A topological approach to evasiveness

Notes by David Galvin

February 18, 2010

Abstract

A graph property $\mathcal{G}$ is a collection of graphs closed under isomorphism. $\mathcal{G}$ is said to be *evasive* if, for every possible local search strategy, there is at least one graph for which membership in $\mathcal{G}$ cannot be decided until the entire graph has been searched.

Most graph properties that one can think of, but not all, turn out to be evasive. A conjecture of Aanderaa, Karp and Rosenberg asserts that every *monotone* (closed under adding edges) property is evasive (except for two trivial exceptions).

I'll explain how topological ideas introduced by Kahn, Saks and Sturtevant, in particular the study of fixed points of simplicial maps, have helped to make significant progress towards the AKR conjecture. These notes are based in part on the original paper of Kahn, Saks and Sturtevant [2] and in part on lecture notes of Lovász and Young [4].

1 Introduction

A *graph* G consists of a set V of *vertices* and a set E of *edges* (unordered pairs of vertices). We think of the edge $\{u, v\}$ as *joining* vertices u and v. The graphs we consider will satisfy $|V| < \infty$. Graphs G and G' on the same vertex set are *isomorphic* if there is a permutation of V that maps $E(G)$ to $E(G')$. A *graph property* is a collection of graphs that is closed under isomorphism. Here are some examples of graph properties:

- $\mathcal{G}_{37}$, the set of graphs with exactly 37 edges.
- $\mathcal{G}_{\delta \geq 3}$, the set of graphs in which every vertex is in at least 3 edges.

- $\mathcal{G}_{\text{tri}}$, the set of graphs containing a triangle (three pairwise joined vertices).
- $\mathcal{G}_{\text{conn}}$, the set of graphs which are connected (between any two vertices there is a path of edges).
- $\mathcal{G}_{\text{Scorp}}$, the set of graphs which are *Scorpions* (with three special vertices, a *sting* joined only to a *tail*, which is also joined to a *body* (but nothing else), with the body joined also to all the other vertices (the *feet*), with no restriction on edges solely involving feet).

A natural notion of the complexity of a graph property $\mathcal{G}$ is the *query complexity*. We imagine that an oracle has access to a graph G, and we can explore the structure of G by asking questions of the form "is $e \in E$?" for any potential edge e. The query complexity of $\mathcal{G}$ is then the smallest k such that, no matter what graph G the oracle has access to, we can determine whether or not $G \in \mathcal{G}$ by asking at most k questions.

A property is said to be *evasive* if its query complexity is $\binom{|V|}{2}$ (the largest it could possibly be). In other words, $\mathcal{G}$ is evasive if for every question-asking strategy we could employ, there is at least one graph G such that under that strategy we cannot determine whether $G \in \mathcal{G}$ without querying every potential edge.

A strategy for establishing evasiveness of a property $\mathcal{G}$ is to think of the oracle being controlled by an adversarial demon, who does not refer to a specific graph G when answering questions, but instead builds G as the questions are being asked in a way that is designed to delay for as long as possible the determination of whether $G \in \mathcal{G}$. We illustrate this idea by using it to demonstrate that the property $\mathcal{G}_{37}$ (having exactly 37 edges) is evasive. The demon has the oracle answer "yes" (that is, "$e \in E$") in response to each of the first 37 queries, and then "no" in response to all the rest. Clearly this adversarial strategy forces us to ask $\binom{|V|}{2}$ questions.

A much more interesting example is $\mathcal{G}_{\text{conn}}$, the property of being connected. Here, the demon has a simple strategy: in response to the question "is $e = \{u, v\} \in E$?" he answers "no" unless that answer would imply that there would be no u-v path in G, even if all remaining unqueried potential edges were included in G; in this case he answers "yes". Suppose (for a contradiction) that there is a strategy that determines whether G is connected in fewer than $\binom{|V|}{2}$ queries. Since throughout the questioning process the demon preserves the possibility that G connected (by adding in all remaining unqueried potential edges), it must be that the strategy concludes that G is connected. Now consider an unqueried edge $e = \{u, v\}$. Since the strategy has concluded that G is connected, there must be a path of queried edges connecting u to v, all of which are known to be in E. Fix such a path, which together with e forms a cycle C, and consider the last edge on it that was queried, $e' = \{u', v'\}$, say. That the demon had answered "yes" to the question "is $e' \in E$?" indicates that at that moment, there was no path of "yes" and unqueried edges connecting u' to v'. But this is a contradiction, since there is such a path, along C.

By a suitably chosen adversarial strategy (often quite intricate), many graph properties, such as $\mathcal{G}_{\text{tri}}$ and $\mathcal{G}_{\delta \geq 3}$ can be shown to be evasive. Indeed, it is hard to find an example of a graph property that is *not* evasive, except the trivial ones, the collection of all graphs and the empty collection (both of which have query complexity 0). But examples do exist; the property $\mathcal{G}_{\text{Scorp}}$ of being a Scorpion, for example, turns out to have query complexity $O(|V|)$.

A feature of many graph properties — such as $\mathcal{G}_{\text{tri}}$, $\mathcal{G}_{\text{conn}}$ and $\mathcal{G}_{\delta \geq 3}$ — is monotonicity. A *monotone* graph property is one that is closed under adding edges. This is a feature that is absent from both $\mathcal{G}_{37}$ and $\mathcal{G}_{\text{Scorp}}$. The following conjecture is often referred to as the Aanderaa-Karp-Rosenberg (AKR) conjecture. As we will discuss presently, a slightly different (and false) form of this conjecture was proposed in [6].

Conjecture 1.1 *Every non-trivial monotone graph property is evasive.*

Originally Aanderaa and Rosenberg conjectured that there is some $\varepsilon > 0$ such that every non-trivial graph property has query complexity at least $\varepsilon |V|^2$. Aanderaa showed this to be false (using $\mathcal{G}_{\text{Scorp}}$), and the ε conjecture was modified to include the word "monotone". It was proved by Rivest and Vuillemin [5], who showed that one may take $\varepsilon = 1/16$; this was later improved to $1/4 - \delta$ for arbitrary $\delta > 0$ by Kahn, Saks and Sturtevant (KSS) [2]. Karp then strengthened the conjecture to its present form.

The AKR conjecture is known to be true if $|V|$ is a prime power (as well as $|V| = 6$). The proof of this fact, by KSS, used topological methods, and the purpose of this note is to explain these methods.

We will focus on a variant of the problem which is simpler in its details, but uses all the relevant concepts. A graph G is said to be *bipartite* with *bipartition* $V = A \cup B$ (or *on* (A, B)) if A and B are disjoint and all edges contain exactly one element from each of A, B. A *bipartite graph property* is a family of bipartite graphs on (A, B) that is closed under permutations of A

and permutations of B; the property is *monotone* if it is closed under adding edges containing exactly one element from each of A, B, and *non-trivial* if it is not the empty family or the family of all bipartite graphs on (A, B). A bipartite graph property is *evasive* if its query complexity, defined in the obvious way, is $|A||B|$. Using the methods of KSS, the AKR conjecture for monotone bipartite graph properties was resolved by Yao [7]. In this note, we give Yao's proof of the following theorem.

Theorem 1.2 *Every non-trivial monotone bipartite graph property is evasive.*

It will be helpful in the sequel to generalize (slightly) the notion of a graph property and bipartite graph property. A *set system* $\mathcal{F}$ on *ground set* X is a collection of subsets of X. When $|X| = \binom{m}{2}$, a set system corresponds exactly to a family of graphs on m vertices (not necessarily closed under isomorphism), via a bijection from X to the set of unordered pairs of vertices. When $|X| = ab$ for some numbers a, b, a set system corresponds to a family of bipartite graphs on (A, B), where $|A| = a$ and $|B| = b$ (again, not necessarily closed under isomorphism), via a bijection from X to $A \times B$.

Extending previous definitions, we say that $\mathcal{F}$ is *monotone* if it is closed under taking supersets and *non-trivial* if there is some subset of $\mathcal{F}$ that is not in $\mathcal{F}$, and some subset that is in $\mathcal{F}$. The query complexity of a set system is defined in the obvious way: an oracle has access to a subset S of X, and we can explore S by asking questions of the form "is $x \in S$?" for any $x \in X$. The query complexity of $\mathcal{F}$ is then the smallest k such that, no matter what set S the oracle has access to, we can determine whether or not $S \in \mathcal{F}$ by asking at most k questions. A set system is then said to be *evasive* if its query complexity is $|X|$.

2 Connection to topology

A *simplicial complex* $\mathcal{K}$ on *ground set* X is a family of subsets of X that is closed under taking subsets. There is a natural bijection between monotone set systems on X and simplicial complexes on X, obtained by sending the monotone set system $\mathcal{F}$ to the simplicial complex $\{X \setminus S : S \in \mathcal{F}\}$. Clearly, determining membership in a monotone set system through queries to an oracle is identical to determining membership in the corresponding simplicial complex, so it makes sense to talk about *evasive* simplicial complexes.

Any simplicial complex on ground set of size n can be realized as a subset $\widehat{\mathcal{K}}$ of $\mathbb{R}^n$: associate a (different) basis vector to each element of X, then to each $K \in \mathcal{K}$ associate the convex hull of the basis vectors associated with elements of K, and finally take the union of all of these sets. We refer to $\widehat{K}$ as a *geometric realization* of $\mathcal{K}$, and say that $\mathcal{K}$ is *contractable* if its realization $\widehat{\mathcal{K}}$ is, that is, if there is a continuous map $\Phi : \widehat{\mathcal{K}} \times [0,1] \to \widehat{\mathcal{K}}$ and a $p \in \widehat{\mathcal{K}}$ such that $\Phi(x,0) = x$ and $\Phi(x,1) = p$ for all $x \in \widehat{\mathcal{K}}$.

The following lemma of Kahn, Saks and Sturtevant [2] shows how topological considerations come into the study of evasiveness.

Lemma 2.1 *If $\mathcal{K}$ is a non-evasive simplicial complex, then $\mathcal{K}$ is contractable.*

The evasiveness of a non-trivial monotone graph or bipartite graph property $\mathcal{G}$ can thus be established via the non-contractability of the associated simplicial complex.

The proof of Lemma 2.1 involves a nice sufficient condition for contractability. We will state it in terms of the more combinatorial notion of collapsibility. A *free face* of a complex $\mathcal{K}$ is a $K \in \mathcal{K}$ that is not maximal (with respect to containment) but is contained in a unique maximal element. An *elementary collapse* of $\mathcal{K}$ is the process of removing from $\mathcal{K}$ some free face K together with all supersets of K, the result of course still being a complex. A complex $\mathcal{K}$ is said to *collapse* to $\mathcal{K}'$ if $\mathcal{K}'$ can be obtained from $\mathcal{K}$ by a sequence of elementary collapses. Finally, say that $\mathcal{K}$ is *collapsible* if it collapses to the empty simplicial complex. A collapsible complex is also contractable (see e.g. [1, p. 49]).

Given a complex $\mathcal{K}$ and an $x \in X$, we define two auxiliary complexes, both on ground set $X \setminus \{x\}$. First, the *link* of $\mathcal{K}$ is

$$\text{link}(x, \mathcal{K}) = \{K \subseteq X \setminus \{x\} : K \cup \{x\} \in \mathcal{K}\}.$$

Next, the *contrastar* of x in $\mathcal{K}$ is

$$\text{cost}(x, \mathcal{K}) = \{K \subseteq X \setminus \{x\} : K \in \mathcal{K}\}.$$

(so $\text{link}(x, \mathcal{K})$ consists of those $K \in \mathcal{K}$ with $x \in K$, but with x removed, and $\text{cost}(x, \mathcal{K})$ consists of those $K \in \mathcal{K}$ with $x \notin K$).

Lemma 2.2 *If there is an $x \in X$ with* $\text{link}(x, \mathcal{K})$ *and* $\text{cost}(x, \mathcal{K})$ *both collapsible, then $\mathcal{K}$ is collapsible.*

Proof: If $K_1, \ldots, K_k$ is a sequence of free faces that can be used to collapse $\mathrm{link}(x, \mathcal{K})$, then $K_1 \cup \{x\}, \ldots, K_k \cup \{x\}$ is a sequence of free faces that can be used to collapse $\mathcal{K}$ to $\mathrm{cost}(x, \mathcal{K})$; the collapsibility of $\mathcal{K}$ then follows from that of $\mathrm{cost}(x, \mathcal{K})$. $\square$

Proof of Lemma 2.1: We proceed by induction on $|X|$ to show that $\mathcal{K}$ is in fact collapsible. If $\mathcal{K}$ is trivial (either empty or a simplex) then it is collapsible (again, see [1]). Note that this includes the base case $|X| = 1$.

If $|X| > 1$ and $\mathcal{K}$ is non-trivial, then there is some $x \in X$ for which "is $x \in K$?" is a *good* first question, that is, the first question in a strategy which always decides membership in $\mathcal{K}$ in fewer than $|X|$ questions. Having asked this question we have two possibilities: if the answer is "no", then $K \in \mathcal{K}$ if and only if $K \in \mathrm{cost}(x, \mathcal{K})$; if the answer is "yes" then $K \in \mathcal{K}$ if and only if $K \setminus \{x\} \in \mathrm{link}(x, \mathcal{K})$. In either case we have (by our choice of x) a strategy which decides membership in the new simplicial complex in fewer than $|X|-1$ questions. In other words, if $\mathcal{K}$ is non-evasive and non-trivial, then there is some $x \in X$ for which both $\mathrm{cost}(x, \mathcal{K})$ and $\mathrm{link}(x, \mathcal{K})$ are non-evasive. By induction and Lemma 2.2, $\mathcal{K}$ is collapsible. $\square$

Here is how our proof of Theorem 1.2 will go. Let $\mathcal{G}$ be a non-trivial monotone bipartite graph property. Because $\mathcal{G}$ is a property, it is invariant under permutations of A and B, and so the associated simplicial complex is invariant under the induced maps. We will examine the set of possible fixed points of these maps, and show that since $\mathcal{G}$ is non-trivial, there cannot be any fixed points. The lack of a fixed point will imply (by some general fixed-point theorems) that the simplicial complex we are working with cannot be contractable, so that by Lemma 2.1, $\mathcal{G}$ must be evasive. The strategy for proving Conjecture 1.1 for m a prime power is similar; the assumption that m is a prime power comes in when we analyze the possible fixed points.

3 Fixed points of simplicial maps

Let $\mathcal{K}$ be a simplicial complex on ground set X. A bijection $\varphi : X \to X$ is an *automorphism* of $\mathcal{K}$ if for each $K \in \mathcal{K}$, we have $\varphi(K)$ $(= \{\varphi(x) : x \in K\}) \in \mathcal{K}$. Such a φ induces continuous linear map $\widehat{\varphi} : \widehat{\mathcal{K}} \to \widehat{\mathcal{K}}$ by $\widehat{\varphi}$ permuting the vertices of $\mathcal{K}$ (the basis vectors associated to each element of X) in correspondence with φ, with the map extended linearly to convex combinations of vertices. A *fixed point* of φ is a $p \in \widehat{\mathcal{K}}$ such that $\widehat{\varphi}(p) = p$.

We write $\text{fix}(\varphi)$ for the set of fixed points of $\mathcal{K}$.

There is a nice combinatorial characterization of $\text{fix}(\varphi)$ in terms of the orbits of φ. If an orbit of φ, X' say, happens to be in $\mathcal{K}$, then (since φ permutes the elements of X'), we have that w', the center of mass of (the geometric realization of) X' is fixed by φ. Let $X_1, \ldots, X_t$ be all such orbits. Then along with the fixed points $w_1, \ldots, w_t$, any convex combination of these points is also fixed by φ, provided, of course, that the point in question is in fact a point of $\widehat{\mathcal{K}}$, which is the same as saying that the union of the corresponding orbits is in $\mathcal{K}$. So $\text{fix}(\varphi)$ contains the set $\widehat{\mathcal{X}}(\varphi)$, the geometric representation of the simplicial complex $\mathcal{X}(\varphi)$ on vertex set $\{X_1, \ldots, X_t\}$, with a set being in $\mathcal{X}$ if the union of the vertices is in $\mathcal{K}$. It turns out that in fact this is all of $\text{fix}(\varphi)$.

Lemma 3.1 $\text{fix}(\varphi) = \widehat{\mathcal{X}}(\varphi)$.

There are many classical fixed-point theorems. The following, a special case of a theorem of Lefshetz [3], generalizes Brouwer's fixed-point theorem.

Theorem 3.2 *Every automorphism of a non-empty contractable simplicial complex has a fixed point.*

We also need a form of the Hopf index formula, for which we need the notion of the *Euler characteristic* of a simplex $\mathcal{K}$. For our purpose this is defined to be

$$\chi(\mathcal{K}) = \sum_{K \in \mathcal{K},\ K \neq \emptyset} (-1)^{|K|}.$$

This is a topological invariant of a simplicial complex, and for all contractable simplicial complexes it is -1.

Theorem 3.3 *If φ is an automorphism of a non-empty contractable $\mathcal{K}$, then the Euler characteristic of* $\text{fix}(\varphi)$ *is* -1.

4 Proof of Theorem 1.2

As a warm-up, we prove the following.

Proposition 4.1 *Let $\mathcal{F}$ be any non-trivial monotone set system closed under a permutation of the ground set that has a single orbit. Then $\mathcal{F}$ is evasive.*

Proof: Let $\mathcal{F}$ be as in the lemma, but not necessarily non-trivial. By a suitable labeling of the ground set, we may assume that $\mathcal{F}$ is invariant under the map $\varphi(x) = x + 1 \pmod{n}$ (here we are identifying the ground set with $\{1, \ldots, n\}$). Suppose that $\mathcal{F}$ is non-evasive. If $\mathcal{F}$ consists of the set of all subsets of X, it is trivial. Otherwise, by Lemma 2.1, the associated (non-empty) simplicial complex $\mathcal{K}$ is contractable. Since $\mathcal{F}$ is invariant under φ, φ is an automorphism of $\mathcal{K}$ and so (by Theorem 3.2) has a fixed point. By Lemma 3.1, the only possible way that φ can have a fixed point is if $\{1, \ldots, n\}$ (the only orbit of φ) is in $\mathcal{K}$; but this implies that $\mathcal{F}$ is empty and so trivial. □

The key idea here is that if we are able to characterize the orbits of a permutation under which $\mathcal{F}$ is invariant, then we also characterize the fixed points of the corresponding simplicial map. Proposition 4.1 is a fairly easy instance of this idea; characterizing the orbits of permutations that fix monotone bipartite families is trickier.

Proof of Theorem 1.2: Let $\mathcal{G}$ be an non-evasive monotone bipartite graph property, with the underlying bipartition (A, B). If $\mathcal{G}$ consists of all bipartite graphs on (A, B) then it is trivial. Otherwise, to $\mathcal{G}$ we correspond, as described, a non-empty, non-evasive simplicial complex $\mathcal{K}$, the elements of whose ground set we refer to as "edges". Let φ be a permutation of the edges that corresponds to a cyclic shift of the vertices of B (leaving A unchanged); this is an automorphism of $\mathcal{K}$. The orbits of φ are exactly the sets of edges that have a vertex of A in common (so there are $|A|$ orbits each of size $|B|$), and we index these orbits in the natural way be elements of A.

Since $\mathcal{K}$ is non-evasive it is contractable and so has at least one fixed point. It follows that there must be at least one orbit of φ, $\{a'\} \times B$ say, that is in $\mathcal{K}$. (Edges are technically unordered pairs, but in the bipartite setting nothing changes if we view them as ordered pairs with the first coordinate in A.) Because $\mathcal{G}$ is a property (closed under isomorphism) we have that $\{a\} \times B \in \mathcal{K}$ for all $a \in A$. The sets in the simplicial complex $\mathcal{X}(\varphi)$ that determines $\text{fix}(\varphi)$ are therefore of the form $A' \times B$ for some $A' \subseteq A$. Again by closure of $\mathcal{G}$ under isomorphism, for a given value of $|A'|$ either all of none of these sets are in $\mathcal{X}$. By monotonicity, then, $\mathcal{X}(\varphi)$ has the form

$$\mathcal{X}(\varphi) = \{A' \times B : |A'| \leq r\}$$

for some $r \leq |A|$. We will argue that we must have $r = |A|$, which says that $A \times B \in \mathcal{X}(\varphi)$ and so $A \times B \in \mathcal{K}$, making $\mathcal{K}$ and thus $\mathcal{G}$ trivial.

The Euler characteristic of $\mathcal{X}(\varphi)$ is

$$\sum_{i=1}^{r}(-1)^i\binom{|A|}{i} = \sum_{i=1}^{r}(-1)^i\left(\binom{|A|-1}{i-1}+\binom{|A|-1}{i}\right)$$
$$= -1+(-1)^r\binom{|A|-1}{r}$$

with the first equality being Pascal's identity and the second the result of telescoping. Since $\mathcal{K}$ is contractable, by Theorem 3.3 the Euler characteristic of $\mathcal{X}(\varphi)$ is -1, so $\binom{|A|-1}{r} = 0$, which can only happen if $|A| = r$. □

References

[1] L. Glaser, *Geometrical Combinatorial Topology*, vol. 1, Van Nostrand, New York, 1970.

[2] J. Kahn, M. Saks and D. Sturtevant, A topological approach to evasiveness, *Combinatorica* **4** (1984), 297–306.

[3] S. Lefschetz, Intersections and transformations of complexes and manifolds, *Trans. Amer. Math. Soc.* **28** (1926), 1-49.

[4] L. Lovász and N. Young, Lecture notes on evasiveness of graph properties, arXiv:cs/0205031v1.

[5] R. Rivest and S. Vuillemin, On recognizing graph properties from adjacency matrices, *Theor. Comp. Sci.* **3** (1978), 371–384.

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

[7] A. Yao, Monotone bipartite graph properties are evasive, *SIAM J. Comput.* **17** (1988), 517-520.