Cops, Robber, and Alarms.

Nancy E. Clarke*
Acadia University
Wolfville, Nova Scotia

Emma L. Connon
University of Waterloo
Waterloo, Ontario

Abstract

The two games considered are mixtures of Searching and Cops and Robber. The cops have partial information, provided first via selected vertices of a graph, and then via selected edges. This partial information includes the robber's position, but not the direction in which he is moving. The robber has perfect information. In both cases, we give bounds on the amount of such information required by a single cop to guarantee the capture of the robber on a copwin graph.

Key words: game, cop, partial information, alarm, pursuit, graph.

2000 Mathematics Subject Classifications: 05C75, 05C99, 91A43.

1 Introduction

The game of Cops and Robber is a pursuit game played on a reflexive graph, i.e. a graph with a loop at every vertex. There are two opposing sides, a set of $k > 0$ cops and a single robber. The cops begin the game by each choosing a vertex to occupy, and then the robber chooses a vertex. The two sides move alternately, where a move is to slide along an edge or along a loop, i.e. pass. There is perfect information, and the cops win if any of the cops and the robber occupy the same vertex at the same time. Graphs on which one cop suffices to win are called *copwin* graphs and are characterized in [7, 9].

In this paper, variations of the Cops and Robber game are considered in which the cops no longer have perfect information, but rather can only

*Thanks are extended to NSERC and MITACS for partial financial support.

get partial information about the robber's movements through the use of alarms, located first on selected vertices of a graph and then on selected edges. As in the original game, the robber has perfect information.

Suppose the game is played on a graph G. In the first variation considered here, the partial information is provided by alarms located on selected vertices of G. These units alert the cops if the robber moves onto a vertex equipped with an alarm. The alarms do not indicate the direction in which the robber is moving when he leaves such a vertex. The minimum number of alarms required by a single cop to guarantee the capture of the robber on (a copwin graph) G is the *alarm number* of G, denoted $A(G)$.

In the second variation considered, the partial information is provided by alarms located on selected edges of G. These units alert the cops if the robber moves along an edge equipped with an alarm, but provide no information about the direction in which the robber is moving. The minimum number of alarms required by a single cop to guarantee the capture of the robber on G is the *edge alarm number* of G, denoted $A^*(G)$.

In general, one can ask for the minimum amount of information needed if there are k cops.

These variations of the game were first introduced in [3, 4], where bounds are given for the alarm and edge alarm numbers of trees. In [3, 4, 6], variations of the game are introduced in which the partial information includes information about the robber's direction, in addition to his position. Again, the information may be provided via selected vertices of a graph or via selected edges. Bounds are given on the amount of such information required by a single cop to guarantee a win on a tree and, with less tight bounds, on a copwin graph. The strategies given for trees are generalized for use on copwin graphs, using the notion of *copwin spanning tree* (defined at the end of this section). Unfortunately, in the case of alarms, the strategies presented in [3, 4] for trees cannot be generalized in this way because of the lack of a directional signal. Thus, a different approach is needed for other copwin graphs. In this paper, we bound the alarm and edge alarm numbers of arbitrary copwin graphs.

We now introduce our notation and basic definitions. Recall that, for us, all graphs will be finite, connected and reflexive. For a graph G, we let $V(G)$ denote the set of vertices of G and $E(G)$, the set of edges. For $e, f \in E(G)$, we use $e \leftrightarrow f$ to indicate that e and f are incident. For $a, b \in V(G)$, we use $a \sim b$ to indicate that a and b are adjacent ($a \neq b$), and $a \simeq b$ if a is adjacent or equal to b. For $x \in V(G)$, $N(x) = \{y \mid y \sim x\}$ is the *open neighbourhood* of x and $N[x] = N(x) \cup \{x\}$ is the *closed neighbourhood*. If $Y \subseteq V(G)$, then $G - Y$ is the induced subgraph with vertex set $V(G) \setminus Y$. If $Z \subseteq E(G)$, then $G - Z$ is the subgraph with vertex set $V(G)$ and edge set $E(G) \setminus Z$.

A mapping $f : V(G) \to V(H)$ is a *homomorphism* if, for $x, y \in V(G)$, $f(x) \simeq f(y)$ whenever $x \sim y$. A subgraph H of a graph G is a *retract* of G if there is a homomorphism $f : V(G) \to V(H)$ such that $f(x) = x$, for all $x \in V(H)$. Note that, since G is reflexive, a homomorphism can send two adjacent vertices to the same vertex.

Sometimes we need to consider the situation where the cops are playing on a retract H, while the robber is playing on the full graph G. If r is the vertex occupied by the robber and f is a fixed retraction map of G onto H, then we refer to $f(r)$ as the robber's *image*. A vertex u of a graph G is *irreducible* if there exists a vertex v in G such that $N[u] \subseteq N[v]$; also we say that v *dominates* u.

A vertex ordering $(x_1, x_2, \ldots, x_n)$ on G is a *domination elimination ordering* [1, 2] if, for each $i \in \{1, 2, \ldots, n-1\}$, there is a $j_i > i$ such that $N_i(x_i) \subseteq N_i[x_{j_i}]$ in $G_i = G - \{x_1, x_2, \ldots, x_{i-1}\}$. If, in addition, for each i, $x_i \sim x_{j_i}$, then this domination elimination ordering is a *copwin ordering* [7]. A main result of [7, 9] is that: *a finite graph G is copwin if and only if G has a copwin ordering.* Structurally, this means that G is copwin if and only if it can be reduced to a singleton by a series of retractions and, in each retraction, there is exactly one vertex that is not fixed and it is irreducible.

Let $(x_1, x_2, \ldots, x_n)$ be a copwin ordering of G. For $j = 1, 2, \ldots, n-1$, define $f_j : V(G_j) \to V(G_{j+1})$ to be the retraction map from G_j to G_{j+1}. If x_k dominates x_j in G_j, so that $f_j(x_j) = x_k$, $k > j$, we write $x_j \to x_k$. Further, if the robber is on vertex r, define $F_i(r) = f_{i-1} \circ f_{i-2} \circ \cdots \circ f_1(r)$, so that F_i is the robber's image on G_i.

The corresponding *copwin spanning tree*, denoted S_{x_n}, is a spanning tree, rooted at x_n, with the property that for vertices $x, y \in V(G)$, $xy \in E(S_{x_n})$ if and only if $f_j(x) = y$ or $f_j(y) = x$, for some j. We say that $x \succeq y$ if $F_i(y) = x$, for some $i = 2, 3, \ldots, n$, and $x \succ y$ if $x \neq y$.

In [5], a strategy is presented that can be used by a single cop to win on a copwin graph. This is known as a *copwin strategy*. Again, consider a fixed copwin ordering $(x_1, x_2, \ldots, x_n)$ of G. The cop begins on vertex x_n, since $F_n(x) = x_n$, for all $x \in V(G)$. Having captured the image of the robber on G_i, $2 \leq i \leq n$, the cop is able to move immediately onto the image on G_{i-1}. Essentially, this is because if $y \simeq z$, $y, z \in V(G)$, then $F_i(y) \simeq F_{i-1}(z)$. The proofs of the main theorems in this paper make use of the copwin strategy.

2 Alarms on Vertices

The schemes given for the placement of alarms are designed to allow the cop to gather sufficient information to be able to play the copwin strategy.

A vertex v of G is *controlled* if v remains unalarmed, but all neighbours of v are alarmed. We define a_G as the minimum number of vertices having alarms such that all vertices of G are either alarmed or controlled, and no alarmed vertex has more than one unalarmed neighbour.

Theorem 1 *Let G be a finite, copwin graph. Then $A(G) \leq a_G$.*

Proof. Let $(x_1, x_2, \ldots, x_n)$ be a copwin ordering of G which realizes a_G. Alarms are placed according to the scheme described in the preamble to the theorem. The cop begins on vertex x_n, and visits each controlled vertex of G. If no alarms sound, then the robber will be found during this phase of the cop's strategy.

So suppose that an alarm sounds. If the cop is on an adjacent vertex, then he moves immediately onto the alarmed vertex and captures the robber. (Note that if the cop is ever on a vertex adjacent to the robber's position, he moves immediately to that vertex.) Otherwise, the cop returns to x_n, and plays the copwin strategy. To do this, the cop must be able to locate the robber's image on G_{n-1} after the robber's next move (i.e. the robber's first move after the cop reaches x_n). If an alarm sounds on this move, then the cop knows the robber's position and proceeds to the image of that vertex on G_{n-1}, as dictated by the copwin strategy. If no alarm sounds, the cop considers the position of the last alarm that sounded. Since any alarmed vertex is adjacent to at most one unalarmed vertex, the alarm that last sounded is on a vertex adjacent to a single unalarmed vertex v, i.e. the robber is on v. The cop proceeds to $F_{n-1}(v)$. Because of the placement of the alarms, the robber's position is known to the cop for the remainder of the game and, thus, the cop wins using the copwin strategy. □

An *independent set* of a graph G is a set of pairwise, nonadjacent vertices of G. In general, a set S of vertices of G is *k-independent* if the distance between any two vertices of S is at least $k+1$. The *k-independence number* of G, which will be denoted $\alpha_k(G)$, is the maximum cardinality of a k-independent set of G.

Let G be a copwin graph, and let S be a 2-independent set of G. Clearly, if all vertices in $G-S$ are alarmed, then all vertices of G are either alarmed or controlled, and no alarmed vertex has more than one unalarmed neighbour. Thus, this placement of alarms satisfies the conditions given in the preamble to Theorem 1. We have proven Theorem 2.

Theorem 2 *Let G be a finite, copwin graph. Then $A(G) \leq |V(G)| - \alpha_2(G)$.*

Example. Consider the copwin graph G shown in Figure 1. The unalarmed vertices are circled. Theorem 1 gives $A(G) \leq 4$. Note that the unalarmed vertices form a 2-independent set of G.

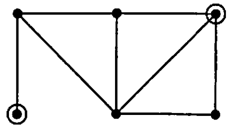

Figure 1: A graph G with $A(G) \leq 4$.

We can improve upon the bounds given by Theorems 1 and 2.

Consider a copwin graph G with copwin ordering $(x_1, x_2, \ldots, x_n)$. A path $x_{i_1}, x_{i_2}, \ldots, x_{i_m}$ of G, $m \geq 2$, is *collapsible* if $f_{i_1}(x_{i_1}) = x_{i_2}$, $f_{i_2}(x_{i_2}) = x_{i_3}, \ldots, f_{i_{m-1}}(x_{i_{m-1}}) = x_{i_m}$. Note that a collapsible path P in G is a path in S_{x_n} and, if $x_n \in V(P)$, then x_n is an endpoint of P. Clearly, $1 \leq i_1 < i_2 < \cdots < i_m \leq n$. Also notice that for $k > i_{m-1}$, $F_k(x_{i_j}) = F_k(x_{i_m})$, for $1 \leq j \leq m-1$.

For a fixed copwin ordering $(x_1, x_2, \ldots, x_n)$ of G, define $\alpha(S_{x_n})$ as the minimum number of vertices having alarms such that all vertices of G are either alarmed or controlled, and if an alarmed vertex x has more than one unalarmed neighbour, then either x is adjacent to a finite number of unalarmed leaves, or else x lies on a collapsible path with exactly two controlled vertices.

Define $\alpha_G = \min_{S_v}\{\alpha(S_v) | S_v \text{ is a copwin spanning tree of } G\}$.

Theorem 3 *Let G be a finite, copwin graph. Then $A(G) \leq \alpha_G$.*

Proof. Let $(x_1, x_2, \ldots, x_n)$ be a copwin ordering of G which realizes α_G. Alarms are placed according to the scheme described in the preamble to the theorem. The cop begins on vertex x_n. If no alarm sounds on the robber's first move, then the cop proceeds to visit all unalarmed vertices of the graph. If no alarm goes off during this search, the robber will be captured during this phase of the cop's strategy.

So suppose that an alarm sounds during the cop's search. Unless the cop is adjacent to the vertex on which the alarm sounds (in which case he catches the robber on his next move), the cop returns to x_n, thereby capturing the robber's image on G_n.

Now suppose that the cop has captured the image of the robber on G_k. If an alarm sounds on the robber's next move, then the cop knows on which vertex the robber is located and moves to the image of that vertex on G_{k-1}. If no alarm sounds, then there are several cases to consider.

Suppose the vertex, x_i say, on which an alarm last sounded is adjacent to exactly one controlled vertex, v say. The cop moves to $F_{k-1}(v)$.

Suppose x_i is adjacent to the unalarmed leaves $x_{i_1}, x_{i_2}, \ldots, x_{i_m}$. If $k-1 > i$, then the cop moves to $F_{k-1}(x_i)$. If $k-1 = i$, then the cop moves to x_i $(= F_{k-1}(x_{i_l}), 1 \leq l \leq m)$. If $k-1 < i$, then $F_k(x_i) = x_i$, and the robber is apprehended!

Otherwise, x_i is adjacent to the controlled vertices x_i' and x_i'', and x_i', x_i, x_i'' is a collapsible path, with $x_i' \to x_i$ and $x_i \to x_i''$. If $k-1 > i$, then the cop moves to $F_{k-1}(x_i'')$. If $k-1 \leq i$, then the cop is on x_i'' if $k-1 = i$, and on x_i if $k-1 < i$. In both cases, the cop moves to $F_{k-1}(x_i')$.

In all cases, the cop captures the image of the robber on G_{k-1}. Thus, the cop wins using the copwin strategy. □

Example. Consider the copwin graph G shown in Figure 2. A copwin spanning tree of G is shown in bold. The unalarmed vertices are circled. Notice that vertices b and d both have two non-leaf, unalarmed neighbours, but a, b, c and e, d, c are collapsible paths.

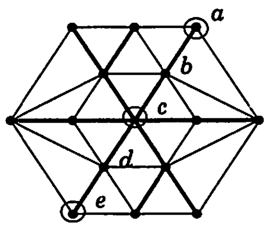

Figure 2: An arrangement of alarms that illustrates how an alarmed vertex can be adjacent to two non-leaf, unalarmed vertices.

Given some arrangement of alarms on the vertices of G, a path P of G is a *freepath* if all vertices of P are unalarmed. A freepath P' of G is *isolated* if all vertices of P' are at distance at least three from any unalarmed vertex not in P'. A collapsible freepath P'', with vertices $x_{i_1}, x_{i_2}, \ldots, x_{i_m}$, is *almost isolated* if P'' is isolated in $G - L(x_{i_2})$, where $L(x_{i_2})$ denotes the set of leaves adjacent to x_{i_2}. (Note that in [3, 4], a set of such freepaths is termed a packing of freepaths. We return to this terminology at the end of the section.) Assuming P'' is an induced subgraph of G, we refer to the subgraph induced by $V(P'') \cup L(x_{i_2})$ as a *free area*.

For a fixed copwin ordering $(x_1, x_2, \ldots, x_n)$ of G, define $\aleph(S_{x_n})$ as the minimum number of vertices having alarms such that

1. free vertices are either controlled or form collapsible freepaths (or free areas),

2. freepaths are induced subgraphs of G,

3. freepaths are (either isolated or) almost isolated, and

4. if an alarmed vertex x_i has more than one unalarmed neighbour, then either

   (a) x_i is adjacent to a finite number of unalarmed leaves,

   (b) x_i lies on a collapsible path with exactly two controlled vertices,

   (c) the unalarmed neighbours of x_i lie on the same freepath $P = x_{i_1}, x_{i_2}, \ldots, x_{i_m}$, and $i > i_m$, or

   (d) x_i has exactly two unalarmed neighbours that lie on the same freepath P.

Define $\aleph_G = \min_{S_v}\{\aleph(S_v) | S_v$ is a copwin spanning tree of $G\}$.

Theorem 4 *Let G be a finite, copwin graph. Then $A(G) \leq \aleph_G$.*

Proof. Let $(x_1, x_2, \ldots, x_n)$ be a copwin ordering of G which realizes $\aleph_G$. Alarms are placed according to the scheme described in the preamble to the theorem. The cop begins on vertex x_n. If no alarm sounds on the robber's first move, then the cop proceeds to visit all unalarmed vertices of the graph, being sure to search the vertices of any free area consecutively. This ensures that the robber cannot move onto previously searched vertices of a freepath undetected, as all freepaths are induced subgraphs of G. If no alarm goes off during this search, the robber will be captured during this phase of the cop's strategy.

So suppose that an alarm sounds during the cop's search. Unless the cop is adjacent to the vertex on which the alarm sounds (in which case he catches the robber on his next move), the cop returns to x_n, thereby capturing the robber's image on G_n.

Now suppose that the cop has captured the image of the robber on G_k. If an alarm sounds on the robber's next move, then the cop knows on which vertex the robber is located and moves to the image of that vertex on G_{k-1}.

If no alarm sounds, then there are several cases to consider. Let x_i be the vertex on which an alarm last sounded. If x_i is adjacent to only one controlled vertex, a finite number of unalarmed leaves, or lies on a collapsible path with exactly two controlled vertices, then the cop proceeds as in the proof of Theorem 3, and captures the robber's image on G_{k-1}.

So suppose x_i is adjacent to vertices in a freepath $P = x_{i_1}, x_{i_2}, \ldots, x_{i_m}$, $m \geq 2$. Further, suppose $i > i_m$. If $i \geq k$, then $F_k(x_i) = x_i$, and the robber is apprehended! Otherwise, $i < k$ and the robber moves from x_i onto P. The cop is on $F_k(x_i)$ and moves to $F_{k-1}(x_{i_1})$ $(= F_{k-1}(x_{i_j}),\ 2 \leq j \leq m)$, and stays with the robber's image. Note that if the cop is on $F_j(x_{i_1})$, $j = 2, 3, \dots, k-1$, and the robber does not move off of P and sound an alarm, then the cop moves to $F_{j-1}(x_{i_1})$ and stays with the robber's image.

Otherwise, $i \leq i_m$ and x_i is adjacent to exactly two vertices of P, x_{i_p} and $x_{i_{p+1}}$ say, with $x_{i_p} \rightarrow x_{i_{p+1}}$. Suppose x_i is retracted onto a vertex of P, i.e. x_i is mapped onto a vertex of P via the retraction map $f_i : V(G_i) \rightarrow V(G_{i+1})$. Then $x_i \rightarrow x_{i_p}$ or $x_i \rightarrow x_{i_{p+1}}$. Suppose $x_i \rightarrow x_{i_{p+1}}$. The other case is similar. If $i_p < k-1$, then the cop is on $F_k(x_i) = F_k(x_{i_{p+1}})$ (unless $F_k(x_i) = x_i$, in which case the robber has already been captured), and moves to $F_{k-1}(x_{i_{p+1}})$. (Note that if the cop is on $F_j(x_{i_q})$, $j = 2, 3, \dots, k-1$, $q = 2, 3, \dots, p+1$, and the robber does not move off of P and sound an alarm, then the cop moves to $F_{j-1}(x_{i_{q-1}})$ and stays with the robber's image.) If $i_p \geq k-1$, then both x_{i_p} and $x_{i_{p+1}}$ are vertices of G_{k-1}. The cop is on $F_k(x_i) = x_{i_{p+1}}$ (unless, once again, $F_k(x_i) = x_i$, in which case the robber has already been captured). If the robber moves to $x_{i_{p+1}}$, he is immediately captured. Otherwise, the robber moves to x_{i_p}, and is captured on the cop's next move.

Suppose now x_i is retracted onto a vertex of $G - P$. If $i_p < k-1$, then the cop is on $F_k(x_i)$, and moves to $F_{k-1}(x_{i_{p+1}})$. (Note again that if the cop is on $F_j(x_{i_q})$, $j = 2, 3, \dots, k-1$, $q = 2, 3, \dots, p+1$, and the robber does not move off of P and sound an alarm, then the cop moves to $F_{j-1}(x_{i_{q-1}})$ and stays with the robber's image.) If $i_p \geq k-1$ then, again, both x_{i_p} and $x_{i_{p+1}}$ are in G_{k-1}. The cop is on $F_k(x_i)$ and moves to $x_{i_{p+1}}$ $(= F_{k-1}(x_{i_{p+1}}))$. Either the robber is immediately captured, or else the cop has lost the robber's image on G_{k-1}. In the latter case, the cop has a modified strategy. If silence follows the robber's moves, the cop moves down P toward the robber. Since P is an induced subgraph of G, the robber will be captured if he does not move off of P and sound an alarm. So suppose the robber moves from $x_{i_l}, l \leq p$, to a vertex $y \notin V(P)$. The cop is on $x_{i_{l+1}}$. Since $x_{i_l} \rightarrow x_{i_{l+1}}$, $N[x_{i_l}] \subseteq N[x_{i_{l+1}}]$, i.e. $y \in N[x_{i_{l+1}}]$. So the cop captures the robber on his next move.

Note that if the cop's strategy takes him to vertex x_{i_1} of P, no further alarms have sounded, and the robber has not been captured, then the cop moves to x_{i_2} and wins by searching the adjacent, unalarmed leaves.

In all cases, the cop captures the image of the robber on G_{k-1}, or captures the actual robber. Thus, the cop wins using (a perhaps modified version of) the copwin strategy. □

Note that for any finite, copwin graph G, $\aleph_G \leq \alpha_G \leq a_G$.

Example. Consider the copwin graph G, shown in Figure 3, with copwin ordering $(x_1, x_2, \dots, x_{19})$. The corresponding copwin spanning tree is shown in bold. The unalarmed vertices are circled. Theorem 4 gives $A(G) \leq 14$.

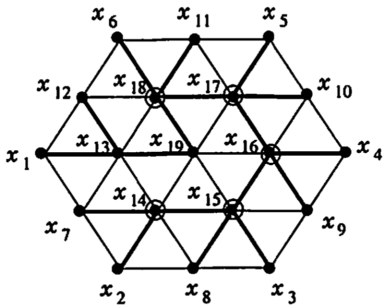

Figure 3: A graph G with $A(G) \leq 14$.

A set of freepaths $\{P_j\}_{j=1}^{n}$ of a graph G is a *packing of freepaths* if, for all $a, b, c \in V(G)$ with $a \sim b \sim c$, $a \in V(P_i)$ and $b, c \notin V(P_i)$, then b and c are alarmed. In [3, 4], it is shown that the alarm number of a tree T is at most $|V(G)|$ less the cardinality of a packing of freepaths on T. The graph G, shown in Figure 4, is used in [3, 4] to show that this result cannot be generalized to copwin graphs. However, a similar result, Theorem 5, follows easily from Theorem 4.

For a fixed copwin ordering $(x_1, x_2, \ldots, x_n)$ of G, define $p_2(S_{x_n})$ as the maximum cardinality of a packing of collapsible freepaths of length at most 1 on G. Define $p_2(G) = \max_{S_v}\{p_2(S_v)|S_v \text{ is a copwin spanning tree of } G\}$. The corresponding placement of alarms satisfies the conditions in the preamble to Theorem 4. We have proven Theorem 5.

Theorem 5 *Let G be a finite, copwin graph. Then $A(G) \leq |V(G)| - p_2(G)$.*

Example. Consider the copwin graph G shown in Figure 4. A copwin spanning tree is shown in bold. The diagonal lines between two branches indicate that every vertex of one branch is adjacent to each of the vertices of the other. The first and last branches (1 and 6) are connected in the same way as branches 2 and 3, 3 and 4, and 5 and 6. The unalarmed vertices are circled, and form a packing of collapsible freepaths of length at most 1.

3 Alarms on Edges

An edge e of G is *controlled* if e remains unalarmed, but all edges incident with e in G are alarmed. Define $L = \{e \in E(G)|e \text{ is a loop of } G\}$. Let $G' = G - L$.

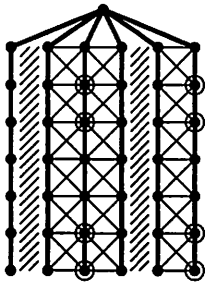

Figure 4: A packing of freepaths of length at most 1 on a copwin graph G.

For a fixed copwin ordering $(x_1, x_2, \ldots, x_n)$ of G, define $a^*(S_{x_n})$ as the minimum number of edges of G having alarms such that all loops $e \in L$ remain unalarmed, all edges in G' are either alarmed or controlled in G' and, for each pair of endpoints x_i, x_j, $i < j$, of an unalarmed edge in G', $f_i(x_i) = x_j$ (i.e. $x_i x_j \in E(S_{x_n})$).

Define $a_G^* = \min_{S_v}\{a^*(S_v) | S_v$ is a copwin spanning tree of $G\}$.

Theorem 6 *Let G be a finite, copwin graph. Then $A^*(G) \leq a_G^*$.*

Proof. Let $(x_1, x_2, \ldots, x_n)$ be a copwin ordering of G which realizes a_G^*. Alarms are placed on the edges of G according to the scheme described in the preamble to the theorem. The cop begins on vertex x_n. The strategy is broken into two phases. During the first phase, the cop forces the robber to sound an alarm to avoid capture. During the second, the cop plays the copwin strategy.

Phase 1. The cop visits all vertices of G, making sure to visit the two endpoints of any unalarmed edge in succession, so that the robber cannot move between these two vertices undetected. It is necessary to search all vertices of G since no vertices are alarmed, and thus the robber may be hiding undetected on any vertex. If no alarm sounds, then the cop will capture the robber during this search. If an alarm sounds during the search, then the cop moves toward the edge e on which the alarm is located.

Case 1. Suppose an alarm sounds on an edge d before the cop reaches e. There are two potentially troublesome cases:

(a) $d \leftrightarrow e$ and, for some controlled edge c, $d \leftrightarrow c \leftrightarrow e$, so that c, d and e form a triangle, and the alarms on d and e did not sound consecutively, and

(b) d and e are each incident with two controlled edges c and c', so that d and e form a 4-cycle with these edges, where opposite sides of the 4-cycle are unalarmed.

If neither of (a) and (b) holds, then the cop knows the robber's location, and proceeds to x_n to begin Phase 2. Otherwise, in order to determine the robber's position, the cop must force the robber to move again. Consider case (a). Let the vertices of the triangle be v_1, v_2 and v_3, where $e = v_1v_2$, $d = v_2v_3$ and $c = v_1v_3$. Note that there may exist an unalarmed edge v_2v_4. The cop searches v_1 and then v_3. If no alarm sounds, the cop searches v_2 and, if applicable, v_4. This forces the robber to move to avoid capture, thereby sounding an alarm. Since the cop now knows the robber's position, he proceeds to x_n. Consider case (b). The cop first searches one of the unalarmed edges, beginning at an endpoint of e. If no alarm sounds, then the cop moves along edge d. If no alarm sounds and the cop does not catch the robber on this move, then the robber is on the remaining unsearched endpoint of e. The cop proceeds to x_n.

Case 2. If no further alarms sound before the cop reaches e, then we consider two cases. If all non-loop edges incident with e are alarmed, then the cop visits the two endpoints of e, thereby forcing the robber to sound another alarm in an attempt to avoid capture. The cop does this, and then proceeds to x_n. If e is incident with controlled edges (there can be at most two, one incident with each endpoint of e), then the cop searches the path consisting of e and these unalarmed edges, thereby forcing the robber to sound a second alarm. If this second alarm satisfies either of properties (a) and (b) above, then the cop proceeds as in Case 1. Otherwise, the cop knows the robber's position and proceeds to x_n.

Phase 2. The cop is on x_n, and has captured the robber's image on G_n. The cop is now able to determine the robber's direction when an alarm sounds.

Suppose that the robber is on vertex v, and the cop is on $F_k(v)$. If an alarm sounds on an edge vv' during the robber's next move, then the cop moves to $F_{k-1}(v')$. Otherwise, the cop must consider the last known position of the robber. This vertex, x_i say, must be an endpoint of the edge where an alarm last sounded. If x_i is not incident with an unalarmed, non-loop edge, then the robber must still be on x_i and the cop moves to $F_{k-1}(x_i)$. Otherwise, x_i is an endpoint of a controlled edge. Let x_j be the second endpoint of this edge, $i < j$. (The case when $i > j$ is similar.) The robber is either on x_i or x_j, and $f_i(x_i) = x_j$. If $k-1 < i$, then the robber is apprehended! If $k-1 = i$, then the cop is on x_j. If the robber is on x_j, he is caught. Otherwise, he is on x_i and is caught on the cop's next move. If $k-1 > i$, then $F_{k-1}(x_i) = F_{k-1}(x_j)$, and the cop moves to $F_{k-1}(x_j)$ and stays with the image of the robber.

Thus, the robber is apprehended using the copwin strategy. □

Example. Consider the copwin graph G shown in Figure 5. One copwin spanning tree is shown in bold. The unalarmed edges are indicated by dashed lines. Theorem 6 gives $A^*(G) \leq 5$. Recall that $|E(G)| = 14$ since G is reflexive.

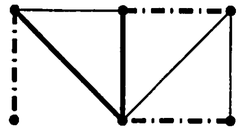

Figure 5: A graph G with $A^*(G) \leq 5$.

Given some arrangement of alarms on the edges of G, a path P of G is now a *freepath* if all edges of P are unalarmed. A freepath P' of G is now *isolated* if all vertices of P' are at distance at least two from the endpoints of any unalarmed edge not in P'. Again, a collapsible freepath P'', with vertices $x_{i_1}, x_{i_2}, \ldots, x_{i_m}$, is *almost isolated* if P'' is isolated in $G - L(x_{i_2})$. Assuming P'' is an induced subgraph of G, we continue to refer to the subgraph induced by $V(P'') \cup L(x_{i_2})$ as a *free area*. A set of freepaths $\{P_i\}_{i=1}^{n}$ is now a *packing of freepaths* if, for all i, P_i is induced, collapsible, and (isolated or) almost isolated.

For a fixed copwin ordering $(x_1, x_2, \ldots, x_n)$ of G, define $\alpha^*(S_{x_n})$ as the maximum cardinality of a packing of freepaths on G.

Define $\alpha_G^* = \max_{S_v}\{\alpha^*(S_v) | S_v \text{ is a copwin spanning tree of } G\}$.

Theorem 7 *Let G be a finite, copwin graph. Then* $A^*(G) \leq |E(G)| - |V(G)| - \alpha_G^*$.

Proof. Let $(x_1, x_2, \ldots, x_n)$ be a copwin ordering of G which realizes α_G^*. Alarms are placed according to the scheme described in the preamble to the theorem. All loops remain unalarmed! The cop begins on vertex x_n. If no alarm sounds on the robber's first move, then the cop proceeds to visit all vertices of the graph, being sure to search the vertices of any free area consecutively. If no alarm goes off during this search, the robber will be captured during this phase of the cop's strategy.

If two distinct alarms sound consecutively, then the robber's position is known to the cop. If two distinct alarms sound, say on edge e and then on edge d, but not consecutively, then there is one potentially troublesome case. (Otherwise, the robber's position is known to the cop.) Suppose $d \leftrightarrow e$, $d = v_1v_2$, $e = v_2v_3$, and there is a (portion of a) freepath P between v_1 and v_3. The cop visits v_2 and then v_1. If the robber is not captured and has not sounded an alarm, giving away his position, then the cop knows that the robber is on (the free area containing) P.

Otherwise, a single alarm sounds, say on edge $f = v_4v_5$. The cop moves toward f. If another alarm sounds, we return to the previous case. Otherwise, the cop visits v_4 and v_5, and either captures the robber or else is able to deduce that the robber is on a freepath P', with $v_4 \in V(P')$ or $v_5 \in V(P')$.

The cop moves to x_n to begin playing the copwin strategy, thereby capturing the robber's image on G_n. Case 1. Suppose first that the robber is known to be on a freepath $P'' = x_{i_1}, x_{i_2}, \ldots, x_{i_m}$, but his position on P'' is unknown. The cop moves to $F_{n-1}(x_{i_1})$, and stays with the image of the robber. In general, if the robber does not move off of P'' and sound an alarm, then the cop is on $F_k(x_{i_1})$ and moves to $F_{k-1}(x_{i_1})$, $k = 2, 3, \ldots, n-1$. Once the robber moves off of P'', the cop proceeds as in Case 2.

Case 2. Now suppose the robber is known to be on v and the cop is on $F_k(v)$. If an alarm goes off on $h = vv'$, the cop moves to $F_{k-1}(v')$. Otherwise, the robber is on a freepath $P'' = x_{i_1}, x_{i_2}, \ldots, x_{i_m}$, i.e. $v = x_{i_p}$, $p \in \{1, 2, \ldots, m\}$. The cop is on $F_k(x_{i_p})$, and moves to $F_{k-1}(x_{i_{p-1}})$.

Note that if the cop's strategy takes him to vertex x_{i_1} of P'', no further alarms have sounded, and the robber has not been captured, then the cop moves to x_{i_2} and wins by searching the adjacent, unalarmed leaves.

Thus, the robber is apprehended using the copwin strategy. □

Example. Consider the copwin graphs G and H, shown in Figure 6, which will be used to illustrate the placement of alarms suggested by Theorem 7. The copwin spanning tree of G used in this example is shown in Figure 4. The copwin spanning tree of H used here is shown in bold. The freepaths are indicated by dashed lines.

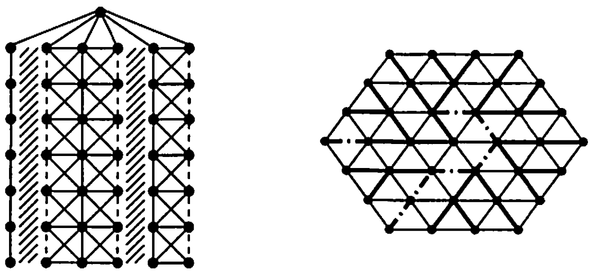

G H

Figure 6: Graphs G and H which illustrate the placement of alarms suggested by Theorem 7.

References

[1] A. Brandstädt, V.B. Le and J.P. Spinrad, *Graph classes: a survey,* SIAM Monographs on Discrete Mathematics and Applications, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1999.

[2] V. Chepoi, *On Distance-Preserving and Domination Elimination Orderings*, SIAM J. Discrete Math. **11** (1998), no. 3, 414-436.

[3] N.E. Clarke, *Constrained Cops and Robber*, Ph.D. Thesis, Dalhousie University, 2002.

[4] N.E. Clarke, *A Game of Cops and Robber Played with Partial Information*, Congressus Numerantium **166** (2004), 145-159.

[5] N.E. Clarke and R.J. Nowakowski, *Cops, Robber, and Traps*, Utilitas Mathematica **60** (2001), 91-98.

[6] N.E. Clarke and R.J. Nowakowski, *Cops, Robber, and Photo Radar*, Ars Combinatoria **56** (2000), 97-103.

[7] R.J. Nowakowski and P. Winkler, *Vertex to Vertex Pursuit in a Graph*, Discrete Math. **43** (1983), 23-29.

[8] T.D. Parsons, *Pursuit Evasion in a Graph*, Lecture Notes in Math. **642**, Springer, Berlin, 1978.

[9] A. Quilliot, *Thèse d'Etat,* Université de Paris VI, 1983.