1+\sqrt{1+8\|G\|}}{2}$ *and* $\alpha(G) \leq \frac{1+\sqrt{1+8\|G^c\|}}{2}$.

Proof. If $\omega(G) = k$, then K_k is a subgraph of G; thus G must contain at least $\|K_k\| = \binom{k}{2} = \frac{k(k-1)}{2}$ edges. Thus, $\|G\| \geq \frac{k}{k-1}$; solving this quadratic yields $\frac{1-\sqrt{1+8\|G\|}}{2} \leq k \leq \frac{1+\sqrt{1+8\|G\|}}{2}$. Since $\omega(G)$ must be positive, only the upper bound is really significant.

The proof for $\alpha(G)$ can e built off of this by noting that $\alpha(G) = \omega(G^c) \leq \frac{1+\sqrt{1+8\|G^c\|}}{2}$. □

This bound is, more often than not, unacceptably awful.

Another useful set of subgraphs to inspect is the presence of cycles as subgraphs of G.

Definition 11. The *girth* of a graph G containing a cycle, denoted $g(G)$, is the smallest k such that $C_k \subseteq G$; the circumference of such a graph is the largest k such that $C_k \subseteq G$.

Under some conventions, the girth and circumference of a forest are defined to be ∞. Circumference is infrequently used, although G having a circumference of $|G|$ is an equivalent condition to G being hamitonian.

Girth can be used in a number of ways, as in this example of a trivial girth-based bound on the clique number and a less trivial bound on the independence number.

Proposition 9. *For G containing a cycle, $g(G) > 3$ if and only if $\omega(G) \leq 2$.*

Proof. If $g(G) > 3$, then by definition $C_3 \not\subseteq G$. Since $C_3 = K_3$, and K_3 is a subgraph of K_n for all $n \geq 3$, $K_n \not\subseteq G$ for $n \geq 3$ and thus $\omega(G) < 3$.

Conversely, if $\omega(G) \leq 2$, then $K_3 \not\subseteq G$, so $C_3 \not\subseteq G$, so $g(G) \neq 3$. However, since G contains a cycle (which must have size at least 3) and no cycle of size 3, its smaller cycle must have size larger than 3 and thus $g(G) > 3$. □

Proposition 10. *For G containing a cycle, $\alpha(G) \geq \lfloor \frac{g(G)}{2} \rfloor$.*

Proof. Let $C = v_1 \ v_2 \ \cdots \ v_{g(G)} \ v_1$ be a cycle in G of length $g(G)$. If v_i and v_j with $i < j$ were adjacent without being consecutive in the structure above, then $v_i \sim v_{i+1} \sim \cdots \sim v_{j-1} \sim v_j \sim v_i$ would be a cycle in G of length $j - i + 1 < g(G)$, which would contradict minimality of $g(G)$. Thus there are no adjacencies among vertices of the cycle except those described in the cycle structure, so $\{v_1, v_3, v_5, \ldots, v_{2\lfloor \frac{g(G)-2}{2} \rfloor + 1}\}$ is an independent set of size $\lfloor \frac{g(G)}{2} \rfloor$, so the largest independent set has at least that magnitude. □