Topology and Evasiveness

April 2, 2022

CS 6820 Analysis of Algorithms
Kimball Strong (ks2424) and Chase Vogeli (cpv29)

Contents

1 Query complexity of graph properties

Let V be a finite set of cardinality n.

Definition 1.1. A *graph property* $\mathcal{G}$ on V is a collection of simple undirected graphs with vertex set V which is closed under isomorphism.

We are interested in measuring the difficulty of determining whether a given graph G belongs to $\mathcal{G}$ in terms of the number of edges of the graph that need to be checked. That is, we imagine the edge set of G is initially unknown and it is only possible to gain information by making edge queries of the form "is $e = \{u, v\}$ an edge of G?" for $u, v \in V$ to an oracle who has access to the entire graph G.

Definition 1.2. The *query complexity* of a graph property $\mathcal{G}$ is the minimum k such that for all graphs G, it is possible to decide whether $G \in \mathcal{G}$ in at most k edge queries. A graph property $\mathcal{G}$ on vertex set V is *evasive* if the query complexity of $\mathcal{G}$ is $\binom{n}{2}$.

A useful conceptual formulation is to imagine a game being played between the player who is trying to determine whether a $G \in \mathcal{G}$ for an unknown graph G and the oracle who answers the player's edge queries. The player wins if they can determine whether $G \in \mathcal{G}$ in at most $\binom{n}{2} - 1$ edge queries, and the oracle wins otherwise. The property $\mathcal{G}$ is then evasive if the oracle has a winning strategy and nonevasive if the player has a winning strategy.

Example 1.3. The most basic example of an evasive graph property is the property of having a nonempty edge set. To see nonemptiness is evasive, observe that the oracle has a winning strategy by answering "no" to the first $\binom{n}{2} - 1$ edge queries. The player cannot determine whether G is nonempty with this information alone, since whether G is nonempty depends on whether the final unqueried pair of vertices is adjacent. Thus, the player is forced to make the final edge query, and the oracle wins.

A useful tool in proving the evasiveness of graph properties is to design an adversarial strategy for the oracle. As seen in the example above, the oracle need not fix a particular graph in advance. Indeed, a winning oracle strategy often involves building the graph as the edge queries come in as to delay the player's determination of whether $G \in \mathcal{G}$ as long as possible. A successful oracle will be able to maintain the ambiguity of whether $G \in \mathcal{G}$ throughout the querying of all of G's edges. We illustrate this principle with another proof of evasiveness.

Proposition 1.4. *Connectivity is an evasive graph property.*

Proof. We describe an adversarial oracle strategy. In response to the query "is $e = \{u, v\}$ in G?," the oracle should answer "no" unless that answer would imply that there exists no path from u to v in G. That is, if adding all the unqueried edges to G results in there being a path from u to v, the oracle answers "no," and "yes" otherwise.

To see that this is a winning strategy for the oracle, suppose that the player can determine whether G is connected without querying all edges. Let $e = \{u, v\}$ be an unqueried edge. Since the oracle strategy preserves the possibility that G is connected after each query, the player must have concluded that G is connected, and there must exist a path from u to v consisting of queried edges. Along with e, this path forms a cycle C in G. Let $e' = \{u', v'\}$ be the most recently queried edge on the cycle C. That the oracle answered "yes" to the query for e' implies that e' being in G was necessary for there to be a path from u to v. However, there exists another path from u to v along the cycle C, which is a contradiction. $\square$

Many other natural graph properties can be shown to be evasive by a cleverly designed oracle strategy. One particularly intricate example is a strategy due to Bollobás [1] which establishes the evasiveness of containing a k-clique for $k \geq 2$. Indeed, it is rather difficult to generate examples of graph properties that are not evasive, which led Rosenberg [6] to conjecture that all graph properties have $\Omega(n^2)$ query complexity. This was however disproven by Aanderaa's discovery of graph properties with linear query complexity. One such example of a graph property with linear query complexity is that of being a *scorpion*, which we discuss here.

Definition 1.5. A *scorpion* is a graph with three distinguished vertices, the *sting*, the *tail*, and the *body*, subject to the following conditions.

- The only neighbor of the sting is the tail.
- The only neighbors of the tail are the sting and the body.
- The body is adjacent to every vertex except for the sting.

No conditions are imposed on the adjacency of the remaining vertices, called the *feet* of the scorpion.

It turns out that the query complexity of the scorpion property is $O(n)$, which implies that being a scorpion is nonevasive for sufficiently large n. We prove this fact here to illustrate an example of a winning strategy against the oracle.

Proposition 1.6 ([2, Homework 2 #3]). *The query complexity of deciding whether a graph is a scorpion is at most* $5n$*.*

Proof. We describe a strategy to determine whether a graph is a scorpion in at most $5n$ edge queries. We begin by selecting an arbitrary vertex $v \in V$ and split into cases based on the degree d_v of v. The degree of a vertex can be determined by querying all of its possible edges, which takes at most n many queries.

- If $d_v = 0$ or $d_v = n-1$, then G cannot be a scorpion. In this case, only the $n-1$ many vertex queries to determine d_v were necessary.

- If $d_v = n-2$, then v is the body if G is a scorpion. To verify G is a scorpion, we proceed as follows. Let $u \in V$ be the unique vertex which is not adjacent to v, which must be the sting if G is a scorpion. With at most n many edge queries, we can verify that $d_u = 1$ and if so, locate its sole neighbor $w \in V$, which must be the tail if G is a scorpion. Finally, in at most n many edge queries, we can verify that w is the tail by checking whether $d_w = 2$. In this case, a total of at most $3n$ edge queries are made.

- If $d_v = 1$, then either v is the sting or a foot if G is a scorpion. Let $u \in V$ be the sole neighbor of v, and we can determine d_u in at most n edge queries. If G is a scorpion, then either u is the body or the tail, and we split into cases based on d_u.

  - If $d_u = 2$, then v was the sting and u is the tail. In at most n more edge queries, we can verify that the other neighbor of u is the body by checking whether it has degree $n-2$.
  - If $d_u = n-2$, then u is the body if G is a scorpion. As above, once we have found the body, we can locate and sting and tail and verify G is a scorpion in at most $2n$ more edge queries.

  In total, this case requires at most $4n$ edge queries.

- If $d_v = 2$, then either v is the tail or a foot if G is a scorpion. In at most $2n$ edge queries, we can determine the degrees of the two neighbors of v. If G is a scorpion, then one of its neighbors $u \in V$ must be the body. There are two possibilities for the other neighbor, $w \in V$.

  - If $d_w = 1$, then w is the sting and v was the tail of the scorpion.
  - If $d_w > 1$, then both v and w are feet. As above, starting from the body $u \in V$, we can locate the sting and tail in at most $2n$ more edge queries.

  In total, this case requires at most $5n$ edge queries.

- If $3 \leq d_v \leq n-3$, then v is a foot if G is a scorpion. We begin by describing a process to locate the sting in at most n edge queries. Let B be the set of neighbors of v and S be the set of non-neighbors. If G is a scorpion, then B must contain the body and S must contain the sting S.

  Start with an arbitrary $u \in B$ and $w \in S$ and repeat the following process. Query the pair $\{u, w\}$ to determine whether they are adjacent. If u and w are adjacent, then v

cannot be the sting, so we delete w from S and pick a new $w \in S$. If u and w are not adjacent, then u cannot be the body unless w is the sting, so we delete u from B and pick a new $u \in B$. If G is a scorpion, then the process terminates with B having been emptied and $w \in S$ the sting. Indeed, the body in B cannot be deleted by any vertex other than the sting in S, and once the sting is chosen from S, B will be emptied.

The above procedure locates the sting in at most n edge queries, since one vertex is deleted after each query. As above, once the sting is located, the tail and body can be located in at most $2n$ more edge queries. In total, this case requires at most $4n$ edge queries.

In the worst case, the above strategy uses at most $5n$ edge queries to determine whether G is a scorpion. $\square$

As a result of the existence of such properties with linear query complexity, the above conjecture was amended to include the condition that the graph property is *monotone*.

Definition 1.7. A graph property $\mathcal{G}$ is *monotone* if $\mathcal{G}$ is closed under adding edges.

Example 1.8. Let $\mathcal{G}_{d=3}$ be the collection of cubic graphs, that is, graphs in which every vertex has three neighbors. Then, $\mathcal{G}_{d=3}$ is not a monotone graph property. Indeed, adding any edges to a cubic graph will result in some vertices having more than three neighbors. On the other hand, the collection $\mathcal{G}_{d\geq 3}$ of graphs in which each vertex has at least three neighbors is monotone.

This amended conjecture was first proven by Rivest and Vullemin [5] who showed that the query complexity of a monotone graph property is at least $n^2/16$. The following stronger conjecture has come to be known as the Aanderaa–Karp–Rosenberg conjecture.

Conjecture 1.9 ([3, Conjecture 2])**.** *Every nontrivial monotone graph property is evasive.*

The Aanderaa–Karp–Rosenberg conjecture in full generality is unresolved, but the following special was proven by Kahn–Saks–Sturtevant [3] as a striking application of topological methods in solving combinatorial problems. Our goal in these notes is the explain these topological methods.

Theorem 1.10 ([3, Theorem 1])**.** *Every nontrivial monotone graph property on a prime power number of vertices is evasive.*

2 Connection to topology

In this section, we recast the above situation into the language of topology using simplicial complexes. Simplicial complexes are discrete models for topological spaces particularly amenable to analysis via combinatorial methods. They should be thought of as being built up out of simplices of various dimensions fitting together in a nice way.

Definition 2.1. An *(abstract) simplicial complex* Δ on a set $\mathcal{X}$ is a collection of subsets of X which is closed under taking subsets, i.e. if $X \in \Delta$ then all the subsets of X are also contained in Δ. The elements of Δ of cardinality $k+1$ are referred to as *k-faces* or *k-simplices.*

Monotone graph properties give rise to simplicial complexes in the following way. Let $\mathcal{G}$ be a monotone graph property on a vertex set V. Since a graph with vertex set V is specified by its edge set, $\mathcal{G}$ can be viewed as a collection of subsets of the set of all possible edges

$$\mathcal{E} = \{\{u, v\} : u, v \in V\}.$$

The monotonicity of $\mathcal{G}$ implies that $\mathcal{G}$ is closed under taking supersets. Thus, taking complements yields a collection of subsets which is closed under taking subsets. Specifically, let $\Delta_{\mathcal{G}}$ be the collection of subsets of $\mathcal{E}$ given by

$$\Delta_{\mathcal{G}} = \{\mathcal{E} \setminus E : E \in \mathcal{G}\},$$

which is a simplicial complex on $\mathcal{E}$ for the reason above.

The simplicial complex $\Delta_{\mathcal{G}}$ of a monotone graph property $\mathcal{G}$ has additional structure arising from the symmetry of graph properties. Recall that graph properties are required to be isomorphism-invariant. That is, if a graph G on vertex set V has property $\mathcal{G}$, then so does the graph obtained from G by relabeling the vertices according to some permutation of the vertex set V. This endows $\mathcal{G}$ and thus $\Delta_{\mathcal{G}}$ with a transitive action of the symmetric group S_V on V. This action will play an important role when we discuss fixed-point theory and its applications to evasiveness.

2.1 Collapsibility of nonevasive complexes

Definition 2.2. Let Δ be a simplicial complex on $\mathcal{X}$. A face $X \in \Delta$ is *free* if there exists a unique maximal face of Δ containing X. An *elementary collapse* of Δ at a free face X removes X along with all faces containing X from Δ. The simplicial complex Δ is *collapsible* if there is a sequence of elementary collapses which takes Δ to the empty simplicial complex.

It turns out that if a graph property $\mathcal{G}$ is not evasive, the existence of a good querying strategy will yield to a sequence of elementary collapses of $\Delta_{\mathcal{G}}$ that fully collapse $\Delta_{\mathcal{G}}$. The goal of this section is to prove this result. It will be more convenient to generalize this result to general simplicial complexes however, not just those that arise from monotone graph properties.

Definition 2.3. A simplicial complex Δ on the set $\mathcal{X}$ is *evasive* if it is not possible to decide whether $X \in \Delta$ for all subsets $X \subseteq \mathcal{X}$ in less than $|\mathcal{X}|$ queries of the form "is $x \in X$?" for $x \in \mathcal{X}$.

Proposition 2.4. *A nonevasive simplicial complex is collapsible.*

The following lemma will aid in the proof, but we must first fix some notation for two simplicial complexes associated to a given simplicial complex, called the link and contrastar. Given a simplicial complex Δ on $\mathcal{X}$ and a point $x \in \mathcal{X}$, the *link* of Δ at x is the simplicial complex

$$L(\Delta, x) = \{X \subseteq \mathcal{X} \setminus \{x\} : X \cup \{x\} \in \Delta\}$$

on the set $\mathcal{X} \setminus \{x\}$. Intuitively, the faces of the link are those that "link" with x to form faces of Δ. Similarly, the *contrastar* of Δ at x is the simplicial complex

$$C(\Delta, x) = \{X \subseteq \mathcal{X} \setminus \{x\} : X \in \Delta\}$$

on the set $\mathcal{X} \setminus \{x\}$. That is, the contrastar consists of faces of Δ that do not contain the point x. We are now ready to state the lemma.

Lemma 2.5. *Let Δ be a simplicial complex on $\mathcal{X}$. If there exists a point $x \in \mathcal{X}$ such that the link $L(\Delta, x)$ and contrastar $C(\Delta, x)$ are collapsible, then Δ is collapsible.*

Proof. That $L(\Delta, x)$ is collapsible gives the existence of a sequence of faces $X_1, \ldots, X_r$ in $L(\Delta, x)$ which collapses $L(\Delta, x)$. That is, X_i is a free face of $L(\Delta, x)$ after the collapses at $X_1, \ldots, X_{i-1}$ have been preformed, and all the collapses at $X_1, \ldots, X_r$ together take $L(\Delta, x)$ to a one-point simplicial complex. The sequence $X_1 \cup \{x\}, \ldots, X_r \cup \{x\}, \{x\}$ of faces of Δ collapses Δ to $C(\Delta, x) \subseteq \Delta$.

That $C(\Delta, x)$ is collapsible gives a sequence $Y_1, \ldots, Y_s$ of faces of $C(\Delta, x)$ whose collapses collapse $C(\Delta, x)$ to a point. Concatenating these sequences as

$$X_1 \cup \{x\}, \ldots, X_r \cup \{x\}, \{x\}, Y_1, \ldots, Y_s$$

gives a sequence of collapses which takes Δ to a point. Thus, Δ is collapsible. □

Proof of Proposition 2.4. Let Δ be a nonevasive graph property on the set $\mathcal{X}$. We induct on $|\mathcal{X}|$. In the base case $|\mathcal{X}| = 1$, any nonempty simplicial complex on a set with one element is collapsible.

Suppose $|\mathcal{X}| > 1$. Since Δ is nonevasive, there exists a querying strategy which determines whether a subset $X \subseteq \mathcal{X}$ is a face in strictly less than $|\mathcal{X}|$ many queries in the worst case. Let $x \in \mathcal{X}$ be the first point queried in such a strategy. If the query returns $x \in X$, then X is a face if $X \setminus \{x\} \in L(\Delta, x)$. If the query returns $x \notin X$, then X is a face if $X \in C(\Delta, x)$. Both $L(\Delta, x)$ and $C(\Delta, x)$ are simplicial complexes on a set of cardinality $|\mathcal{X}| - 1$ for which there exists a querying strategy that decides membership in strictly less than $|\mathcal{X}| - 1$ many queries. The inductive hypothesis implies that $L(\Delta, x)$ and $C(\Delta, x)$ are both collapsible. By the previous lemma, Δ is collapsible as well. □

Corollary 2.6. *If $\mathcal{G}$ is a nonevasive graph property, then $\Delta_{\mathcal{G}}$ is collapsible.*

2.2 Simplicial homology

In this section, we explore what concequences the collapsability of a simplicial complex has on one of its most important algebraic invariants, its *homology*. For completeness sake, we briefly review the construction of simplicial homology here.

In what follows, we will implicitly fix a prime number p and work over the finite field $\mathbb{F}_p$ of cardinality p. For much of the purely topological results in this section, this choice of prime number is arbitrary, but in our desired application, p will be such that the number of vertices of our graph property will be a power of p.

Let Δ be a simplicial complex on the set $\mathcal{X}$. An *orientation* of a face $X \in \Delta$ is an equivalence class of orderings of the set X which is closed under even permutations. This implies each simplex has two possible orientations. If we fix an ordering of the vertex set $\mathcal{X}$ of Δ, this induces an ordering of each face, and thus a choice of orientation for each face.

Denote by Δ_k the set of k-faces of Δ, which we recall are those of cardinality $k + 1$. A k-chain is a formal linear combination of elements of Δ_k with coefficients in the field $\mathbb{F}_p$.

More formally, the *group of k-chains* Δ is $C_k(\Delta) = \mathbb{F}_p\Delta_k$, the free $\mathbb{F}_p$-vector space with basis Δ_k. We think of each face $X \in \Delta_k$ as coming with a prescribed orientation, and declare that X with the reverse orientation is equal to $-X$ in the group $C_k(\Delta)$.

We define a collection of homomorphisms $d_i : C_k(\Delta) \to C_{k-1}(\Delta)$ for $0 \leq i \leq k$ as follows. For an oriented simplex $X = (x_0, \ldots, x_k) \in \Delta_k$, we set

$$d_i(X) = (x_0, \ldots, x_{i-1}, x_{i+1}, \ldots, x_k) \in C_{k-1}(\Delta)$$

to be the oriented simplex obtained from X by deleting the ith vertex. We extend d_i to all of $C_k(\Delta)$ by linearity. Now, we define the *boundary map* $d : C_k(\Delta) \to C_{k-1}(\Delta)$ by setting

$$d(X) = \sum_{i=0}^{k} (-1)^i d_i(X)$$

for k-simplices $X \in \Delta_k$ and extending to all of $C_k(\Delta)$ by linearity.

We now define two subgroups of the group of k-chains as follows. A k-chain $c \in C_k(\Delta)$ is said to be a *k-cycle* if $d(c) = 0$ or in other words, if c "has no boundary." A k-chain $c \in C_k(\Delta)$ is a *k-boundary* if there is a $c' \in C_{k+1}(\Delta)$ such that $d(c') = c$. Accordingly, we define the *group of k-cycles* to be

$$Z_k(\Delta) = \ker\big(d : C_k(\Delta) \to C_{k-1}(\Delta)\big).$$

and the *group of k-boundaries* to be

$$B_k(\Delta) = \operatorname{im}\big(d : C_{k+1}(\Delta) \to C_k(\Delta)\big).$$

An important property[1] of the boundary map is $d^2 = 0$, which implies that

$$B_k(\Delta) \subseteq Z_k(\Delta),$$

or in other words, boundaries have no boundary themselves. This inclusion enables us to define the *kth homology group of* Δ as the quotient

$$H_k(\Delta) = Z_k(\Delta)/B_k(\Delta).$$

Another way to view this group is as follows. Say two k-cycles $c, c' \in Z_k(\Delta)$ are *homologous* if their difference is a boundary, that is $c - c' \in B_k(\Delta)$. The homology group $H_k(\Delta)$ is then the group of *homology classes* of k-cycles.

We make one final definition before returning to the proofs at hand. The simplicial complex Δ is *acyclic* if

$$H_k(\Delta) \cong \begin{cases} \mathbb{F}_p & k = 0, \\ 0 & k \neq 0. \end{cases}$$

These homology groups are the simplest possible that a nonempty simplicial complex can have. Indeed, we will soon see that these homology groups are the homology groups of a simplicial complex that consists of a single 0-simplex.

Our goal in this section is to show that a collapsible simplicial complex is acyclic. As a start, we examine how elementary collapses affect homology.

[1] Indeed, Henri Cartan is famously quoted as saying: "If I could only understand the beautiful consequences following from the concise proposition $d^2 = 0$."

Proposition 2.7. *The homology of a simplicial complex is preserved under elementary collapses.*

The proof of this proposition will follow immediately from the following two lemmas. In the proof, it will be useful to introduce a way to decompose elementary collapses into pieces which are easier to analyze with respect to homology. Call the elementary collapse induced by a free face X contained in the unique maximal face Y a *codimension one collapse* if $\dim Y - \dim X = 1$.

Lemma 2.8. *Any elementary collapse can be written as a sequence of codimension one collapses.*

Proof. Let Δ be a simplicial complex, and let X be a free face of Δ contained in the unique maximal face Y. We induct on the codimension $\ell = \dim Y - \dim X$ of X in Y. In the base case $\ell = 1$, the collapse at X is already codimension one.

Suppose $\ell > 1$. Pick a maximal proper face $Y' \subseteq Y$, and set $X' = X \cup (Y \setminus Y')$. Note that by construction, X' is uniquely contained in the maximal face Y, and that $\dim Y - \dim X' = \ell - 1$, so by the inductive hypothesis, the collapse at the face X' can be written as a sequence of codimension one collapses. After the collapse at X', the face Y' is maximal and is the unique maximal face containing X. Furthermore, we have by construction that $\dim Y' - \dim X = \ell - 1$, so by the inductive hypothesis, the collapse at X can be written as a sequence of codimension one collapses. Observe that collapsing at X' with maximal face Y followed by collapsing at X with maximal face Y' is equivalent to collapsing at X with maximal face Y. We conclude that the collapse at X with maximal face Y can be written as a sequence of codimension one collapses. $\square$

Lemma 2.9. *Homology is preserved under codimension one collapses.*

Proof. Let Δ be a simplicial complex, and let X be a free face contained in the unique maximal face Y. Furthermore, suppose that X is codimension one in Y and that $\dim Y = k$. Let $\Delta' = \Delta \setminus \{X, Y\}$ be the simplicial complex obtained by preforming the collapse at X.

We show the homology groups of Δ' are equal to those of Δ in each dimension. Since Y is a maximal face of Δ, Y is not the face of any higher-dimensional simplex, and thus removing Y from Δ will not affect the homology in dimensions greater than k. In dimension k, observe that $Z_k(\Delta') = Z_k(\Delta)$ since Y is not a k-cycle in Δ and that that $B_k(\Delta') = B_k(\Delta)$ since $C_{k+1}(\Delta') = C_{k+1}(\Delta)$. Therefore, $H_k(\Delta') = H_k(\Delta)$. In dimension $k-1$, removing Y will remove the subspace of $B_k(\Delta)$ spanned by $d(Y)$, however this subspace will also be removed from $Z_k(\Delta)$, as X is a summand of $d(Y)$. Therefore,

$$H_{k-1}(\Delta') = Z_{k-1}(\Delta')/B_{k-1}(\Delta') = Z_{k-1}(\Delta)/B_{k-1}(\Delta) = H_{k-1}(\Delta).$$

In dimension $k-2$, $Z_{k-2}(\Delta') = Z_{k-2}(\Delta')$ since the simplices in dimensions $k-2$ and lower are left unchanged. We also have that $B_{k-2}(\Delta') = B_{k-2}(\Delta)$. Indeed, if we write $X_1, \dots, X_k$ for face of Y other than X, we have that

$$d^2(Y) = d(X) + \sum_{i=1}^{k} (-1)^i d_i(X_i) = 0$$

from which it follows that $d(X) \in B_{k-2}(\Delta')$ even though X itself is removed from Δ. We conclude that $H_{k-2}(\Delta') = H_{k-2}(\Delta)$. As noted, the simplicies in dimension $k-2$ and lower are preserved, so the homology in dimensions below $k-2$ is as well. □

We are now ready to turn to the proof of the main result of this section.

Proposition 2.10. *A collapsible simplicial complex is acyclic.*

Proof. Let Δ be a collapsible simplicial complex a set $\mathcal{X}$. There is a sequence of elementary collapses that takes Δ to a point, and by Proposition 2.7, homology is preserved under elementary collapses. It follows that the homology of Δ is isomorphic to the homology of a point, so it thus suffices to show that a one-point simplicial complex is acyclic.

By a one-point simplicial complex, it is meant a simplicial complex on $\mathcal{X}$ for which the only simplex is the singleton $\{x\}$ for a point $x \in \mathcal{X}$. Denote such a one-point simplicial complex by $*$. The only nonzero chain group of $*$ is in degree zero, where we have $C_0(*) \cong \mathbb{F}_p$. Since all boundary maps are zero for $*$, we have that

$$Z_0(*) = \ker d = C_0(*) \cong \mathbb{F}_p \qquad B_0(*) = 0$$

from which it follows that $H_0(*) \cong \mathbb{F}_p$. Since $C_k(*) = 0$ for $k > 0$, we have $H_k(*) = 0$ for $k > 0$. We conclude that $*$ is acyclic, and thus Δ is acyclic. □

To summarize, we have thus far proven:

Proposition 2.11. *If $\mathcal{G}$ is a nonevasive graph property, then $\Delta_{\mathcal{G}}$ is acyclic.*

2.3 Reduction to fixed-point theory

We conclude this section by setting up the arguments involving fixed-point theory in the following section.

Definition 2.12. A group G is *p-by-cyclic*[2] if there exists a normal p-subgroup $H \lhd G$ such that the quotient G/H is cyclic.

Theorem 2.13. *If Δ is a nonempty acyclic simplicial complex which admits a transitive action by a p-by-cyclic group, then Δ is a simplex.*

We concern ourselves with the proof of this theorem in the next section. For now, we see how we can obtain our desired result using this theorem.

Proof of Theorem 1.10. Let $\mathcal{G}$ be a monotone nonevasive graph property on V where $n = |V|$ is a prime power. We will show that $\mathcal{G}$ is empty, and is therefore trivial.

Let $n = p^r$ for p a prime, and identify V with the finite field $\mathbb{F}_{p^r}$. As above, let $\mathcal{E}$ denote the set of all possible edges in a graph on the vertex set V. Observe that the group

$$\mathrm{AGL}_1(\mathbb{F}_{p^r}) = \{x \mapsto ax + b : a \in \mathbb{F}_{p^r}^\times, b \in \mathbb{F}_{p^r}\}$$

[2]This is nonstandard terminology, but is inspired by actual group properties in the literature, such as the notion of a free-by-cyclic group. If you can guess what a free-by-cyclic group is based on the definition above, you may agree that "p-by-cyclic" is an appropriate name.

of 1-dimensional affine linear transformations of $\mathbb{F}_{p^r}$ acts doubly transitively on V and thus transitively on $\mathcal{E}$. Indeed, given $\{a_1, b_1\}, \{a_2, b_2\} \in \mathcal{E}$, the affine transformation

$$\alpha(x) = \frac{b_2 - a_2}{b_1 - a_1}(x - a_1) + a_2$$

satisfies $\alpha(a_1) = a_2$ and $\alpha(b_1) = b_2$. Observe that the subgroup

$$H = \{x \mapsto x + b : b \in \mathbb{F}_{p^r}\}$$

consisting of translations in $\mathrm{AGL}_1(\mathbb{F}_{p^r})$ is a normal p-subgroup, and the quotient $\mathrm{AGL}_1(\mathbb{F}_{p^r})/H$ is exactly $\mathrm{GL}_1(\mathbb{F}_{p^r}) \cong \mathbb{F}_{p^r}^\times$, which is cyclic. It follows that $\mathrm{AGL}_1(\mathbb{F}_{p^r})$ is a p-by-cyclic group which acts transitively on $\Delta_{\mathcal{G}}$. As $\Delta_{\mathcal{G}}$ is acyclic by Proposition 2.11, Theorem 2.13 gives that $\Delta_{\mathcal{G}}$ is a simplex. We conclude that $\mathcal{G}$ is empty. □

3 Fixed-point theory

In the last section, we reduced proving the evasiveness of monotone graph properties on prime power numbers of vertices to Theorem 2.13, a purely topological result that acyclic simplicial complexes with transitive actions of p-by-cyclic groups are simplices. We reduce even further to proving the mere existence of a fixed point of said group action, that is:

Theorem 3.1. *If Δ is a nonempty acyclic simplicial complex and G a p-by-cyclic group acting on Δ, then there is simplex in Δ fixed by G.*

Before proving this theorem, we demonstrate how we can use it to prove Theorem 2.13.

Proof of Theorem 2.13. Let Δ be a simplicial complex on the base set $\mathcal{X}$, and let it be acted on transitively by the p-by-cyclic group G. Then by Theorem 3.1 there is a simplex of Δ fixed by G. By transitivity this simplex must be the entire set $\mathcal{X}$, so Δ is a simplex. □

Now, before we proceed to the proof of Theorem 3.1, we will need several definitions and lemmas. Let Δ be an acyclic simplicial complex, acted on by an automorphism T of period p. Define two endomorphisms of the chain complex of abelian groups obtained from Δ:

$$\alpha = 1 - T \qquad\qquad \beta = 1 + T + T^2 + \cdots + T^{p-1}$$

Note that $\alpha\beta = \beta\alpha = 0$

Definition 3.2. We call a chain $c \in C_n(\Delta)$ **type** α if it is in the image of α, and similarly define **type** β.

Lemma 3.3. *A chain c is of type α iff $\beta(c) = 0$ (and similarly with α and β swapped).*

We omit this lemma's proof, as it is a little tedious.

Definition 3.4. Suppose that the set of simplices fixed by T forms a sub-simplicial complex, called K. We call a boundary $c \in C_n(\Delta)$ a *typed boundary* if it is type α and, mod K,[3] is also the boundary of a chain of type α (or the analogous statement with β).

[3]By 'mod K' we mean that $c = d(\alpha) + k$ for some chain k of simplices in K.

Definition 3.5. Suppose c_n is an n-chain of type α and c_{n-1} is a chain of type β. If there exists a chain x_n of dimension n such that $c_n = \alpha(x_n)$ and $c_{n-1} = d(x_n)$, we write $c_n : c_{n-1}$. We write the same for the same statement with α and β swapped.

Lemma 3.6. *Suppose $c_n : c_{n-1}$. If c_n is a typed boundary, then so is c_{n-1}.*

Proof. Suppose the types of c_n and c_{n-1} are of types α and β, respectively. Then we have that there exist chains x_n and x_{n+1} such that, mod K:

$$c_n = \alpha(x_n) \qquad c_{n-1} = d(x_n) \qquad c_n = d(\alpha(x_{n+1}))$$

Let $z = -d(x_{n+1}) + x_n$. Then $\alpha(z) = -c_n + c_n = 0$, so z is of type β. Note $d(z) = d^2(x_{n+1}) + d(x_n) = c_{n-1}$, so c_{n-1} is a typed boundary, as desired. The proof is identical in the case that the types of c_n and c_{n-1} are swapped. $\square$

Lemma 3.7. *Let e be a vertex of Δ, and the action have no fixed points. Then $\alpha(e) \neq d(\alpha(x))$ for any chain x.*

Proof. Suppose to the contrary that such an x exists. Let $z = d(x) - e$. Then $\alpha(z) = 0$, so $z = \beta(w)$ for some w. But for any 0-simplex f, $\beta(f)$ is a cycle mod p (because it is the sum of p points mod p, and Δ is connected), so z is a cycle. It follows that $-e$ is a cycle, which is impossible. $\square$

Theorem 3.8. *If Δ is a nonempty acyclic, finite-dimensional simplicial complex acted on by an automorphism T of period p, then the simplicial complex of fixed points of this action is nonempty.*

Proof of Theorem 3.8. Suppose that there are no fixed points. Let x_0 be a vertex in Δ and define $c_0 := \alpha(x_0) = x_0 - T(x_0)$. By acyclicity, c_0 is a boundary. So let x_1 be a chain such that $d(x_1) = c_0$, and define $c_1 := \beta(x_1)$ Note that $d(c_1) = \beta(d(x_1)) = \beta(\alpha(x_1)) = 0$, so c_1 is a boundary. Let x_2 be such that $d(x_2) = c_1$, and define $c_2 := \alpha(x_2)$. Continuing on in this manner, we get two sequences $x_0, ..., x_d$ and $c_0, , .., c_d$ (where d is the dimension of Δ) such that

$$d(x_{i+1}) = c_i \qquad\qquad c_i = \begin{cases} \alpha(x_i) & \text{if } i \text{ is even} \\ \beta(x_i) & \text{if } i \text{ is odd} \end{cases}$$

By construction, we therefore have $c_d : c_{d-1} : \cdots : c_0$. Note that, by acyclicity, c_d must be zero, and is therefore trivially a typed boundary. It follows by Lemma 3.6 that c_0 is a typed boundary. But this contradicts Lemma 3.7. $\square$

To prove acyclicity, we will need to refine our simplicial complex Δ: note that the fixed simplices of an action on a simplicial complex can fail to form a simplicial complex. For a simple example, consider the simplicial complex on two vertices v_1 and v_2 given by $\{\{v_1, v_2\}, \{v_1\}, \{v_2\}\}$. If we consider the action of C_2 which swaps v_1 and v_2, this preserves the 2-simplex $\{v_1, v_2\}$ but it does not preserve its two boundary simplices. To rectify this, we introduce barycentric subdivision:

Definition 3.9. Let Δ be a simplicial complex. The barycentric subdivison $\mathcal{B}(\Delta)$ is the simplicial complex whose vertex set is given by the simplices of Δ, and whose n-simplices are ascending sequences of $n+1$ simplices of Δ.

This construction has the following useful property, which we do not prove (but is elementary enough to be left as an exercise):

Lemma 3.10. *If Δ is acyclic, then so is $\mathcal{B}$.*

Note furthermore that given an action on Δ, this naturally induces an action on $\mathcal{B})(\Delta)$. Furthermore, the induced action is such that the fixed simplices form a simplicial complex: Let T be an endomorphism of Δ, and denote by $T_{\mathcal{B}}$ the induced endomorphism of $\mathcal{B}(\Delta)$. Let s_n be an n-simplex of $\mathcal{B}(\Delta)$ which is fixed by $T_{\mathcal{B}}$. By definition, s_n is a set of simplices of Δ $\{t_0, t_1, ..., t_n\}$ where each t_i is in the boundary of t_{i+1}. Since s_n is fixed, we have that $\{t_0, ..., t_n\} = \{T(t_0), ..., T(t_n)\}$. But because i-simplices are not k-simplices for $i \neq k$, and because T takes i-simplices to i-simplicies, this means that $T(t_i) = t_i$ for each i, so T fixes each vertex of s_n and therefore fixes each face.

With this tool, we can now state and prove the following:

Theorem 3.11. *If Δ is a nonempty acyclic, finite-dimensional simplicial complex acted on by an automorphism T of period p, then the simplicial complex of fixed points of the induced action on $\mathcal{B}(\Delta)$ is nonempty and acyclic.*

Proof. Nonemptiness follows from the previous theorem, so we need only prove acyclicity. Let c_n be an n-cycle of the fixed point set $\mathcal{B}(\Delta)^{T_{\mathcal{B}}}$. By acyclicity of $\mathcal{B}(\Delta)$, there exists some $n+1$ chain x_{n+1} such that $d(x_{n+1}) = c_n$. As before, we construct sequences of simplices $c_d : \cdots : c_n$ and $x_d : \cdots : x_n$ and conclude that c_n is a typed boundary - in particular, it is equal to $d(\beta(x_{n+1}))$. Since the image of β consists of fixed simplices, we conclude that c_n is a boundary in the sub-simplicial complex of fixed points. □

And to conclude the section, we prove Theorem 3.1: because one of the authors got mixed up and forgot to write up a self-contained proof, we will need to take the following lemma [[4], Lemma 1] without proof:

Lemma 3.12. *Let $\mathbb{Z}/n\mathbb{Z}$ act on the acyclic complex Δ. Then this action has a fixed point.*

Proof of Theorem 3.1. Let G be a p-by-cyclic group acting on the acyclic simplicial complex Δ. Note that G fixes a simplex of Δ if and only if it fixes a simplex of $\mathcal{B}(\Delta)$, so it suffices to consider the action of G on $\mathcal{B}(\Delta)$. Let H be a normal p-subgroup of G with G/H cyclic. We prove that H fixes a simplex in $\mathcal{B}(\Delta)$: Let H be of order p^k. Pick an order p element $T \in H$ such that $\langle T \rangle$ is a normal subgroup of H. By Theorem 3.11, the fixed points of the action of this normal subgroup form a subcomplex which is again acyclic, and we now have an action of $H/\langle T \rangle$, which is of order p^{k-1} on this fixed subcomplex. Thus by inducting on the order k of H, we find that H fixes a subcomplex $K \subseteq \mathcal{B}(\Delta)$ such that K is acyclic. Note that G/H acts on on K; by the lemma above this action has a fixed point. Thus the action of G on $\mathcal{B}(\Delta)$ has a fixed point, thus G fixes a simplex of Δ. □

References

[1] Béla Bollobás. *Complete Subgraphs Are Elusive.* Journal of Combinatorial Theory, Series B **21** (1976), 1–7.

[2] Dexter Kozen. *The Design and Analysis of Algorithms.* Texts and Monographs in Computer Science (1992), Springer, New York.

[3] Jeff Kahn, Michael Saks and Dean Sturtevant. *A topological approach to evasiveness.* Combinatorica **4** (1984), 297–306.

[4] Robert Oliver *Fixed-Point Sets of Group Actions on Finite Acyclic Complexes* Commentarii mathematici Helvetici **50** (1975), 155–177.

[5] Ronald Rivest and Jean Vuillemin. *On recognizing graph properties from adjacency matrices.* Theoretical Computer Science **3** (1976), 371–384.

[6] Arnold Rosenberg. *On the time required to recognize properties of graphs: A problem.* ACM SIGACT News **5** (1973), 15—16.

[7] Paul Althaus Smith. *Fixed-Point Theorems for Periodic Transformations* American Journal of Mathematics **63** (1941), 1–8