Chapter 1

Introduction

1.1. The Game

We all grew up playing games. Some of us are lucky enough to play them while working. Such is the case with Cops and Robbers: it is at once a game you can play for fun on a piece of paper with some spare coins and a deep mathematical research topic containing hard conjectures and problems. The purpose of this chapter is a kind of *mezze*: readers will gain the requisite notation and background to tackle the harder topics in later chapters, and also gain some insight into the heart of the game.

To set the stage, do you remember the video game Pac-Man? If you are not a member of the video game generation, then let us recall how it is played. You, Pac-Man, are stuck in a maze. You can move up and down, and across, but not through walls. Unfortunately, there are some attackers in the form of ghosts who are trying to capture you. They do this by touching you, or by *occupying your position in the maze.* Your goal is to eat dots set throughout the maze while avoiding capture. We do not care as much about the dot eating. In some sense, the real goal is to move about the maze unfettered by the ghosts. This is fairly easy with one ghost, but the more ghosts, the greater chance you have of being captured sooner. You can see all

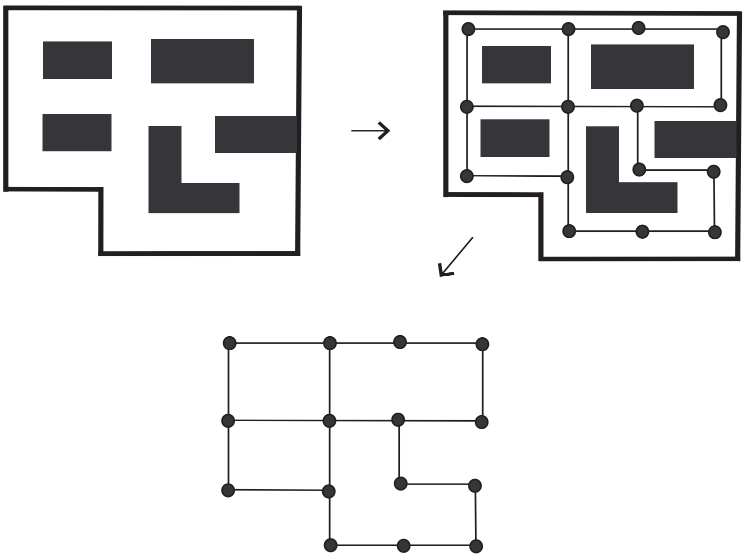

Figure 1.1. A maze and its corresponding graph.

the players loose in the maze and remember all the moves of ghosts (and they can see you).

We may think of the maze as a set of discrete cells, each joined to one above, below, or beside it, assuming there is no wall blocking your way. To help visualize this, see Figure 1.1. For more on this approach in artificial intelligence and so-called *moving target search*, see [**157**] and also [**120**]. (In moving target search, octile connected maps which allow diagonal moves are often studied. In this case, a cell becomes a clique of order four.) Analyzing the movements of players in Pac-Man then becomes a problem about certain kinds of graphs. We focus on a particular view that deviates from the original game somewhat: how many ghosts are needed to ensure they can always capture you, by some strategy? Some mazes require more ghosts, some fewer. For example, think of a very simple maze consisting of a rectangle. One ghost would eternally chase you to no avail, but two can corner you. The game of Cops and Robbers is—in some sense—a discretized version of Pac-Man, and the cop number corresponds to the minimum number of ghosts needed to capture you. You are the intruder, or robber, and the cops are the ghosts.

To be more precise, *Cops and Robbers* (or, as it is sometimes called, *Cops and Robber*) is a game played on a reflexive graph; that is, the vertices each have at least one loop. Multiple edges are allowed but make no difference to the game play, so we always assume there is exactly one edge between adjacent vertices. There are two players consisting of a set of *cops* and a single *robber*. The game is played over a countable sequence of discrete time-steps or *rounds*, with the cops going first in round 0. The cops and robber occupy vertices; for simplicity, we often identify the player by the vertex they occupy. We refer to the set of cops as C and the robber as R. The rules of the game are straightforward: when a player is ready to move in a round they must move to a neighboring vertex. Because of the loops, players can *pass* or remain on their own vertex. This may or may not be a wise strategy for the robber, depending on the graph. Note that if we play on irreflexive graphs, then we still allow passes. Also observe that any subset of C may move in a given round.

The cops win if after some finite number of rounds, one of them can occupy the same vertex as the robber (in a reflexive graph, this is equivalent to the cop landing on the robber). This is called a *capture*. The robber wins if he (usually the cops are considered female and the robber male) can evade capture indefinitely. A *winning strategy for the cops* is a set of rules that, if followed, result in a win for the cops. A *winning strategy for the robber* is defined analogously. Cops and Robbers is often called a *vertex-pursuit game* on graphs, for reasons that should now be apparent to the reader.

As an elementary but instructive example, consider the game played on a 5-cycle C_5. We label the vertices $1, 2$, 3, 4, and 5, as in Figure 1.2, and place a cop on vertex 1. If the robber chooses 1, then that would be suicide, and choosing vertex 2 or 5 would result in his losing in round 1. The robber chooses 3 and can evade capture in round 1. It is straightforward to see the robber has a winning strategy (just move to $i \pm 1 \pmod 5$ in order to maintain distance two from the cop).

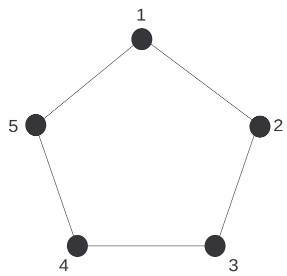

Figure 1.2. A labeled 5-cycle.

Two cops are enough, however, to win. If a second cop occupies 3, then the robber will be caught in round 0 or 1, depending on his initial move. Cycles of size 4 or larger are similar with respect to the game (note that cycles correspond to discretized versions of the simplified rectangular maze we discussed above), because two cops are necessary and sufficient to guarantee a win for the cops.

If we place a cop at each vertex, then the cops are guaranteed to win. Therefore, the minimum number of cops required to win in a graph G is a well-defined positive integer (or infinite cardinal) called the *cop number* (or *copnumber*) of the graph G. We write $c(G)$ for the cop number of a graph G. If $c(G) = k$, then we say G is *k-cop-win*. In the special case $k = 1$, we say G is *cop-win* (or *copwin*). A graph with $c(G) > 1$ is sometimes called *robber-win* (since one cannot capture the robber).

The game of Cops and Robbers was first considered by Quilliot [**167**] in his doctoral thesis, and was independently considered by Nowakowski and Winkler [**165**]. Although [**167**] predates [**165**], the latter reference is sometimes referred to as the starting point of the literature on the topic. The authors of [**165**] were told about the game by G. Gabor. Both [**167**] and [**165**] refer only to one cop. The introduction of the cop number came in 1984 with Aigner and Fromme [**2**]. Many papers have now been written on the cop number of graphs since these three early works; see the surveys [**7**] and [**103**]. For example, at least a dozen theses (at the master's and doctoral level) have been written on the topic; see [**14**], [**51**], [**52**], [**79**], [**97**], [**113**], [**122**], [**156**], [**161**], [**167**], [**168**], [**179**], [**184**], and [**188**].

As an introduction to the topic of Cops and Robbers, we begin this chapter by first covering some notation and definitions from graph theory in Section 1.2. The more advanced reader can skip this, although a casual perusal may eliminate any confusion with notation when reading later sections and chapters. We discuss some examples of cop number in Section 1.3, and include the elementary but helpful Theorem 1.3 which provides a lower bound on the cop number in terms of the minimum degree for graphs without small cycles. In Section 1.4 we prove Frankl's upper bound for the cop number; see Theorem 1.6. Along the way, we will show that one cop can guard an isometric path. We finish with a discussion of retracts in Section 1.5, which play a critical role in the structure of cop-win graphs.

1.2. Interlude on Notation

As we stated in the Preface, we assume (although it is not essential) that the reader has some background in graph theory, such as a first course on the topic. Two good references on the topic are [**68**] and [**195**]. However, as an aid to the reader, we summarize at least some of the notation used as well as some of the requisite background here. As such, the present section is short and may be safely skipped by more advanced readers.

We will use the following notation throughout. The set of natural numbers (which contains 0) is written $\mathbb{N}$, while the rationals and reals are denoted by $\mathbb{Q}$ and $\mathbb{R}$, respectively. The cardinality of $\mathbb{N}$ is $\aleph_0$, while the cardinality of $\mathbb{R}$ is $2^{\aleph_0}$. If $n > 0$ is a natural number, then define

$$[n] = \{1, \ldots n\}.$$

The Cartesian product of two sets A and B is written $A \times B$. The difference of two sets A and B is written $A \backslash B$.

As we will present a number of asymptotic results, we give some corresponding notation. Let f and g be functions whose domain is some fixed subset of $\mathbb{R}$. We write $f \in O(g)$ if

$$\limsup_{x\to\infty} \frac{f(x)}{g(x)}$$

exists and is finite. We will abuse notation and write $f = O(g)$. This is equivalent to saying that there is a constant $c > 0$ (not depending on x) and an integer N such that for $x > N$, $f(x) \leq cg(x)$.

We write $f = \Omega(g)$ if $g = O(f)$, and $f = \Theta(g)$ if $f = O(g)$ and $f = \Omega(g)$. If $\lim_{x\to\infty} \left|\frac{f(x)}{g(x)}\right| = 0$, then $f = o(g)$ (or $g = \omega(f)$). So if $f = o(1)$, then f tends to 0. We write $f \sim g$ if

$$\lim_{x\to\infty} \left|\frac{f(x)}{g(x)}\right| = 1.$$

If x is a real number, then $1 + x \leq e^x$. We will sometimes write e^x as $\exp(x)$, especially if x is a complicated expression. We write $\log x$ for the logarithm in base e (other bases will be made explicit). If $4 \leq m \leq n$ are non-negative integers, then

$$\binom{n}{m} \leq \frac{n^m}{2^m} \leq n^m. \tag{1.1}$$

For a graph G, we often write $G = (V(G), E(G))$, or if G is clear from context, $G = (V, E)$. The set E may be empty. Elements of $V(G)$ are *vertices*, and elements of $E(G)$ are *edges*. Vertices are sometimes referred to as *nodes*. We write uv for an edge $\{u, v\}$, and say that u and v are *joined* or *adjacent* (we use both terms interchangeably); we say that u is *incident* with v, and that u and v are *endpoints* of uv. All the graphs we consider are reflexive unless otherwise stated.

The cardinality $|V(G)|$ is the *order* of G, while $|E(G)|$ is its *size*. Given a vertex u, define its *neighbor set* $N(u)$ to be the set of vertices joined and not equal to u (also called *neighbors* of u). The *closed neighbor set* of u, written $N[u]$, is the set $N(u) \cup \{u\}$. We write $G \restriction S$ (or as either $\langle S\rangle_G$ or $G[S]$) for the subgraph of G *induced by the set of vertices* S; that is, the graph with vertices in the set S, with two vertices joined if and only if they are joined in G. If S is a set of vertices, then $G - S$ is the subgraph induced by $V(G)\backslash S$; if $S = \{x\}$, then we write this as $G - x$. If H is an induced subgraph of G, then we sometimes write $G - H$ for $G - V(H)$.

The *degree* of a vertex is the cardinal $|N(u)|$, and is written $\deg_G(u)$ or simply $\deg(u)$. A graph is *k-regular* if each vertex has degree k. A *path* is a sequence of vertices such that each vertex is joined to the next vertex in the sequence; the length of a path is the number of its edges. A path of order n is written P_n. A graph is *connected* if there is a path between any two vertices. The relation of being connected by a path is an equivalence relation on V, and the equivalence classes are the *connected components* of G. A graph which is not connected is called *disconnected*; a connected component consisting of a single vertex is called an *isolated vertex*. A *cut vertex* is one whose deletion results in a disconnected graph. A vertex joined to all other vertices is called *universal*. A vertex of degree one will be called an *end-vertex*.

A *homomorphism* f from G to H is a function $f : V(G) \to V(H)$ which *preserves edges*; that is, if $xy \in E(G)$, then $f(x)f(y) \in E(H)$. We abuse notation and simply write $f : G \to H$. An *embedding* from G to H is an injective homomorphism $f : G \to H$ with the property that $xy \in E(G)$ if and only if $f(x)f(y) \in E(H)$. We will write $G \leq H$ if there is some embedding of G into H and say that G *embeds in* H. An *isomorphism* is a bijective embedding; if there is an isomorphism between two graphs, then we say they are *isomorphic*. We write $G \cong H$ if G and H are isomorphic. The relation $\cong$ is an equivalence relation on the class of all graphs, whose equivalence classes are *isomorphism types* or *isotypes*. We will always identify a graph with its isomorphism type. An *automorphism* of G is an isomorphism from G to itself. A graph is *vertex-transitive* G if for all pairs of vertices u and v of G, there is an automorphism f of G, so that $f(u) = v$. Note that every vertex-transitive graph is k-regular for some integer $k > 0$.

The *distance* between u and v, written $d_G(u, v)$ (or just $d(u, v)$), is either the length of a shortest path connecting u and v (and 0 if $u = v$) or ∞ otherwise. Note that $d(u, v)$ turns each graph into a metric space. The *diameter* of a connected graph G, written $\mathrm{diam}(G)$, is the supremum of all distances between distinct pairs of vertices. If the graph is disconnected, then $\mathrm{diam}(G)$ is ∞.

The *complement* of G, written $\overline{G}$, has vertices $V(G)$ with two distinct vertices joined if and only if they are not joined in G. A *complete graph of order* n or n*-clique* has all edges present and is written K_n. A set of vertices S is *independent* if $\langle S\rangle_G$ contains no edges. A *co-clique of order* n is $\overline{K_n}$. The graph of order n with no edges is $\overline{K_n}$.

A *wheel of order* n, written W_n, consists of a cycle C_n along with one universal vertex. A *hypercube* of dimension n, written Q_n, has vertices elements of $\{0,1\}^n$ with two vertices joined if they differ in exactly one coordinate.

The *chromatic number* of G, written $\chi(G)$, is the minimum cardinal n with the property that $V(G)$ may be partitioned into n many independent sets; that is, the minimum n so that G has *proper* n*-coloring*. If $\chi(G)=2$, then G is *bipartite*. A *complete bipartite graph* has all possible edges present between the two colors, and is written $K_{m,n}$, where m and n are the orders of the vertex classes. A *star* is a graph $K_{1,n}$, for some positive integer n.

In a graph G, a set S of vertices is a *dominating set* if every vertex not in S has a neighbor in S. The *domination number* of G, written $\gamma(G)$, is the minimum cardinality of a dominating set. Since placing a cop on each element of a dominating set ensures a win for the cops in at most two rounds, we have that $c(G)\leq\gamma(G)$.

Although our primary focus is on undirected graphs, we may sometimes assign orientations to edges. A *directed graph* or *digraph* is defined identically as a graph, except that $E(G)$ consists of ordered pairs of vertices. As with graphs, we assume our directed graphs are reflexive. The edges are then called *directed edges* or *arcs* (u,v), where u is the head and v is the tail. The vertex v is an *out-neighbor* of u, while u is an *in-neighbor* of v. The *in-degree* of u, written $\deg^-(u)$ is the number of vertices v such that (v,u) are directed edges; the *out-degree* $\deg^+(u)$ is defined dually. Subgraphs, induced subgraphs, and isomorphisms are defined analogously to graphs.

A digraph is *oriented* if it is antisymmetric: if (u,v) is a directed edge, then (v,u) is not a directed edge. An *orientation of a graph* is an assignment of directions to the edges resulting in an oriented graph. A *tournament* is an orientation of a clique.

An *order* (or *partially ordered set* or *poset*) is an oriented digraph that is *transitive*: whenever (u, v) and (v, w) are arcs, then so is (u, w). We write $u \leq v$ if (u, v) is an arc in an order. We say that v *covers* u if $u \leq v$, $u \neq v$, and there is no x such that $u \leq x \leq v$. A vertex u is *minimal* if $v \leq u$ implies that $v = u$; *maximal* elements are defined dually. Orders are often represented by *Hasse diagrams*, which are drawings in the plane (although edge crossings are allowed) so that u is below and adjacent to v if v covers u. Note that reflexive and transitive arcs are not shown in Hasse diagrams. See Figure 1.3.

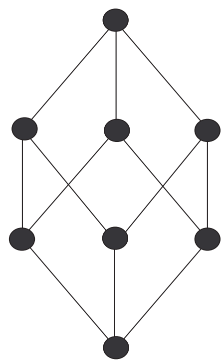

Figure 1.3. The Hasse diagram of an order.

A *directed path* is a path with all directed edges pointing in one direction (so all vertices internal to the path have in- and out-degree equaling one). A *directed cycle* is a cycle with all arcs directed in the same direction. A digraph is *strongly connected* if there is a directed path connecting every pair of vertices. A *weakly connected* digraph has its underlying undirected graph (with no orientations on edges) connected. A digraph is *acyclic* if it contains no directed cycle.

The cop number of directed graphs is defined in the analogous way to the undirected case. The only difference, of course, is that the players can only move following the orientation of a directed edge. A version of Cops and Robbers played on orders will be explored in Exercise 27.

1.3. Lower Bounds

When graph theorists see a new graph parameter, they usually first attempt to compute it for the most common graphs such as cycles, paths, and cliques. The following lemma—whose proof is left as an exercise—does just that.

Lemma 1.1. (1) *For $n > 0$ an integer we have that*

$$c(P_n) = c(W_n) = c(K_n) = 1,$$

and for $n \geq 4$,

$$c(C_n) = 2.$$

(2) *If G is the disjoint union of G_1 and G_2 written $G_1 + G_2$, then*

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

In particular,

$$c(\overline{K_n}) = n.$$

Owing to Lemma 1.1 (2), we usually restrict our attention to connected graphs. For example, one cop is needed for each isolated vertex, since the robber can occupy one and pass indefinitely. Trees, which are connected and contain no cycles, are a favourite graph class. An infinite one-way path (that is, the vertices of the path are just the non-negative integers, with i joined to $i + 1$ for all $i \in \mathbb{N}$) is called a *ray*, and a graph with no ray is called *rayless*.

Lemma 1.2. (1) *A finite tree is cop-win.*

(2) *The cop number of an infinite tree is either 1 or infinite. It is 1 exactly when the tree is rayless.*

Proof. For item (1), we use the fact that each finite tree contains an end-vertex (finite trees always contain at least two end-vertices; see Exercise 4a). Place the cop on an arbitrary vertex. The strategy of the cop is to move towards the robber on the unique path connecting the cop and robber. Note that with this strategy, $d(C, R)$ never increases. A simple induction establishes that this is possible in any connected graph (roughly put, the robber can never "move around" the cop). However, after some number of rounds (bounded above by diam(T)), the robber will move to an end-vertex. After that round, $d(C, R)$ decreases by one since there is a unique path connecting R and C. Repeating this argument after at most diam$(T) - 1$ many rounds results in $d(R, C) = 0$, and the cop wins.

A tree of any order with no ray has an end-vertex (see Exercise 9). Now, in a rayless tree, apply the same winning strategy as the one used by the cop in a finite tree. If the tree has a ray and only a finite number of cops are at play, then the robber can always stay a distance of at least one away from any cop. Hence, no winning strategy exists for the cops, and the robber wins. □

End-vertices play a critical role in the proof of Lemma 1.2 (1). They are the simplest examples of *corners*: vertices x with the property that there is some vertex y such that $N[x] \subseteq N[y]$. Corners play a major role in characterizing finite cop-win graphs. See Section 1.5 and Chapter 2 for more discussion.

See Figure 1.4 for an example illustrating Lemma 1.2 (2). This tree is formed by attaching a path of each finite length to a root vertex. The rayless tree in Figure 1.4 has an unusual and vaguely morbid property: the robber in round 0, by choosing which branch to occupy, decides how long he wants to live! Infinite graphs demonstrate many pathological properties, as demonstrated by this example. They therefore deserve special attention and form the focus of Chapter 7. We therefore make the following assumption for the remainder of this

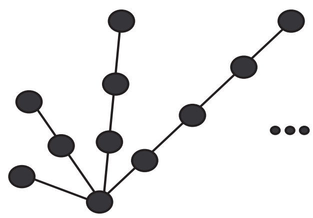

Figure 1.4. A rayless tree.