1 Leaves in a Tree

A *leaf* in a tree is a vertex with degree 1.

(a) Consider a tree with $n \geq 3$ vertices. What is the largest possible number of leaves the tree could have? Prove that this maximum m is possible to achieve, and further that there cannot exist a tree with more than m leaves.

(b) Prove that every tree on $n \geq 2$ vertices must have at least two leaves.

Solution:

(a) We claim the maximum number of leaves is $n-1$. This is achieved when there is one vertex that is connected to all other vertices (this is called the *star graph*).

We now show that a tree on $n \geq 3$ vertices cannot have n leaves. Suppose the contrary that there is a tree on $n \geq 3$ vertices such that all its n vertices are leaves. Pick an arbitrary vertex x, and let y be its unique neighbor. Since x and y both have degree 1, the vertices x, y form a connected component separate from the rest of the tree, contradicting the fact that a tree is connected.

(b) We give a proof by contradiction. Consider the longest path $\{v_0, v_1\}, \{v_1, v_2\}, \ldots, \{v_{k-1}, v_k\}$ between two vertices $x = v_0$ and $y = v_k$ in the tree (here the length of a path is how many edges it uses, and if there are multiple longest paths then we just pick one of them). We claim that x and y must be leaves. Suppose the contrary that x is not a leaf, so it has degree at least two. This means x is adjacent to another vertex z different from v_1. Observe that z cannot appear in the path from x to y that we are considering, for otherwise there would be a cycle in the tree. Therefore, we can add the edge $\{z, x\}$ to our path to obtain a longer path in the tree, contradicting our earlier choice of the longest path. Thus, we conclude that x is a leaf. By the same argument, we conclude y is also a leaf.

The case when a tree has only two leaves is called the *path graph*, which is the graph on $V = \{1, 2, \ldots, n\}$ with edges $E = \{\{1,2\}, \{2,3\}, \ldots, \{n-1, n\}\}$.

2 Coloring Trees

(a) What is the minimum number of colors needed to color a tree? Prove your claim.

(b) Prove that all trees are *bipartite*: the vertices can be partitioned into two groups so that every edge goes between the two groups.

[*Hint:* How does your answer to part (a) relate to this?]

Solution:

(a) The minimum number of colors needed to color a tree is 2. Prove this using induction.

When the tree is just a single vertex, we can indeed color it with a single color.

Assume that all trees with k vertices can be 2-colored. Now consider a tree T with $k+1$ vertices. We know that every tree must have at least 2 leaves, so remove one leaf u and the edge connected to u. The resulting graph $T-u$ is a tree with k vertices and can be 2-colored by our inductive hypothesis.

Now when we add u back to the graph, u has a single neighbor which is colored with one color. Simply color u with the other color. Thus, any tree with $k+1$ vertices can be 2-colored.

We have shown that any tree of any size can be colored with only 2 colors.

(b) Since every tree can be colored with 2 colors, every tree is bipartite. If the vertices can be colored such that no two vertices connected by an edge share the same color, this means that every edge connects a vertex of color c_0 with a vertex of color c_1. Thus, the tree can be split into two groups: all vertices colored c_0 and all vertices colored c_1. By the rules of coloring, every edge must connect a vertex in the first group with a vertex in the second group. Every tree can therefore be considered a bipartite graph.

3 Edge-Disjoint Paths in a Hypercube

Prove that between any two distinct vertices x, y in the n-dimensional hypercube graph, there are at least n edge-disjoint paths from x to y (i.e., no two paths share an edge, though they may share vertices).

Solution:

We use induction on $n \geq 1$. The base case $n = 1$ holds because in this case the graph only has two vertices $V = \{0, 1\}$, and there is 1 path connecting them. Assume the claim holds for the $(n-1)$-dimensional hypercube. Let $x = x_1x_2\dots x_n$ and $y = y_1y_2\dots y_n$ be distinct vertices in the n-dimensional hypercube; we want to show there are at least n edge-disjoint paths from x to y. To do that, we consider two cases:

1. Suppose $x_i = y_i$ for some index $i \in \{1, \dots, n\}$. Without loss of generality (and for ease of explanation), we may assume $i = 1$, because the hypercube is symmetric with respect to the indices. Moreover, by interchanging the bits 0 and 1 if necessary, we may also assume $x_1 = y_1 = 0$. This means x and y both lie in the 0-subcube, where recall the 0-subcube (respectively, the 1-subcube) is the $(n-1)$-dimensional hypercube with vertices labeled $0z$ (respectively, $1z$) for $z \in \{0, 1\}^{n-1}$.

Applying the inductive hypothesis, we know there are at least $n-1$ edge-disjoint paths from x to y, and moreover, these paths all lie within the 0-subcube. Clearly these $n-1$ paths will still be edge-disjoint in the original n-dimensional hypercube. We have an additional path from x to y that goes through the 1-subcube as follows: go from x to x', then from x' to y' following any path in the 1-subcube, and finally go from y' back to y. Here $x' = 1x_2 \dots x_n$ and $y' = 1y_2 \dots y_n$ are the corresponding points of x and y in the 1-subcube. Since this last path does not use any edges in the 0-subcube, this path is edge-disjoint to the $n-1$ paths that we have found. Therefore, we conclude that there are at least n edge-disjoint paths from x to y.

2. Suppose $x_i \neq y_i$ for all $i \in \{1, \dots, n\}$. This means x and y are two opposite vertices in the hypercube, and without loss of generality, we may assume $x = 00\dots0$ and $y = 11\dots1$. We explicitly exhibit n paths $P_1, \dots, P_n$ from x to y, and we claim they are edge-disjoint.

   For $i \in \{1, \dots, n\}$, the i-th path P_i is defined as follows: start from the vertex x (which is all zeros), flip the i-th bit to a 1, then keep flipping the bits one by one moving rightward from position $i+1$ to n, then from position 1 moving rightward to $i-1$. For example, the path P_1 is given by

   $$000\dots0 \to 100\dots0 \to 110\dots0 \to 111\dots0 \to \cdots \to 111\dots1$$

   while the path P_2 is given by

   $$000\dots0 \to 010\dots0 \to 011\dots0 \to \cdots \to 011\dots1 \to 111\dots1$$

   Note that the paths $P_1, \dots, P_n$ don't share vertices other than $x = 00\dots0$ and $y = 11\dots1$, so in particular they must be edge-disjoint.

Alternative for Case 2:
We can also reduce case 2, in which x and y have no bits in common, to case 1.

Suppose that $x_i \neq y_i$ for all $i = 1, \dots, n$. Let $\tilde{x}$ be x with the first bit flipped, so now $\tilde{x}$ and y both lie on a subcube together. From the inductive hypothesis, there are $n-1$ edge-disjoint paths from $\tilde{x}$ to y along the shared subcube (these paths only flip bits 2 through n because the paths lie entirely on a subcube, and these paths are of length at least $n-1$ since $\tilde{x}$ and y differ by that many bits). We would like to make these into edge-disjoint paths from y to x.

Starting at y, take the first of the $n-1$ paths, but before starting, flip the first bit. Then follow the first path to get a path from y to x. Then to use the second path starting at y, we travel one step along the second path, then flip the first bit, and then continue along the second path (carry out the sequence of bit flips in the second path, giving us another path from y to x). Continuing this way, we take the $(n-1)$st path, travel $n-2$ steps along the path, flip the first bit, then continue following the path.

It can be seen that this gives $n-1$ edge-disjoint paths from x to y. Where do we get the last path? Well we take the first path again, and now we follow the path for $n-1$ steps, and then flip the first bit – this gives us yet another edge-disjoint path for a total of n edge-disjoint paths! Cool!

4 Triangulated Planar Graph

In this problem you will prove that every triangulated planar graph (every face has 3 sides; that is, every face has three edges bordering it, including the unbounded face) contains either (1) a vertex of degree 1, 2, 3, 4, (2) two degree 5 vertices which are adjacent, or (3) a degree 5 and a degree 6 vertices which are adjacent. Justify your answers.

(a) Place a "charge" on each vertex v of value $6 - \text{degree}(v)$. What is the sum of the charges on all the vertices? (*Hint*: Use Euler's formula and the fact that the planar graph is triangulated.)

(b) What is the charge of a degree 5 vertex and of a degree 6 vertex?

(c) Suppose now that we shift $1/5$ of the charge of a degree 5 vertex to each of its neighbors that has a negative charge. (We refer to this as "discharging" the degree 5 vertex.) Conclude the proof under the assumption that, after discharging all degree 5 vertices, there is a degree 5 vertex with positive remaining charge.

(d) If no degree 5 vertices have positive charge after discharging the degree 5 vertices, does there exist any vertex with positive charge after discharging? If there is such a vertex, what are the possible degrees of that vertex?

(e) Suppose there exists a degree 7 vertex with positive charge after discharging the degree 5 vertices. How many neighbors of degree 5 might it have?

(f) Continuing from Part (e). Since the graph is triangulated, are two of these degree 5 vertices adjacent?

(g) Finish the proof from the facts you obtained from the previous parts.

Solution:

(a) Let V be the vertex set , E be the edge set, F be the faces in the graph, we have

$$\begin{aligned} \sum_{v \in V} 6 - \text{degree}(v) &= 6|V| - \sum_{v \in V} \text{degree}(v) && (1) \\ &= 6|V| - 2|E|. && (2) \end{aligned}$$

The last step is because that we count each edge twice as degree for each end vertex. And since the graph is triangulated, each face uses exactly three edges and each edge is shared by two faces, so we can substitute $|F| = 2|E|/3$ in Euler's formula to get

$$\begin{aligned} |V| + |F| &= |E| + 2 && (3) \\ |V| + \frac{2|E|}{3} &= |E| + 2 && (4) \\ 3|V| + 2|E| &= 3|E| + 6 && (5) \\ |E| &= 3|V| - 6. && (6) \end{aligned}$$

Substitute (6) into (2) to get that the sum of charge is 12.

(b) The charge is 1 for degree 5 vertex, and 0 for degree 6 vertex.

(c) If there is a degree 5 vertex with positive remaining charge, that means at least one of its neighbors is not negatively charged. In other words, at least one of its neighbors has degree 1, 2, 3, 4, 5, or 6, which proves the statement.

(d) Yes, there exists a vertex with positive charge. Since we know that the sum of charge of the entire graph is 12, it is impossible to have no positively charged vertex.

The possible degrees of vertex that have positive charge after discharging are $1,2,3,4,7$. Vertices with degree at least 8 will have initial charge -2, and will not have positive charge even if all their neighbors are of degree 5.

Side note: one can use Part (e,f) to rule out the possibility of degree-7 vertices. But for simplicity we only require students to rule out degrees 5 and 8 and beyond. The idea is that, we can't have two adjacent degree-5 vertices to have zero-charge after discharging. Suppose that there exists a degree-7 vertex with positive charges after discharging, then by Part (e,f), out of the 7 neighbors, 6 of them have to be of degree 5, and two of these degree-5 vertices have to be adjacent. Therefore, if none of the degree-5 vertices have positive charges, then we get two adjacent degree-5 vertices that have zero-charge, which is a contradiction.

(e) At least 6 out of the 7 are degree 5.

(f) Since the graph is triangulated, observe that fixing a drawing of the planar graph, we can order neighbors of the degree 7 node clockwise. And every two consecutive neighbors (defined by the ordering) form a triangle with the degree 7 vertex. From Part (e) we know that at least 6 out of the 7 are of degree 5. Therefore, it is impossible that none of these degree-5 vertices are adjacent to another degree-5 node.

(g) We split the proof into several cases. First note that there is always a vertex with degree at most 5. Suppose the contrary that every vertex is of degree at least 6, then the total charge would not have been positive, contradicting Part (a), where we showed the total charge is always 12.

If there are no vertices with degree at least 5, then we see that this is case (1). Therefore we consider the case where there is always a vertex of degree 5 from now on.

When there is a degree-5 vertex with a positive remaining charge, the statement is true by Part (c). When there is no degree-5 vertex with positive remaining charge, we know from Part (d) that either there is a positively charged vertex with degree $1,2,3,4$ or with degree $6,7$. For a degree-6 vertex to have a positive charge after discharging, it must be adjacent to a degree-5 vertex. For a degree-7 vertex, we know that at least two degree 5 vertices are adjacent from Part (f) which concludes our proof.

Side note: alternatively, one could simply rule out the possibility of a degree-7 vertex as explained in Part (d).

5 Euclid's Algorithm

(a) Use Euclid's algorithm from lecture to compute the greatest common divisor of 527 and 323. List the values of x and y of all recursive calls.

(b) Use extended Euclid's algorithm from lecture to compute the multiplicative inverse of 5 mod 27. List the values of x and y and the returned values of all recursive calls.

(c) Find $x \pmod{27}$ if $5x + 26 \equiv 3 \pmod{27}$. You can use the result computed in (b).

(d) Assume a, b, and c are integers and $c > 0$. Prove or disprove: If a has no multiplicative inverse mod c, then $ax \equiv b \pmod{c}$ has no solution.

Solution:

(a) The values of x and y of all recursive calls are (you can get full credits without the column of x mod y):

Function Calls	(x, y)	x mod y
#1	$(527, 323)$	204
#2	$(323, 204)$	119
#3	$(204, 119)$	85
#4	$(119, 85)$	34
#5	$(85, 34)$	17
#6	$(34, 17)$	0
#7	$(17, 0)$	—

Therefore, $\gcd(527, 323) = 17$.

(b) To compute the multiplicative inverse of 5 mod 27, we first call `extended-gcd(27,5)`. Note that `(x div y)` in the pseudocode means $\lfloor x/y \rfloor$. The values of x and y of all recursive calls are (you can get full credits without the columns of x div y and x mod y):

Function Calls	(x, y)	x div y	x mod y
#1	$(27, 5)$	5	2
#2	$(5, 2)$	2	1
#3	$(2, 1)$	2	0
#4	$(1, 0)$	—	—

The returned values of all recursive calls are:

Function Calls	(d, a, b)	Returned Values
#4	—	$(1, 1, 0)$
#3	$(1, 1, 0)$	$(1, 0, 1)$
#2	$(1, 0, 1)$	$(1, 1, -2)$
#1	$(1, 1, -2)$	$(1, -2, 11)$

Therefore, we get $1 = (-2) \times 27 + 11 \times 5$ and

$$1 = (-2) \times 27 + 11 \times 5 \equiv 11 \times 5 \pmod{27},$$

so the multiplicative inverse of 5 mod 27 is 11.

(c)

$$\begin{aligned} 5x + 26 \equiv 3 \pmod{27} &\Rightarrow 5x \equiv 3 - 26 \pmod{27} \\ &\Rightarrow 5x \equiv -23 \pmod{27} \\ &\Rightarrow 5x \equiv 4 \pmod{27} \\ &\Rightarrow 11 \times 5x \equiv 11 \times 4 \pmod{27} \\ &\Rightarrow x \equiv 44 \pmod{27} \\ &\Rightarrow x \equiv 17 \pmod{27}. \end{aligned}$$

(d) False. We can have a counterexample: $a = 3$, $b = 6$, and $c = 12$, so a has no multiplicative inverse mod c (because $a = 3$ and $c = 12$ are not relatively prime). However, $3x \equiv 6 \pmod{12}$ has solutions $x = 2, 6, 10$ mod 12.

6 Fibonacci GCD

The Fibonacci sequence is given by $F_n = F_{n-1} + F_{n-2}$, where $F_0 = 0$ and $F_1 = 1$. Prove that, for all $n \geq 0$, $\gcd(F_n, F_{n-1}) = 1$,

Solution:

Proceed by induction.

Base Case: We have $\gcd(F_1, F_0) = \gcd(1, 0) = 1$, which is trivially true.

Inductive Hypothesis: Assume we have $\gcd(F_k, F_{k-1}) = 1$ for some $k \geq 1$.

Inductive Step: Now we need to show that $\gcd(F_{k+1}, F_k) = 1$ as well.

We can show that:

$$\gcd(F_{k+1}, F_k) = \gcd(F_k + F_{k-1}, F_k) = \gcd(F_k, F_{k-1}) = 1$$

Note that the second expression comes from the definition of Fibonacci numbers. The last expression comes from Euclid's GCD algorithm, in which $\gcd(x, y) = \gcd(y, x \bmod y)$, since

$$F_k + F_{k-1} \equiv F_{k-1} \pmod{F_k}$$

Therefore the statement is also true for $n = k + 1$.

By the rule of induction, we can conclude that $\gcd(F_n, F_{n-1}) = 1$ for all $n \geq 1$, where $F_0 = 0$ and $F_1 = 1$ and $F_n = F_{n-1} + F_{n-2}$.