Graph theory

Po-Shen Loh

June 2012

At first, graph theory may seem to be an *ad hoc* subject, and in fact the elementary results have proofs of that nature. The methods recur, however, and the way to learn them is to work on problems. Later, when you see an Olympiad graph theory problem, hopefully you will be sufficiently familiar with graph-theoretic arguments that you can rely on your own ingenuity to invent an approach.

1 Warm-up

1. Are these two graphs the same?

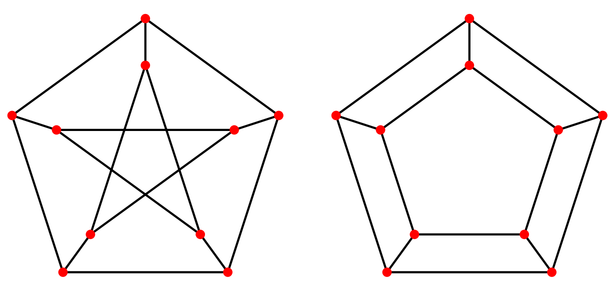

   Solution: No. The Petersen graph (to the left) has no cycles of length 4.

2. There are 21 friends at a party. Each person counts how many of the other people are he/she has met before. Is it possible for the sum of these counts to be an odd number?

3. Show that every graph has at least two vertices with equal degree.

   Solution: Pigeonhole: all degrees between 0 and $n-1$, but if we have a 0, we cannot have an $n-1$. So there are $n-1$ available degrees and n vertices.

2 Terminology

Definition. *A **graph** $G = (V, E)$ is a collection V of **vertices** and $E \subset V \times V$ of **edges**.*

Remark. Informally, we think of the edges as linking the pairs of vertices that they correspond to, and typically represents graphs by drawings in which we connect the endpoints by a curve.

Another remark. If every edge links a unique pair of distinct vertices, then we say that the graph is *simple*. Most of our work will be with simple graphs, so we usually will not point this out.

Glossary of terms.

- Two vertices are *adjacent* if there is an edge that has them as endpoints.

- A graph is *bipartite* if the vertex set can be partitioned into two sets $V_1 \cup V_2$ such that edges only run between V_1 and V_2.
- The *chromatic number* of a graph is the minimum number of colors needed to color the vertices without giving the same color to any two adjacent vertices.
- A *clique* on n vertices, denoted K_n, is the n-vertex graph with all $\binom{n}{2}$ possible edges.
- A *complete* graph on n vertices, denoted K_n, is the n-vertex graph with all $\binom{n}{2}$ possible edges.
- A graph is *connected* if there is a path between every pair of distinct vertices.
- A *cycle* is a path for which the first and last vertices are actually adjacent.
- The *degree* $d(v)$ of a vertex v is the number of edges that are incident to v.
- An *Eulerian circuit* is a walk that traverses every edge exactly once, and returns to its starting point.
- A *forest* is a not-necessarily-connected graph with no cycles. See also the definition of *tree* below.
- If U is a subset of the vertices, then the *induced subgraph* $G[U]$ is the graph obtained by deleting all vertices outside U, keeping only edges with both endpoints in U.
- A *Hamiltonian path* is a path that includes every vertex. A *Hamiltonian cycle* is a cycle that includes every vertex.
- We say that an edge e is *incident* to a vertex v if v is an endpoint of e.
- A *path* is a sequence of distinct, pairwise-adjacent vertices. See also the definition of *walk* below.
- A graph is *planar* if it is possible to draw it in the plane without any crossing edges.
- A *tree* is a connected graph with no cycles. See also the definition of *forest* above.
- A *walk* is a sequence of not-necessarily-distinct, pairwise-adjacent vertices. See also the definition of *path* above.

3 Classical results

1. The sum of all of the degrees is equal to twice the number of edges.

   Solution: By counting in two ways, we see that the sum of all degrees equals twice the number of edges.

2. Every bipartite graph has chromatic number ≤ 2.

3. Every n-vertex tree has exactly $n - 1$ edges.

4. Every connected graph contains a *spanning tree*, i.e., a tree which includes every vertex.

5. Every connected graph with all degrees even has an *Eulerian circuit*, i.e., a walk that traverses each edge exactly once.

   Solution: Start walking from a vertex v_1 without repeating any edges, and observe that by the parity condition, the walk can only get stuck at v_1, so we get one cycle. If we still have more edges left to hit, connectivity implies that some vertex v_2 on our current walk is adjacent to an unused edge, so start the process again from v_2. Splice the two walks together at v_2, and repeat until done.

6. (Dirac.) Let G be a graph on n vertices with all degrees at least $n/2$. Show that G has a Hamiltonian cycle.

   Solution: Suppose the longest path has t vertices $x_1, \ldots, x_t$. We will show there is a cycle of t vertices as well. Suppose not. All neighbors of x_1 and x_t must lie on the path or else it is not longest. Minimum degree condition implies that both have degree $\geq t/2$. But if $x_1 \sim x_k$, then $x_t \not\sim x_{k-1}$ or else we can re-route to get a cycle. So, each of x_1's $t/2$ neighbors on the path prohibit a potential neighbor of x_t. Yet x_t's neighbors come from indices $1 \ldots t-1$, so there is not enough space for x_t to have $t/2$ neighbors there, avoiding the prohibited ones.

   Now if this longest path is not the full n vertices, then we get a cycle C missing some vertex x. But min-degree $n/2$ implies that the graph is connected (smallest connected component is $n/2+1$), so there is a shortest path from x to C, and adding this to the cycle gives a longer path than t, contradiction.

7. (Hall's Marriage Theorem.) For any set $S \subset A$, let $N(S)$ denote the set of vertices (necessarily in B) which are adjacent to at least one vertex in S. Then, A has a perfect matching to B if and only if $|N(S)| \geq |S|$ for **every** $S \subset A$.

8. (Euler characteristic.) Every connected planar graph satisfies $V - E + F = 2$, where V is the number of vertices, E is the number of edges, and F is the number of faces.

9. (Kuratowski.) A graph is planar if and only if it contains no subdivision of K_5 or $K_{3,3}$.

10. (Four color theorem.) Every planar graph can be properly colored using at most 4 colors.

11. (Vizing.) If the maximum degree of a graph is Δ, then at least Δ colors are required to properly color its edges, but $\Delta + 1$ colors are sufficient.

12. (Turán.) For $r \geq 3$, the Turán graph $T_{r-1}(n)$, with $r-1$ parts of as equal size as possible, and all edges between distinct parts, is the unique n-vertex graph with the maximum number of edges subject to having no K_r subgraphs.

4 Well-known facts

1. Let G be a graph. It is possible to partition the vertices into two groups such that for each vertex, at least half of its neighbors ended up in the other group.

   Solution: Take a max-cut: the bipartition which maximizes the number of crossing edges.

2. Let δ be the minimum degree of G, and suppose that $\delta \geq 2$. Then G contains a cycle of size $\geq \delta + 1$. In particular, it contains a path with $\geq \delta$ edges.

   Solution: Take a longest path. Let v be its last endpoint. By maximality, every one of v's $\geq \delta$ neighbors lie on the path. So path has length $\geq \delta + 1$.

3. Consider a graph where every vertex has degree exactly $2k$. Show that it is possible to orient each edge such that the maximum in-degree is exactly k.

   Solution: Direct along an Eulerian circuit.

4. Every k-regular bipartite graph can have its edges partitioned into k edge-disjoint perfect matchings.

   Solution: Suffices to find one perfect matching. Every set S expands because it has k edges out, and each vertex on the other side can only absorb up to k of them in.

5. (Petersen, 1891.) A *2-factor* of a graph is a 2-regular spanning subgraph (i.e., containing all vertices, and having all degrees equal to 2). For every positive integer k, show that every $2k$-regular graph can be partitioned into k edge-disjoint 2-factors.

Solution: Suffices to find one 2-factor. Take an Eulerian orientation. Split each vertex v into v^+, v^-. This gives a bipartite graph with twice as many vertices. If there was an edge $\overrightarrow{vw}$, now put it from v^- to w^+. It is a k-regular bipartite graph, so it has a perfect matching by above. Collapsing back the v^+, v^-, we get a 2-factor.

6. (Sperner.) An *antichain* in the lattice of subsets $2^{[n]}$ is a collection $\mathcal{F}$ of subsets of $\{1, 2, \ldots, n\}$ such that no subset in F completely contains the other. Prove that the size of the largest antichain is $\binom{n}{\lfloor n/2 \rfloor}$.

   Solution: We can partition the entire lattice into that many chains. To see this, start from the bottom, and go up to level $\lfloor n/2 \rfloor$. At each step, when we go up, we need to find a perfect matching from the lower layer to the upper layer which sends sets to supersets. This is a regular bipartite graph, and since degrees in the lower layer are $\geq$ degrees in the upper layer, we can always proceed.

7. (Birkhoff, Von Neumann.) Let A be a square $n \times n$ matrix of nonnegative integers, in which each row and column sum up to the positive integer m. Prove that A can be expressed as a sum of m *permutation matrices* $A = P_1 + \cdots P_m$. Here, a permutation matrix is an $n \times n$ matrix of zeros and ones, such that each row contains a single "1", and each column contains a single "1."

   Solution: Pull out permutation matrices inductively via Hall's theorem. LHS is the rows, RHS is the columns, and there is an edge between i on the LHS and j on the RHS if a_{ij}, the (i, j) entry of A, is > 0. Actually, write the number a_{ij} next to that edge.

   For any subset S of the LHS, note that the sum of all numbers on edges coming out of S is precisely $m|S|$. But those edges are a subset of the edges coming out of $N(S)$, so their sum is $\leq m|N(S)|$. Thus $|N(S)| \geq |S|$ for any S.

8. Every planar graph can be properly colored using at most 6 colors.

   Solution: Since $E \leq 3V - 6$ for connected planar graphs, they are 5-degenerate.

9. Every *tournament* (complete graph with every edge oriented in one of the two directions) contains a vertex of degree at least $\frac{n-1}{2}$.

10. Every graph G with average degree d contains a subgraph H such that all vertices of H have degree at least $d/2$ (with respect to H).

    Solution: Condition on G is that the number of edges is at least $nd/2$. If there is a vertex with degree $< d/2$, then delete it, and it costs 1 vertex and $< d/2$ edges, so the condition is preserved. But it can't go on forever, because once there is 1 vertex left, average degree is 0.

5 Problems

1. (Romania, 2006.) Each edge of a polyhedron is oriented with an arrow such that every vertex has at least one edge directed toward it, and at least one edge directed away from it. Show that some face of the polyhedron has its boundary edges coherently oriented in a circular direction.

   Solution: A directed cycle in the graph exists by simply following out-edges until we repeat vertices. Take a shortest directed cycle. If this is not a face, then it has a chord. Yet no matter which way the chord is oriented, one of the two sub-cycles will be strictly shorter, and still be coherent.

2. (Hungary 2010.) There were n people at a party. Sometimes three persons played a game of cards. At the end of the party it turned out that any three persons played in at most one game together and any two persons played exactly twice together. For which values of n is this possible if $3 < n < 9$?

   Solution: Answer: for $n = 4, 6, 7$. We are covering each edge of K_n twice with triangles. Thus we must have $3 \mid n(n-1)$, which leaves only the three options above. It remains to find the constructions. For $n = 4$, it suffices to take every one of the 4 possible triangles.

For $n = 6$, we will need 10 triangles. Start with the "Cayley" triangles $(x, x+1, x+3)$, identifying the vertices with $\mathbb{Z}_6$ and using all 6 starting points x. This has the effect of covering every edge exactly once, except for the 3 main diagonals, which are double-covered already. To hit the rest, use the three triangles $(0, 1, 2)$, $(2, 3, 4)$, and $(4, 5, 0)$, and the triangle $(1, 3, 5)$.

For $n = 7$, we will need 14 triangles. Use the Cayley triangles $(x, x+1, x+3)$ and $(x, x-1, x-3)$.

3. (Belarus 2010/C7.) Let $n \geq 3$ distinct points be marked on a plane so that no three of them lie on the same line. All points are connected with the segments. All segments are painted one of the four colors so that if in some triangle (with the vertices at the marked points) two sides have the same color, then all its sides have the same color (each of the four colors is used). What is the largest possible value of n?

   Solution: Answer: 9. This makes each color class an equivalence relation, so that each color class is the vertex-disjoint union of cliques. Consider a maximal clique in color 1. It can't be everybody, so there's someone else, v. Since it's maximal, every edge from v to it has color not 1. But equivalence classes, so the edges from v to it are all different colors. There are only 3 colors left. Thus the maximal clique has order at most 3. In particular, every color class's cliques are at most triangles, and hence the maximum degree in each color class is at most 2.

   There are only 4 colors, so the maximum degree is at most 8, and thus $n \leq 9$. Construction for $n = 9$: split the vertices into three triples for the first color class: (a_1, a_2, a_3), (b_1, b_2, b_3), (c_1, c_2, c_3). Next color class makes 3 cliques (a_i, b_i, c_i) for constant i. Next color class makes 3 cliques (a_i, b_{i+1}, c_{i+2}), wrapping mod 3. Fourth color class makes 3 cliques (a_i, b_{i+2}, c_{i+4}), wrapping mod 3.

4. (Sweden 2010.) A town has $3n$ citizens. Any two persons in the town have at least one common friend in this same town. Show that one can choose a group consisting of n citizens such that every person of the remaining $2n$ citizens has at least one friend in this group of n.

   Solution: The codegree condition implies that the diameter of the graph is at most 2. We prove that every n-vertex graph with diameter ≤ 2 has a dominating set (a subset S of vertices such that every other vertex is either in, or has a neighbor in S) of size only $\leq \sqrt{n \log n} + 1$. To see this, let $p = \sqrt{\frac{\log n}{n}}$.

   Observe that since the diameter is at most 2, if any vertex has degree $\leq np$, then its neighborhood already is a dominating set of suitable size. Therefore, we may assume that all vertices have degree strictly greater than np. It feels "easy" to find a small dominating set in this graph because all degrees are high. Consider a random sample of np vertices (selected uniformly at random, with replacement), and let S be their union. Note that $|S| \leq np$. Now the probability that a particular fixed vertex v fails to have a neighbor in S is strictly less than $(1-p)^{np}$, because we need each of np independent samples to miss the neighborhood of v. This is at most $e^{-np^2} \leq e^{-\log n} = n^{-1}$. Therefore, a union bound over the n choices of v produces the result.

5. (Sweden 2010.) Some of n students in a class ($n \geq 4$) are friends. Any $n-1$ students in the class can form a circle so that any two students next to each other on the circle are friends, but all n students cannot form a similar circle. Find the smallest possible value of n.

   Solution: Answer: $n = 10$, the Petersen graph. It is well-known that the Petersen graph is not Hamiltonian, but it is easy to see that if one deletes any vertex, one can find a Hamilton cycle. This is easy to check by symmetry, because all outer vertices are the same, and all inner vertices are the same.

   To see why Petersen is not Hamiltonian, observe that it is two disjoint 5-vertex graphs linked by a single perfect matching. Any Hamilton cycle must cross back and forth between the parts. If it just goes across once, and then back, then on each side it must visit all 5 vertices in one go. Those are paths of length 4, and it's easy to see that if one takes 4 consecutive edges along the outer cycle, then it doesn't complete to an H-cycle. Otherwise, the H-path must go across, back, across, and back. It can't do more times because there are only 5 matching edges. In this case, WLOG start with two consecutive outer edges, and take forced moves until we are stuck without an H-cycle.

Now we must show that no $n \le 9$ will work. First key observation: if there is a vertex of degree 2 or less, then it's impossible. Indeed, if so, then delete a neighbor of it; the remainder must be Hamiltonian, but now this vertex has degree ≤ 1, contradiction. Thus the minimum degree is at least 3, which already disposes of all cases $n \le 6$.

Next observation: let v be the max-degree vertex. Since $G-v$ is still Hamiltonian, take an H-cycle of the remainder. If two adjacent vertices of the cycle are neighbors of v, then we can extend, contradiction. So v's neighbors on the cycle are separated by at least one vertex each.

For $n = 7$, we can't have all degrees equal to 3, because the sum of degrees must be even. Thus there is a degree-4 vertex. But the remainder cycle has 6 vertices, so it's impossible to alternate. (Another way to see that $n = 7$ fails: using an H-cycle from $G - v$, and then adding v, we get an H-path with v as an endpoint, but the other endpoint has degree ≥ 3, and $4+3 \ge n$, so the Ore-type condition wins.)

For $n = 8$, the above alternating condition shows that it is over if there is a vertex of degree 4. Hence the graph is 3-regular. Pull out a vertex, and look on the remaining 7-cycle. Add back its neighbors, and there is actually only one way to do so with proper spacing. Then there are only 2 ways to complete to a 3-regular graph, and in both cases the entire thing is Hamiltonian.

For $n = 9$, we can't have all odd degrees, so there's a vertex of degree 4. It must interact with the remaining 8-cycle in exactly one way: alternating neighbors. Now the rest of the vertices on the 8-cycle need degrees ≥ 3. Let A be the set of neighbors of v, and let B be the others. If two consecutive (separated only by a single vertex of A) vertices of B are adjacent, then one can see that the whole thing is Hamiltonian. If two opposite vertices of B are adjacent, then also the whole thing is Hamiltonian. Thus the only way to relieve the degrees is to have the vertices of B adjacent to vertices of A that they are not already adjacent to. And we can't create any vertices of degree 5, or else done by failing to alternate as above. Then there is exactly one way to make this graph, and already, all vertices of A have degree 4 and all vertices of B have degree 3, and v has degree 4. No more edges can be added because connecting two vertices of B wins already as above. Thus this is the whole graph. But deleting any vertex of A we can see that the remainder is not Hamiltonian.

6. (Hungary 2010.) Prove that the edges of the complete graph with 2009 vertices can be labeled with $1, 2, \ldots, \binom{2009}{2}$ such that the sum of the labels corresponding to all edges having a given vertex is different for any two vertices.

   Solution: Let $n = 2009$. Consider the random labeling. Let u and v be two fixed vertices. The label of the edge between u and v is irrelevant for the bad event that the label sums are equal at u and v. Expose the $n-2$ labels of edges from u to $[n] \setminus \{u, v\}$. Let their sum be S. Next we want to expose the $n-2$ labels of the edges from v to $[n] \setminus \{u, v\}$. It suffices to show that conditioned on the $n-2$ labels we already saw from u, their sum equals S with probability less than $1/\binom{n}{2}$, because then a union bound implies that there is a labeling that avoids all bad events.

   Intuitively, the probability is actually of order $n^{-5/2}$. To see this, suppose that the $n-2$ new random labels are sampled independently with replacement from the full set $I = \{1, \ldots, \binom{n}{2}\}$. If we sample one integer from I, the variance is of order n^4. Therefore, the variance of this slightly different sum is of order n^5, and since we are adding i.i.d. random variables, the distribution is "nice", and the probability that the sum is any particular number is of order at most $n^{-5/2}$.

   Now we formalize this. We make exactly $n-3$ i.i.d. samples from I. After this, we look to see whether we ever got the same label multiple times, or if we repeated a label we saw from u. For each of these occurrences, we re-sample uniformly from I until we find new labels, and ultimately build a set of $n-3$ distinct new labels. Finally, we repeat this procedure until we get a final new label, and that produces a set of $n-2$ labels which are distinct from those from u, while also uniformly distributed over all possibilities.

The first observation is that during the first round, the probability that we hit a repeat label is less than $(2n-2)/\binom{n}{2} = \frac{4}{n}$. Therefore, the number of times we will have to resample in the second round is stochastically dominated by Bin $\left[n, \frac{4}{n}\right]$, and the probability that such a Binomial exceeds $\log n$ is at most

$$\binom{n}{\log n}\left(\frac{4}{n}\right)^{\log n} < \left(\frac{4e}{\log n}\right)^{\log n} \ll n^{-3} .$$

Now note that success (getting a sum of exactly S) comes in one of two ways: (1) if the Binomial exceeds $\log n$, and then we get lucky, or (2) the Binomial stays below $\log n$, the sum of the $n-3$ labels after the first round is within $n^2 \log n$ of S, and then after the second round, the final label in the third round makes the sum exactly S. The chance of winning from (1) is at most n^{-3} from above. The chance of winning from (2) is at most the probability that the sum of the $n-3$ labels after the first round is within $n^2 \log n$ of S, and the final label makes the sum exactly S.

We calculate this by multiplying upper bounds of the probabilities that (a) the first round sum is within $n^2 \log n$ of S and (b) the conditional probability that the third round label makes the sum exactly equal to S. The latter probability is obviously at most $\frac{3}{n^2}$ because there is only one choice for it which would make the sum S. The probability of (a) can be bounded by the Central Limit Theorem, because the first round sum is precisely the sum of i.i.d. random variables with bounded second moment. So, by the CLT, the probability that this sum lies within any given window of at most $n^2 \log n$, given that the variance of the sum should be of order n^5, is $o(1)$.

7. (USAMO Awards 2012.) Find a 3-coloring of the edges of the dodecahedron.

   Solution: First find a Hamilton cycle, and then use two colors alternating on it. The remainder is a perfect matching.