Cops and Robbers

Crystal Xie

December 15, 2024

Abstract

We explore the pursuit-evasion game known as Cops and Robbers, particularly the criteria on the graph that allows a single cop to win, as well as a strategy the cop can use to guarantee capture in a finite number of moves. We will assume no prior knowledge other than basic understanding of graph theory terminology.

1 Introduction

In the game of Cops and Robbers, there is a set of cops C and a robber R, as well as a graph. The graph must be (in our version) a reflexive graph, i.e. an undirected graph where the vertices each have at least one loop. Every graph we look at in this paper will be assumed reflexive unless otherwise stated.

To start the game, the cops choose the vertices they will start on, and then the robber chooses a vertex he will start on. Starting with the cops, the two sides alternate turns, and in each turn, a side of players from that side (namely either the robber, or a subset of the cops) can each move along an edge to a neighboring vertex. The cops win if, at some point, one of them can occupy the same vertex as the robber, resulting in a *capture*. The robber wins if he can avoid capture forever.

Definition. *For a given graph G, we will let the cop-number $c(G)$ denote the minimum number of cops in C that will allow the cops to win.*

There is a lot of work regarding how to calculate the cop-number for various graphs. We will only focus on when this number is 1.

Definition. *The graph G is a cop-win graph if the cop-number is 1, and G is a k-cop-win graph if the cop-number is k for some natural number k.*

Definition. *The graph G is a robber-win graph if the cop number is greater than 1.*

We will also use the following notations in our discussion of various graphs. Readers interested in learning more about graph theory can refer to [2].

- A path of order n, i.e. a sequence of $n+1$ vertices such that each vertex is joined to the next vertex in the sequence (resulting in n edges), will be written as P_n. An infinite one-way path is a ray.
- A cycle of order n, i.e. a path of order n along with an edge joining the last and first vertices in the sequence, will be written as C_n.
- A wheel of order n, i.e. a cycle C_n along with one universal vertex, will be written W_n.
- A complete graph of order n, i.e. a graph where every pair of vertices has an edge joining them, will be written K_n.

- The complement of a graph G, i.e. the graph with the same set of vertices but with pairs of vertices joined if and only if they are not joined in G, will be denoted $\overline{G}$. In particular, the graph with n vertices and no edges is $\overline{K_n}$.

We know that $c(G)$ is a well-defined, finite positive integer, because the cops will win if there are enough cops to assign one to each vertex. In other words, $c(G) \leq |V(G)|$. Moreover, we only need to worry about connected graphs, because:

Lemma 1.1. *Let the disjoint sum of G_1 and G_2 be written as $G = G_1 + G_2$. Then*

$$c(G_1 + G_2) = c(G_1) + c(G_2).$$

An immediate result of this is that $c(\overline{K_n}) = |V(\overline{K_n})| = n$.

The value of $c(G)$ is easy to compute for many common graphs, such as paths, cycles, complete graphs, and trees.

- For an integer $n > 0$, we have
$$c(P_n) = c(W_n) = c(K_n) = 1.$$
- For $n \geq 4$, we have
$$c(C_n) = 2.$$
- A finite tree is cop-win. The cop number of an infinite tree is 1 when the tree is rayless, and infinite otherwise.

Infinite graphs have many pathological properties, so we will ignore them and only focus on finite graphs for this paper.

2 Finite Cop-win Graphs

The graphs that can be won with one cop alone have a beautiful characterization that we will explore in this section.

Definition. *A corner is a vertex x such that there is some vertex y with $N[x] \subseteq N[y]$. The vertex y is said to dominate the vertex x.*

Corners are useless for a robber, since if he is in a corner and the cop is on a vertex that dominates that corner, then wherever the robber moves in his turn, the cop will be able to move to his vertex in her subsequent turn.

For instance, consider the game shown in Figure 1. Note that vertex c is a corner, since vertices d_1 and d_2 both dominate it. If the robber is in vertex c, and the cop is on d_1, then if the robber goes to d_2, the cop can follow to d_2. If the robber goes to d_1, then because we are assuming all graphs are reflexive as part of the ruleset, the cop can catch the robber at d_1.

So, we might want to ignore the useless vertex and consider the graph without that corner. We can do this with the useful notion of retracts:

Definition. *Let H be an induced subgraph of G, formed by deleting one vertex. If there is a homomorphism f from G onto H such that $f(x) = x$ for all $x \in V(H)$ (the identity on H), then we say that H is a one-point retract of G, and the map f is called a one-point retraction.*

For example, the subgraph formed by deleting a leaf from a tree is a retract of the original tree. Notice that the distance between any two vertices doesn't increase when a retract is taken.

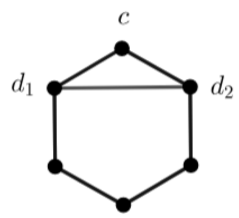

Figure 1: A game with corner c and neighboring vertices d_1 and d_2.

There is a more general notion of retracts that aren't restricted to deleting one vertex only. They can be considered as a composition of a set of one-point retractions. We won't really need general retractions, so we'll refer to one-point retractions as simply retractions from now on.

Theorem 2.1. *If H is a retract of G, then $c(H) \leq c(G)$.*

It follows from this theorem that if G is a cop-win graph (i.e. $c(G) = 1$), then so are its retracts. In the other direction, if there is a retract of G that is a cycle with at least 4 vertices, then $c(G) \geq 2$ and thus G must be robber-win.

We can take a retract that removes (from our consciousness) a corner u dominated by v, using the retraction

$$f(x) = \begin{cases} v & \text{if } x = u, \\ x & \text{else.} \end{cases}$$

We will say $u \to v$ here. By Theorem 2.1, if G is cop-win, deleting the corner still results in a cop-win graph. Furthermore, if G is cop-win, then we can always do a corner removal.

Lemma 2.2. *If G is a cop-win graph, then G contains at least one corner.*

Therefore, we say that a graph is *dismantlable* if there is some sequence of corner removals that results in the graph K_1 (the graph that is a single vertex).

Surprisingly, the dismantlable graphs completely characterize cop-win graphs. This theorem was proven by Nowakowski and Winkler, and independently by Quilliot, and it is the main result of this section:

Theorem 2.3. *A graph is cop-win if and only if it is dismantlable.*

Proof. For the "only if" direction, we use induction on the number of vertices of G. The base case of K_1 is trivial. In general, each cop-win graph has a corner u dominated by some vertex v, by Lemma 2.2. Then $G-u$ is a retract of G, which is also cop-win. Also, $G-u$ has one less vertex, so it is dismantlable, meaning G is dismantlable.

For the "if" direction, we also use induction on the number of vertices of G. The base case is again the trivial case K_1. In general, if G is dismantlable and has $n+1$ vertices for some fixed $n \geq 1$, then G has a corner u, with $u \to v$ for some vertex v such that $G-u$ is dismantlable. Note that $G-u$ is cop-win by the induction hypothesis.

The strategy for the cop is as follows: she can play in $G-u$ using her winning strategy here, but whenever the robber moves to u, the cop moves as if he had moved to v, which is possible since $u \to v$. Say that f is a retraction that maps u to v and fixes everything else. Then the cop will eventually capture the image of the robber, $f(R)$, with her winning strategy on $G-u$. At this point, either $R = f(R)$, or R is on u and C is on v, and the cop wins either this round or the next round, respectively. $\square$

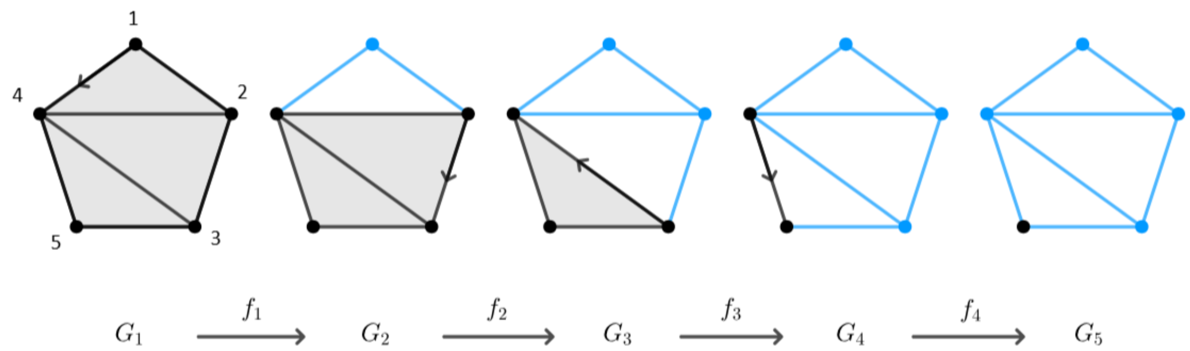

Figure 2: A sequence of retracts of a cop-win graph.

3 The Cop-win Strategy for Dismantlable Graphs

Let there be a graph G, and let $n = |V(G)|$. This graph is dismantlable if we can label the vertices using the integers $1, 2, \ldots, n$ such that for each $i < n$, the vertex i is a corner in the subgraph of G induced by $\{i, i+1, \ldots, n\}$. In other words, this is a way to make a sequence of retractions that remove vertices in the order given by their labels, such that every time we are removing a corner from the remaining subgraph. We call this ordering of $V(G)$ a *cop-win ordering*. There can be more than one cop-win ordering for the same graph.

There is a winning strategy for the cops that utilizes this somewhat linear structure on the vertices of a dismantlable graph, called the **Cop-win Strategy**. To set this up, we need to define some notation. For each i, let $S_i = \{n, n-1, \ldots, i\}$ be the set of vertices remaining after $i-1$ retractions, and let G_i represent the retract of G removing those first $i-1$ corners. So G_i is the subgraph of G induced by S_i, i.e. the graph with vertices in S_i with pairs of vertices joined if and only if they are joined in G.

For each $1 \leq i \leq n-1$, let $f_i : G_i \to G_{i+1}$ be the retraction map from G_i to G_{i+1}. Note that, as before, f_i maps the vertex i to a vertex that dominates i in G_i. Then for each $1 \leq i \leq n$, let F_i be the mapping formed by retracting corners $1, 2, \ldots, i-1$, i.e.

$$F_i = f_{i-1} \circ \cdots \circ f_2 \circ f_1,$$

and let F_1 be the identity mapping on G. The F_i are all homomorphisms, since the f_i are (and because all our graphs are assumed to be reflexive). We demonstrate this notation, along with a cop-win ordering, in Figure 2, with the corresponding images under F_i shown in the table below.

v	$F_1(v)$	$F_2(v)$	$F_3(v)$	$F_4(v)$	$F_5(v)$
1	1	4	4	4	5
2	2	2	3	4	5
3	3	3	3	4	5
4	4	4	4	4	5
5	5	5	5	5	5

The **Cop-win Strategy** dictates that the cop begins on the vertex $G_n = n$, which is the image of the robber's position (and all of $V(G)$) under F_n. During the game, if the robber is at vertex u and the cop is on $F_i(u)$ (the image of the robber's position in G_i), then if the robber moves to the vertex v, the cop moves to the image $F_{i-1}(v)$ of the robber in the larger graph G_{i-1}. In other words, the cop keeps moving to the robber's image in progressively less "advanced" retractions, and this strategy says that eventually the robber's position will coincide with his image, with the cop on it.

We demonstrate this using the example given in Figure 2 and the corresponding table. First, the cop will start on the vertex 5. If the robber chooses to start on vertex 3, then the cop is on $F_5(3)$, so in the cop's first move, she will move to vertex $F_4(3) = 4$. If the robber then moves to vertex 2, then the cop moves to vertex $F_3(2) = 3$. If the robber then moves to vertex 1, then the cop moves to vertex $F_2(1) = 4$. Finally, if the robber moves to vertex 2, then the cop moves to vertex $F_1(2) = 2$, catching the robber. The Cop-Win Strategy guarantees that in each turn, the cop can always move this way, and will eventually catch the robber.

Theorem 3.1. *Using the Cop-win Strategy, the cop will capture the robber in at most n moves.*

In our proof we'll use the notion of neighbor sets to simplify the reasoning.

Definition. *For a vertex x in a given graph, the set of vertices joined to x but not equal to x is the (open) neighbor set $N(u)$. The closed neighbor set, i.e. the set containing both x and its neighbors, is $N[u] = N(u) \cup \{x\}$.*

Proof. We will use induction to show that the cop can always keep moving to the image of the robber. More precisely, say the robber is on u and the cop is on $F_i(u)$, where $i \leq n$. Suppose the robber moves to v. Then we need to show that the cop can move to $F_{i-1}(v)$ from $F_i(u)$, or in neighborhood notation, that

$$F_{i-1}(v) \in N[F_i(u)]. \tag{*}$$

Note that because v is joined with u, and retractions never increase distance between pairs of vertices, we have

$$F_{i-1}(v) \in N[F_{i-1}(u)]. \tag{†}$$

In other words, because v is in the neighborhood of u, all the images of v are in the closed neighborhood of u as well.

First, if $F_{i-1}(u) = F_i(u)$, i.e. if $F_{i-1}(u)$ is not the corner removed in the $(i-1)$-th retraction, then from (†) we have

$$F_{i-1}(v) \in N[F_{i-1}(u)] = N[F_i(u)],$$

so $(*)$ is true. In the other case, where $F_{i-1}(u)$ is the corner removed in the $(i-1)$-th retraction, we know from the definition of corners that $N[F_{i-1}(u)] \subseteq N[F_i(u)]$. Combining this with (†), we get

$$F_{i-1}(v) \in N[F_{i-1}(u)] \subseteq N[F_i(u)],$$

so we have shown $(*)$ again. Therefore, in each step, the cop can catch the robber's image, and in at most n moves of the cop, she will move to $F_1(v) = v$ and catch the robber. □

References

[1] Anthony Bonato and Richard J. Nowakowski. *The Game of Cops and Robbers on Graphs.* American Mathematical Society, 2011.

[2] Radu Bumbacea. *Graphs: An Introduction.* XYZ Press, 2020.