EVASIVENESS OF BIPARTITE GRAPH PROPERTIES

JIAYANG ZHAO

ABSTRACT. This paper gives a brief overview of the evasiveness of graph properties and its relationship to topology. We begin by introducing the concept of evasiveness and give several examples of evasive and non-evasive graph properties. Next, we introduce several key topological concepts that will be important in proving the evasiveness of certain graph properties. We subsequently define simplicial complexes and describe how to view these objects as topological spaces. Finally, we give a proof of the fact that all non-trivial bipartite graph properties are evasive.

CONTENTS

1. Evasiveness of Graph Properties

We are given a finite graph G and a particular graph property P, such as whether the graph is empty or whether it contains a triangle. We know the vertices of G, but we only know some, possibly zero, of the edges that are or are not present in G. We are allowed to ask, one at a time, whether or not an edge is present in G. The goal of this game is to minimize the number of edges we must query in order to determine whether or not G satisfies P.

Some graph properties require us, in the worst cases, to query all possible edges. These graph properties are known as evasive properties. For example, if G has n vertices and we start with no information on the edges of G, an evasive graph property requires us to query $n(n-1)/2$ edges. If G is bipartite with A and B the sets of vertices in each of its partitions, an evasive property requires us to query all possible $|A| \cdot |B|$ edges. This paper will examine when graph properties are evasive.

Example 1.1. (Emptiness) Whether or not a given graph is empty is an evasive graph property. In the worst case, one can query every edge except one, finding that each queried edge is not present. In this case, the emptiness of the graph is determined by the final unqueried edge.

Example 1.2. (Triangle-Containing) Whether or not a given graph contains a triangle is also an evasive property. It is possible to show that we can always arrive at the case where we have one edge left to query, and we know that there are two adjacent edges adjacent to the unknown edge. In this case, whether or not the graph contains a triangle is dependent on the presence of the last unqueried edge.

Example 1.3. (At Least One Edge per Vertex in A) Given a bipartite graph G with partitions A and B, whether or not each vertex in A is connected to at least one edge is an evasive property. We note that no matter what order we query the edges, it is always possible that after $|A| \cdot |B| - 1$ queries, we arrive at a graph where $|A| - 1$ vertices in A are connected to exactly one edge, but the last vertex in A is not. For example, this occurs when the answers to all of our queries is "not present" unless we are asking for the last possible edge of a vertex in A.

In the year 1973, it was conjectured that all non-trivial graph properties on graphs of n vertices require at least $O(n^2)$ queries. However, this conjecture was soon disproven by counterexamples such as the following scorpion graph property, which only requires $O(n)$ queries.

Example 1.4. (Scorpion Graph) A scorpion graph is a graph containing a three-vertex path such that one endpoint of the path is connected to all remaining vertices while the other two path vertices have no incident edges besides the ones in the path.

First, we note that any graph can contain at most one scorpion, uniquely defined by its three-vertex tail. Proof: Suppose there are two scorpions in the same graph. The scorpions must have different tails. However, by definition, one of the endpoints of the first scorpion's tail must be connected to at least two of the vertices of the second scorpion's tail. This is impossible, as two of the vertices of the second tail should have no other incident edges besides the ones connecting them.

We will label the three vertices of the tail a, b, and c, where a is only connected to b, b is only connected to c, and c is connected to all remaining vertices. Given a graph G with n vertices, select a vertex p in G and query all of its edges.

If p is connected to either none or all other vertices, then G is not a scorpion.

If p is connected to all vertices except one, then it must be the c of the tail. In this case, there is exactly one vertex to which p is not connected, and that vertex must be the a of the tail. Query all of a's edges to see that it is connected to exactly one other vertex b, and query all the edges of b to confirm that G is indeed a scorpion. This process should take $n - 1 + n - 2 + n - 3 = 3n - 6$ queries.

Otherwise, if $1 \leq deg(p) \leq n - 3$, then we know that p is not c. If $deg(p) = 1$ or 2, we can query all of p's adjacent vertices to check for the scorpion. Up to this point, we will have used $3n - 6$ queries. If we do not find the scorpion, or if originally $3 \leq deg(p) \leq n - 3$, then we know p is not on the tail.

Separate the remaining $n - 1$ vertices into two sets: A, the set of vertices not connected to p; and C, the set of vertices connected to p. We note that a must be in A, and c must be in C. Take a vertex x in A and a vertex z in C and query their edge. If they are connected, then x cannot be a, so remove x from A. If they are not connected, then select another vertex x' in A and query the edge $\{x', z\}$. If these two are connected, then we similarly remove x' from A. If they are not connected, then we remove z from C, since $deg(z) < n - 2$ thus it cannot be c. Note that in this process, after every two queries we can remove at least one vertex from either A or C. Continue this process until only one vertex remains in either A or C. Since $|A| + |C| = n - 1$, the process will end within $2n - 2$ queries. At the end of this process, we will have identified either a or c of the tail.

Query all edges along the potential tail to check for the scorpion. This process should take at most another $n - 4 + n - 5 + n - 6 = 3n - 15$ queries. Thus in total, we can determine whether or not a graph is a scorpion within $3n - 6 + 2n - 2 + 3n - 15 = 8n - 23$ queries. $\square$

After the discovery of these counterexamples, the conjecture was later weakened by strengthening the hypothesis to require that the graph property be monotone as well. A monotone graph property is a property such that if a graph G satisfies the property, then any graph to which G is a subgraph will also satisfy the property. We note that the scorpion graph property is not a monotone property, but all of the prior examples are. This conjecture, known as the Aanderaa-Rosenburg conjecture, was proven in 1975, when it was shown that all non-trivial, monotone properties require $n^2/16$ queries. This lower bound was later improved to $n^2/9$, then to $n^2/4$, and finally to $n^2/3 - o(n^2)$ in 2013.

Amidst these developments, it was conjectured, in what would be known as the Aanderaa-Karp-Rosenburg conjecture, that all non-trivial, monotone graph properties are evasive. This conjecture remains unresolved. However, it is known to be true when $|G|$ is prime, as well as when G is bipartite.

2. Contractability and Deformation Retracts

In the following section, we will define and explain some basic notions and examples regarding relevant topological ideas. As we shall later see in this paper, the evasiveness of a graph property is intimately connected to the contractibility of a certain topological space associated with that property.

Definition 2.1. Given continuous $f, g : X \to Y$, we say that f and g are homotopic if there exists a continuous function $H : X \times [0, 1] \to Y$ such that $H(x, 0) = f(x)$ and $H(x, 1) = g(x)$.

Definition 2.2. A topological space X is contractible if the identity function on X is homotopic to a constant function valued at some point x_0 in X.

Example 2.3. The unit disk $\mathbb{D}$ is contractible. Let $H(x, t) = (1 - t)x$. Then $H(x, 0) = (1 - 0)x = x$ and also $H(x, 1) = (1 - 1)x = 0$. We note that H is continuous on $\mathbb{D}$, and we are done.

We can view the construction of this homotopy to a constant function as selecting a fixed center and moving all points in the space toward the center along the path of least distance. Using this interpretation, it becomes clear that we can extend this result to all convex subsets of $\mathbb{R}^2$. In fact, we can pick any point within the convex subset to be the point of contraction.

Example 2.4. The unit circle $\mathbb{S}^1$ is not contractible. Suppose $\mathbb{S}^1$ is contractible. Then there exists a continuous homotopy $H : \mathbb{S}^1 \times [0, 1] \to \mathbb{S}^1$ such that

$$H((\cos s, \sin s), 0) = (\cos s, \sin s) \text{ and } H((\cos s, \sin s), 1) = (\cos s_0, \sin s_0).$$

Define $\tilde{H} : \mathbb{D} \to \mathbb{S}^1$ such that

$$\tilde{H}(t \cos s, t \sin s) = H((\cos s, \sin s), 1 - t).$$

We note that every point in $\mathbb{D}$ except $(0, 0)$ can be uniquely represented in the form $(t \cos s, t \sin s)$, for t in $(0, 1]$ and s in $[0, 2\pi]$. Also, at $t = 0$, we have $\tilde{H}(0) = H((\cos s, \sin s), 1) = (\cos s_0, \sin s_0)$, which is independent of the particular value of s we choose. Thus our function $\tilde{H}$ is well defined, and

$$\tilde{H}(\cos s, \sin s) = (\cos s, \sin s).$$

Thus we have a continuous $\tilde{H} : \mathbb{D} \to \mathbb{S}^1$ such that for all x in $\mathbb{S}^1$, we know $\tilde{H}(x) = x$, i.e. there exists a continuous function from the disk to the circle that also fixes the boundary of the disk.

Divide $\mathbb{S}^1$ into three congruent arcs A, B and C. Because $\tilde{H}$ is continuous, it can be shown that there exists a δ such that there are no a, b, c satisfying $\tilde{H}(a) \in A, \tilde{H}(b) \in B, \tilde{H}(c) \in C$ and

$$|a - b|, |b - c|, |c - a| < \delta.$$

Tile the disk $\mathbb{D}$ with triangles whose side lengths are all less than δ. At the boundary of the disk, let arcs of the unit circle be the third "sides" to the tiling triangles. For each triangle, color red the vertices that $\tilde{H}$ maps to points in A, blue the vertices that $\tilde{H}$ maps to points in B, and green the vertices that $\tilde{H}$ maps to points in C. By the above claim, since the side lengths are less than δ, no triangle should have three different colored vertices. Let edges that lie between vertices of the same color have the value 0, and let edges that lie between vertices of different colors have the value 1.

Let the value of a triangle be the sum of its three edges. We note that for any triangle, if all three vertices are different, the sum of the edges of the triangle is $3 \equiv 1 \pmod 2$. If there are two of the same color, then the sum of the edges is $2 \equiv 0 \pmod 2$, and if there are three of the same color, then the sum of the edges is also $0 \pmod 2$. Therefore, since we have no triangles all three different colors, the total sum of all the triangles in our tiling should be $0 \pmod 2$.

We note that in our sum of the triangle values of the tiling, each edge that is not an arc of the unit circle, is double counted, thus their sum does not contribute to the total sum $\pmod 2$. The edges on the boundary of $\mathbb{D}$ are single-counted, however, and thus they do contribute to our sum. By our definition of $\tilde{H}$, the points on the boundary of the disk map to themselves, thus all of the edges on the boundary have value 0, except for the three edges that contain the intersections of the three arcs, which have value 1. Thus the boundary edges contribute a total of $3 \equiv 1 \pmod 2$ to our sum. Thus our total should be $\equiv 1 \pmod 2$. This is a contradiction, thus $\mathbb{S}^1$ is not contractible. $\square$

Definition 2.5. A subspace A of X is a deformation retract of X if there exists a function $r : X \to A$ such that for all a in A, $r(a) = a$, and r is homotopic to the identity function on X.

Alternatively, we can say that a subspace A is a deformation retract of X if there exists an $H : X \times [0,1] \to X$ such that for all x in X and for all a in A, we have $H(x,0) = x, H(x,1) \in A$, and $H(a,1) = a$.

Definition 2.6. Let X be a topological space. Then we define the cone of X to be $C(X) = X \times [0,1]/(x,1) \sim (y,1)$.

We note that $\mathbb{D}$ is the cone of $\mathbb{S}^1$. The argument used in the proof that $\mathbb{S}^1$ is not contractible can be generalized to the following: a space X is contractible only if it is a deformation retract of $C(X)$.

Lemma 2.7. *If A is a deformation retract of X and A is contractible, then X must also be contractible.*

Proof. Since A is a deformation retract of X, we know there exists a homotopy $F : X \times [0,1] \to X$ from the identity function on X to a function $r : X \to A$. Moreover, since A is contractible, there exists a homotopy $G : A \times [0,1] \to A$ from the identity function on A to a constant function valued at some point a_0 in A. Let

$$H(x,t) = \begin{cases} F(x,2t) & if \quad 0 \le t \le 1/2 \\ G(r(x), 2t-1) & if \quad 1/2 \le t \le 1 \end{cases}$$

Due to the fact that $F(x,1) = r(x) = G(f(x),0)$ for all x, the two continuous functions F and G glue together, and thus form a continuous function. This is because $X \times [0,1/2]$ and $X \times [1/2,1]$ are both closed sets, and gluing functions along closed sets preserves continuity.

We note $H(x,0) = F(x,0) = x$, and $H(x,1) = G(r(x),1) = a_0$, thus H is a homotopy from the identity function on X to the constant function to a_0. $\square$

3. Graphs

Within this paper, we will want to consider graphs and other graph-like objects, namely simplicial complexes, as topological spaces. These objects are combinatorial in nature, however, thus in order to treat them as topological spaces, we will give a method of embedding them into $\mathbb{R}^n$.

Definition 3.1. A graph G is a collection of sets of pairs of elements. Each set A in G is known as an edge. A vertex of G is any element of A, where A is a set in G.

Usually, graphs are defined by a set of vertices V and a set of pairs of vertices, or edges, E. We present the above nonstandard definition of graphs in order to emphasize the fact that graphs are special cases of simplicial complexes, on which we will further elaborate later in this paper.

Although some graphs cannot be embedded into $\mathbb{R}^2$, i.e. they cannot be drawn in two dimensions in such a way that no edges intersect, all graphs can be embedded in $\mathbb{R}^3$. One way to achieve this is to represent all of the vertices of the graph as distinct points on the x-axis. For each edge that is in the graph, draw a semicircle between the corresponding vertices on the x-axis, and rotate the semicircle about the x-axis by some angle α_i. As long as all of the α_i are different, then the edges of our embedding will never intersect except at the vertices, as desired.

Definition 3.2. Let T be a finite tree. We define T' to be the tree obtained by removing all vertices of degree one and their corresponding edges. These vertices and their edges are the leaves of T.

Lemma 3.3. *If a graph is a finite tree, then it is contractible.*

Proof. Let T be a finite tree. Consider the tree T'. It is clear that T' is a deformation retract of T, since by definition, each leaf of T has only edge, thus we can always construct a homotopy such that all vertices and edges shared by T and T' remain constant, and all points on the leaves of T move at a constant rate along their only edges to their only adjacent vertices.

Assume for the sake of strong induction that all trees with k or less vertices are contractible. Consider a tree T with $k+1$ vertices. Construct the tree T' by removing all of T's leaves. We note that T' must have k or less vertices, thus, by the inductive hypothesis, T' is contractible. But T' is a deformation retract of T, thus, by a previous lemma, T is contractible also.

We note that in the base case of $k = 1$, the tree is simply a single vertex, which is trivially contractible. By strong induction, the lemma is thus true for all trees. $\square$

Lemma 3.4. *All automorphisms of trees have fixed points and/or fixed edges.*

Proof. Let T be a finite tree. Consider the tree T'. We note that any automorphism of a tree preserves the degree of the vertices being moved, thus leaf vertices can only be moved to other leaf vertices.

This implies that for any automorphism of T, there exists an automorphism of T' such that the vertices of T' move to the same locations as their corresponding vertices do in the original tree T.

Thus we can inductively reduce our automorphism to automorphisms of T', T'', and so on until we obtain an automorphism of a single edge or vertex. The tree is finite, so this process will necessarily end. The edge or vertex maps to itself, thus it maps to itself in the original automorphism of T. $\square$

4. Simplicial Complexes

We will now extend the previous ideas involving graphs to more general objects known as simplicial complexes. In particular, we note that graphs can be seen as simplicial complexes whose largest sets are restricted in size to two elements. For the intents of this paper, we need only consider finite simplicial complexes, thus whenever we mention simplicial complexes, finiteness can be assumed.

Definition 4.1. A simplicial complex is a non-empty collection of sets K such that if a set B is in K and A is a subset of B, then A is in K. A vertex of K is any element of A, where A is in K.

As noted in the previous section, we will often consider simplicial complexes as topological spaces. However, simplicial complexes are abstract combinatorial objects with no topology. Thus we will consider the geometric realization of simplicial complexes, defined as follows: let K be a simplicial complex with vertex set $\{v_1, \ldots, v_n\}$. Consider the standard basis $\{\vec{e}_1, \ldots, \vec{e}_n\}$ of $\mathbb{R}^n$. The geometric realization $\tilde{K}$ of K is the subset of $\mathbb{R}^n$ that is the union over all $A = \{v_{i_1}, \ldots, v_{i_k}\} \in K$ of all points

$$\{\lambda_{i_1} \cdot \vec{e}_{i_1} + \cdots + \lambda_{i_k} \cdot \vec{e}_{i_k} | \lambda_{i_1} + \cdots + \lambda_{i_k} = 1, 0 \leq \lambda_{i_1}, \ldots, \lambda_{i_k} \leq 1\}.$$

Definition 4.2. Let K be a simplicial complex, and let v be a vertex of K. Then the set K minus v, and the Link of K at v are, respectively,

$$K_v = \{A \in K | v \notin A\} \text{ and } Link_v K = \{A \in K | v \notin A, A \cup \{v\} \in K\}.$$

Example 4.3. In the specific case where the simplicial complex K is simply a graph, we note that K_v is the graph without the vertex v or any of the edges connected to v. The $Link_v K$ is the set of all vertices that are connected to the vertex v.

Later in this paper, we shall find that it is incredibly useful to know if the topology on a given simplicial complex is contractible. Although there is no algorithm to determine whether or not a simplicial complex is contractible, we do have tools at our disposal, such as the following lemma.

Lemma 4.4. *If there exists a v such that K_v and $Link_v K$ are contractible, then K is contractible.*

Proof. We first note that any simplicial complex K can be represented as the union of $K_v \cup C(Link_v K)$. If $Link_v K$ is contractible, then we know there exists a function $r : C(Link_v K) \to Link_v K$ such that r is homotopic to the identity function on $C(Link_v K)$ and $r(x) = x$ for all x in $Link_v K$. Let $\tilde{r} : K \to K_v$ be a function such that

$$\tilde{r}(x) = \begin{cases} x & if \quad x \in K_v \\ r(x) & if \quad x \in C(Link_v K). \end{cases}$$

We note that $Link_v K$ is a subset of K_v, and $r(x) = x$ for x in $Link_v K$, thus our $\tilde{r}$ is continuous. We note that $\tilde{r}$ is also homotopic to the identity function on K, thus we see that K_v is a deformation retract of K. By a previous lemma, we see that if K_v is contractible, then so is K. $\square$

Definition 4.5. The Euler characteristic of a simplicial complex K is

$$\mathcal{X}(K) = \sum\nolimits_{A \in K} (-1)^{|A|-1}.$$

Example 4.6. It is easy to see that the Euler characteristic of a single vertex is simply 1. We will note, but not prove, that given two simplicial complexes K and L, if K is homotopic to L, then $\mathcal{X}(K) = \mathcal{X}(L)$. In particular, this means that if a simplicial complex K is contractible, then $\mathcal{X}(K) = 1$.

Definition 4.7. Given simplicial complexes K and L, a simplicial map $f : K \to L$ is a map from vertices of K to vertices of L such that if A is in K, then $f(A)$ is in L.

We note that every simplicial map $K \to L$ gives a continuous map $\tilde{K} \to \tilde{L}$ from the geometric realization of K to the geometric realization of L. The converse of this statement is false. However, it is true in a loose sense if we consider equivalence of maps under homotopy.

Definition 4.8. Given a simplicial complex K and a simplicial map $f : K \to K$, we define a new simplicial complex $Fix(f, K)$ whose simplices are unions of the orbits of f that are also simplices in K.

The key thing to note about the $Fix(f, K)$ is that it's geometric realization is essentially the collection of fixed points of the geometric realization of K under f.

Theorem 4.9. *If K is contractible, then for all simplicial maps $f : K \to K$, it is true that $\mathcal{X}(Fix(f, K)) = 1$.*

Corollary 4.10. *If K is contractible, then every automorphism f of K has a fixed point. As noted before, the geometric realization of $Fix(f, K)$ is essentially the collection of all fixed points of the geometric realization of K under f. But since K is contractible, we know that $\mathcal{X}(Fix(f, K)) = 1$ which implies that the simplicial complex $Fix(f, K)$ is non-empty, thus f must have a fixed point.*

Corollary 4.11. *Let K be a contractible simplicial complex with automorphism $f : K \to K$ which is transitive on vertices. Then K is a simplex. Since f is transitive on vertices, the only possible vertex of $Fix(f, K)$ is the orbit of f containing all vertices of K. But $\mathcal{X}(Fix(f, K)) = 1 \neq 0$ thus the orbit of all vertices must be a vertex of $Fix(f, K)$. Thus all of K's vertices are contained in one simplex.*

5. Simplicial Complexes from Monotone Graph Properties

The goal of this section is to prove the bipartite case of the Aanderaa-Karp-Rosenburg conjecture. The primary trick in doing so is to note that given any monotone graph property, we are always able to construct a simplicial complex corresponding to that property. In doing so, we translate the combinatorial problem of determining the evasiveness of a property into a topological problem of determining whether a given space is contractible.

Definition 5.1. Let G be a graph with a set of unknown possible edges E. Let P be a monotone graph property. We define the simplicial complex representation of P, which we shall denote by $K(P, E)$, as following: a subset S of E is in $K(P, E)$ if the graph with S's edges does not satisfy P.

We note that $K(P, E)$ is indeed a simplicial complex due to the fact that P is monotone. Thus if S is in $K(P, E)$, then any subset of S must also be in $K(P, E)$.

Theorem 5.2. *If P is non-evasive then $K(P, E)$ is contractible.*

Proof. Assume, for the sake of strong induction, that all simplicial complexes with k or less vertices of non-evasive graph properties are contractible. Consider a simplicial complex K with $k+1$ vertices constructed from a graph G with unknown edges E and a non-evasive graph property P.

We first note that each time we query and receive an answer, we essentially reset the game with a new simplicial complex. In other words, we can combine any graph property and known sets of present and non-present edges into a new game with property P' and graph G' which is equivalent to playing the game with P and G knowing that certain edges are or are not present.

Pick a vertex v in K and query for it. Since P is non-evasive, the new property P' we obtain is still non-evasive, regardless of what answer we receive. If the edge v is present in the graph, then the new simplicial complex we must consider is $Link_v K$. Otherwise, the edge v is not in the graph, and we are left with K_v. In either case, the new simplicial complex has necessarily k or less vertices. By our inductive hypothesis, both $Link_v K$ and K_v are contractible, which implies, by a previous lemma, that the entire simplicial complex K must also be contractible. $\square$

Theorem 5.3. *Every bipartite graph property is evasive.*

Proof. Let G be a bipartite graph with partitions $A = \{a_1, \ldots, a_n\}, B = \{b_1, \ldots, b_m\}$, and unknown edges E. Suppose a graph property P is non-evasive. Construct the simplicial complex $K = K(P, E)$.

Let $f : K \to K$ be a simplicial map such that $f(a_i b_j) = a_i b_{j+1}$. For any i such that $1 \le i \le n$, the sequence of edges $a_i b_1, \ldots, a_i b_m$ is an orbit of f. If for a particular i_0, the set $\{a_{i_0} b_1, \ldots, a_{i_0} b_m\}$ is a simplex of the complex K, then the set $\{a_i b_1, \ldots, a_i b_m\}$ must also be a simplex of K, since the associated graphs must all be isomorphic to each other regardless of what value of i we choose, thus if one of the graphs did not satisfy P, then they would all not satisfy P. Therefore if the set $\{a_i b_1, \ldots, a_i b_m\}$ is a vertex of $Fix(f, K)$ for $i = i_0$, then it is a vertex for all i.

This result can be generalized as follows: we note that any orbit of f is entirely defined by an initial set of edges $\{a_{i_1} b_{j_1}, \ldots, a_{i_k} b_{j_k}\}$. Given this set of edges, the orbit is obtained by simply iteratively adding 1 to all of the indices of the b vertices, yielding $\bigcup_{i=1}^{n} \{a_{i_1} b_{j_1+i}, \ldots, a_{i_k} b_{j_k+i}\}$. Thus let us define the degree of an orbit by

the number of edges in its initial set of edges. We note that by this definition, the orbits mentioned in the previous paragraph are the degree 1 orbits of f.

By similar logic to before, if one orbit of degree r is in $Fix(f, K)$, then all orbits of degree r must be in $Fix(f, K)$ as well, since their corresponding graphs are all isomorphic to each other. In fact, because P is a monotone property, all orbits of degree $\leq r$ must also be in $Fix(f, K)$, since if the graph of a degree r orbit did not satisfy P, any subgraph of that graph cannot satisfy P either.

Let r be the largest degree of an orbit. We note that the Euler characteristic of $Fix(f, K)$ is

$$\mathcal{X}(Fix(f, K)) = n - \binom{n}{2} + \binom{n}{3} - \cdots + (-1)^{r-1} \binom{n}{r}.$$

By initial assumption, P is non-evasive, thus $\mathcal{X}(Fix(f, K)) = 1$. Using basic properties of binomial coefficients, we have

$$1 - n + \binom{n}{2} - \binom{n}{3} + \cdots + (-1)^r \binom{n}{r} = 0$$

$$\implies 1 - \left[1 + \binom{n-1}{1}\right] + \left[\binom{n-1}{1} + \binom{n-1}{2}\right] - \cdots + (-1)^r \left[\binom{n-1}{r-1} + \binom{n-1}{r}\right] = 0$$

$$\implies (-1)^r \binom{n-1}{r} = 0.$$

But this is impossible for $r < n$, thus we know $r = n$. But if this is true, then the orbit of f containing all vertices of K is also a simplex of K, i.e. the set of all edges in G is in K, which implies that the complete graph does not satisfy P. But since the property is monotone, this means that no graph satisfies P, which means that P can only be the trivial graph property. Therefore, all non-trivial, monotone bipartite graph properties are evasive. $\square$

6. Acknowledgments

It is a pleasure and an honor to thank my mentor, Carlos "Charly" di Fiore, for his great contributions to my understanding and writing of this paper. Throughout the entire process of finding a topic, learning the topic, and conveying the ideas I learned, Charly has been there for me every step of the way. I am incredibly grateful for his time, commitment, and patience during this experience.

7. References

"Aanderaa-Karp-Rosenberg Conjecture." Wikipedia 11 Feb. 2015. Web. 14 Aug. 2015.

Kahn, Jeff, Michael Saks, and Dean Sturtevant. A Topological Approach to Evasiveness. Combinatorica, 1984. Web.

Lovasz, Laszlo, and Neal Young. "Lecture Notes on Evasiveness of Graph Properties." arXiv.org. N.p., 18 May 2002. Web. 14 Aug. 2015.