Pursuit-evasion games and visibility

Danny Dyer

Department of Mathematics and Statistics
Memorial University of Newfoundland

Graphs, groups, and more:
celebrating Brian Alspach's 80th and Dragan Marušič's 65th birthdays
Koper, Slovenia

Speleotopology

An Intuitive Approach To
SPELEOTOPOLOGY

INTRODUCTION

—— by Richard L. Breisch

Speleotopology is the word I coined for the application of topology to the study of caves. To the best of my knowledge this is the first serious work on speleotopology.[1]

The problem is this: A person is lost in a particular cave and is wandering aimlessly. Is there any efficient way for the rescue party to search for the lost person? What is the minimum number of searchers required to explore a cave so that it is impossible to miss finding the victim if he is in the cave?

(Breisch, SW Cavers, 1967)

Two main models:

Edge-searching (or sweeping) is a pursuit-evasion model where a fast, invisible robber that can stop on vertices or edges tries to elude slow, visible cops that move on vertices. Can be thought of as analogous to trying to find a child lost in a cave.

(Parsons, 1978)

Cops and robber is a pursuit-evasion where a slow, visible robber that can only move on vertices tries to elude slow, visible cops, also moving on vertices. Analogous to Pac-Man, or "tag."

(Quilliot, 1978/Nowakowski & Winkler, 1983)

Simultaneous Edge-searching basics:

The cops...

- ...have complete knowledge of the graph.
- ...move slowly, from vertex to vertex.
- ...cannot see the robber.
- ...can all simultaneously move.
- ...can remain in their position.

The robber...

- ...has complete knowledge of the graph.
- ...can move arbitrarily fast, stopping on edges, at any time.
- ...can see the cops.
- ...can remain in its position.

On a graph X, the minimum number of cops needed to guarantee capture of the robber is the **edge-search number**, $s(X)$.

An edge searching example

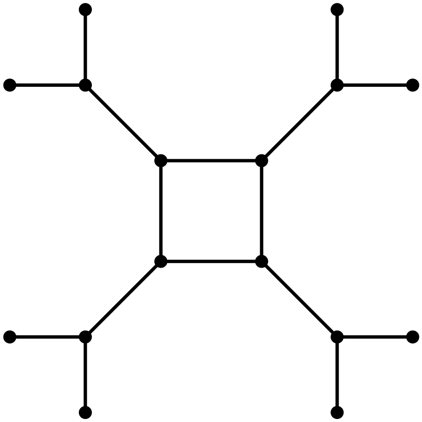

An edge searching example

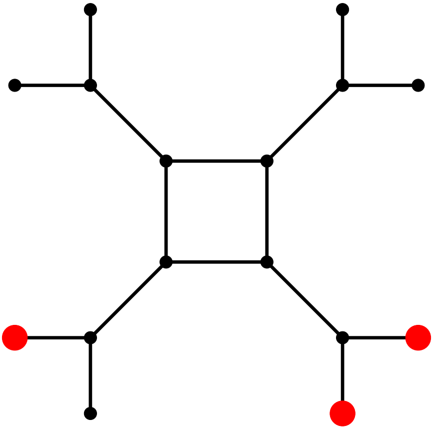

An edge searching example

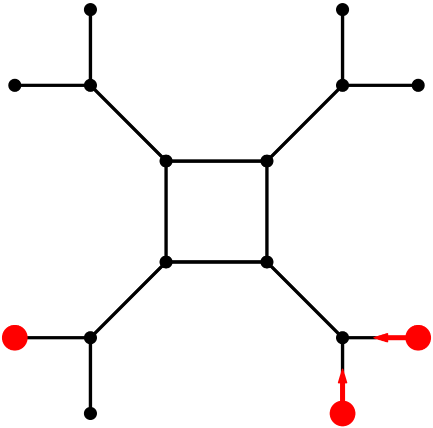

An edge searching example

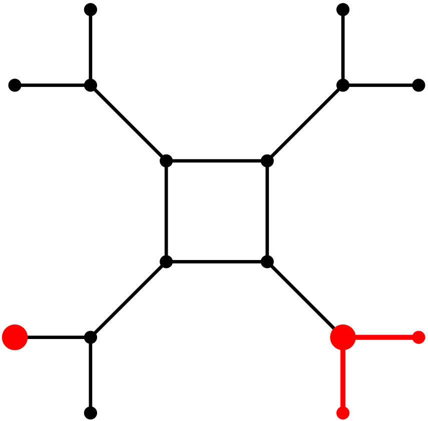

An edge searching example

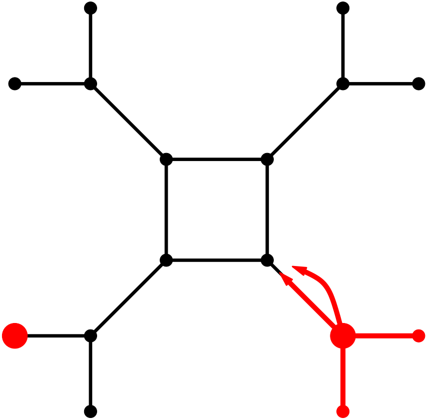

An edge searching example

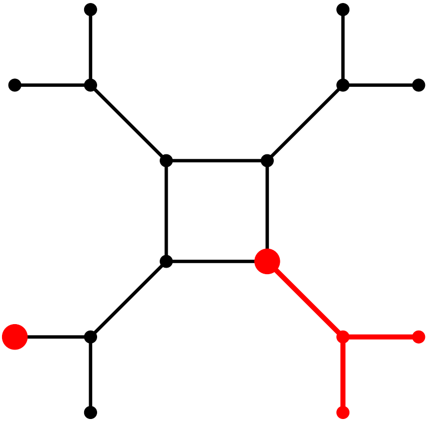

An edge searching example

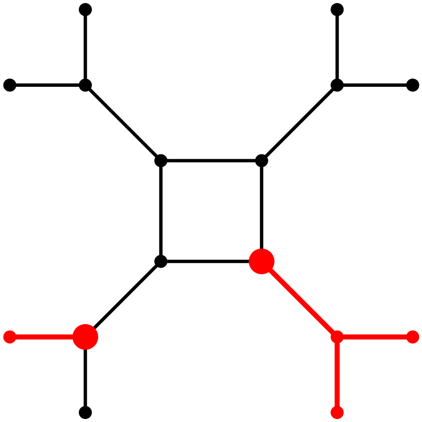

An edge searching example

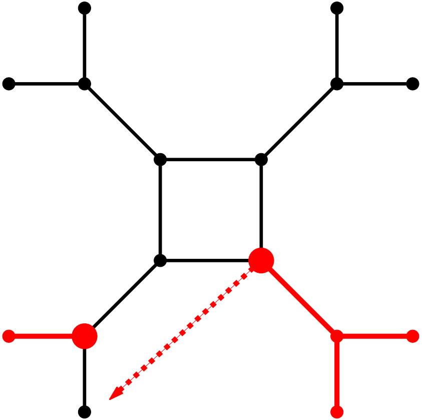

An edge searching example

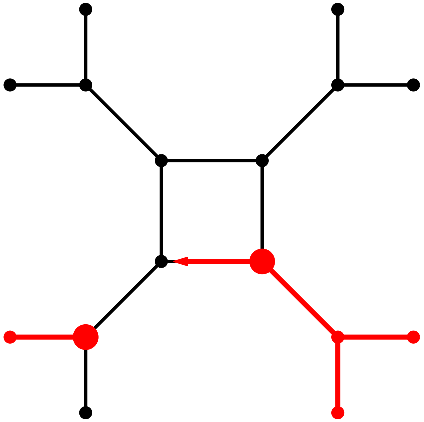

An edge searching example

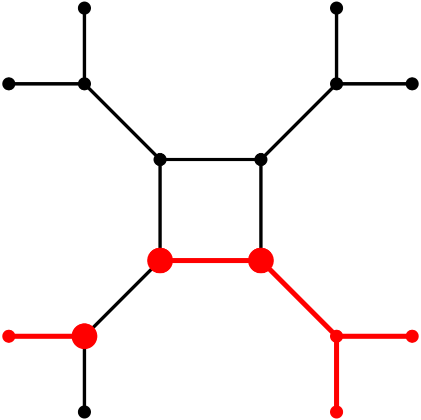

An edge searching example

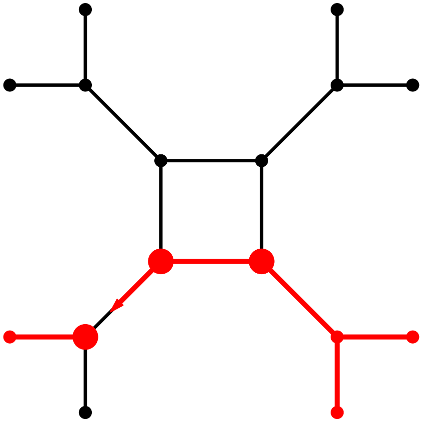

An edge searching example

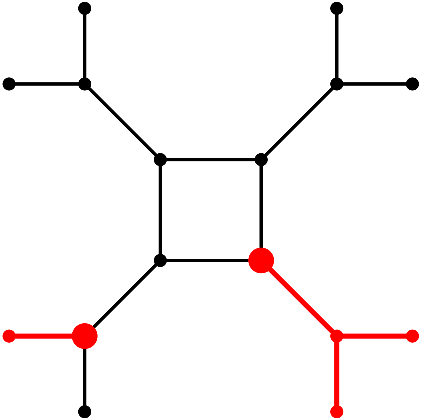

An edge searching example

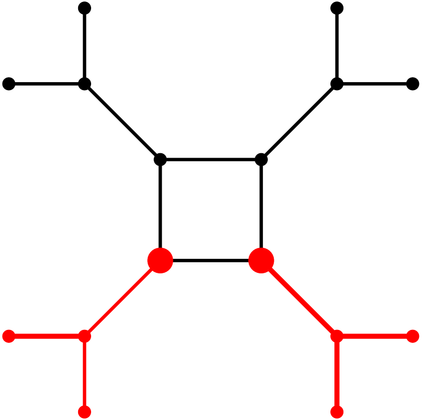

An edge searching example

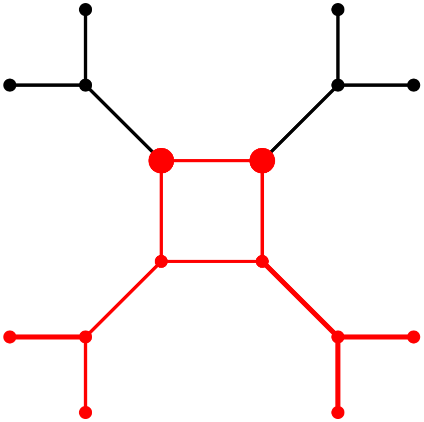

An edge searching example

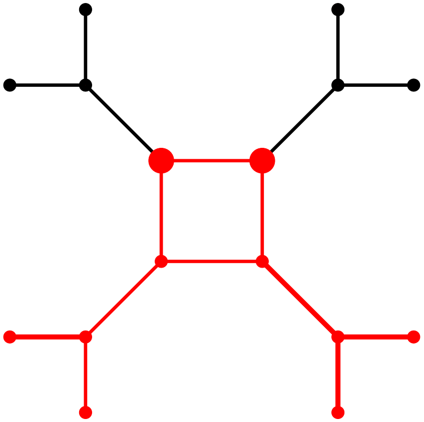

So, $s(X) \leq 3$. In fact, $s(X) = 3$.

The cop and robber model

The cops...

- ...have complete knowledge of the graph.
- ...move slowly, from vertex to vertex.
- ...can see the robber.
- ...can all simultaneously move.
- ...can remain in their position.

The robber...

- ...has complete knowledge of the graph.
- ...moves slowly, from vertex to vertex.
- ...can see the cops.
- ...can remain in its position.

On a graph X, the minimum number of cops needed to guarantee capture of the robber in a finite number of turns is the **cop number** $c(X)$.

A cops and robber example

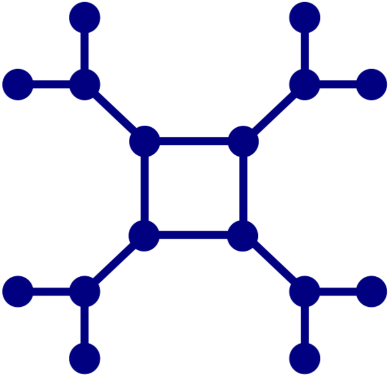

A cops and robber example

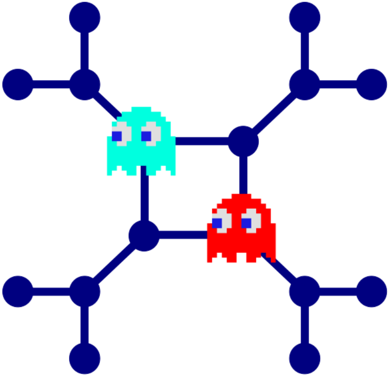

A cops and robber example

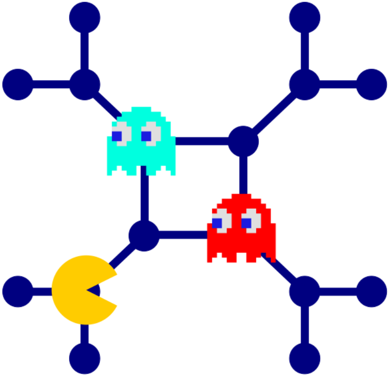

A cops and robber example

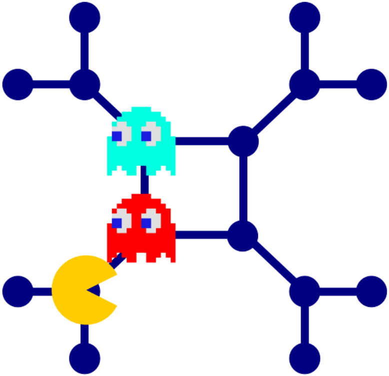

A cops and robber example

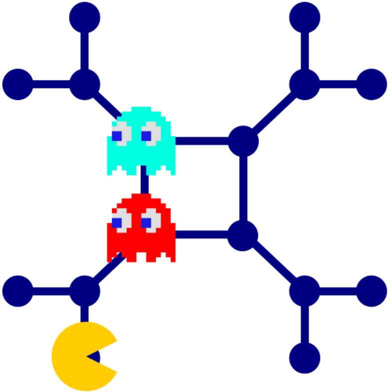

A cops and robber example

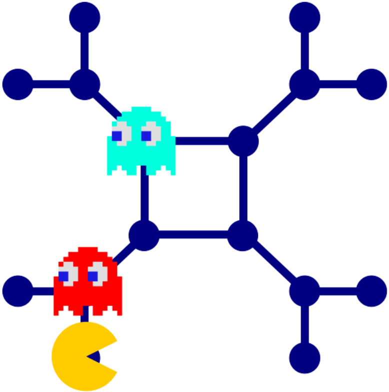

A cops and robber example

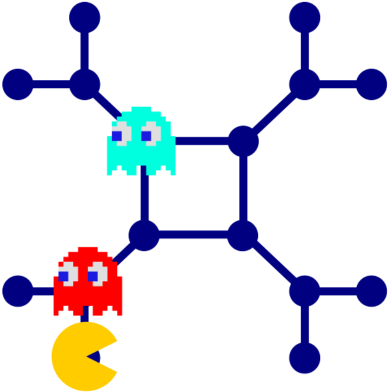

So, $c(X) \leq 2$. In fact, $c(X) = 2$.

Time constraints

Theorem (Alspach, Dyer, Hanson, Yang 2008)

In the cops and robber model, if X is reflexive multigraph on n vertices, then the minimum number of cops needed to guarantee capture of the robber in a single move is $\gamma(X)$.

Theorem (ADHY 2008)

In the simultaneous edge-searching model, if X is a reflexive multigraph, then the minimum number of searchers needed to guarantee capture of the robber in a single move is $|E(X)| + m$, where $n - m$ is the largest order induced bipartite submultigraph of X.

The **zero-visibility** cop and robber model

The cops...

- ...have complete knowledge of the graph.
- ...move slowly, from vertex to vertex.
- ...CANNOT see the robber.
- ...can all simultaneously move.
- ...can remain in their position.

The robber...

- ...has complete knowledge of the graph.
- ...moves slowly, from vertex to vertex.
- ...can see the cops.
- ...can remain in its position.

On a graph X, the minimum number of cops needed to guarantee capture of the robber in a finite number of turns is the zero visibility cop number $c_0(X)$.

Basic differences

$c(K_2) = 1$

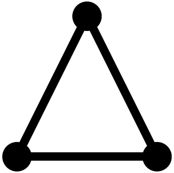

$c(K_3) = 1$

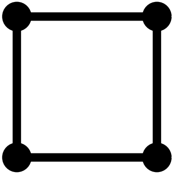

$c(C_4) = 2$

$c_0(K_2) = 1$

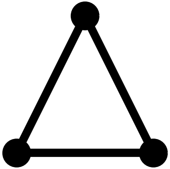

$c_0(K_3) = 2$

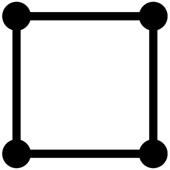

$c_0(C_4) = 2$

Basic differences

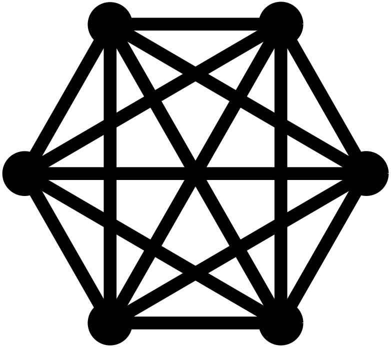

So, $c(K_n) = 1$.

Basic differences

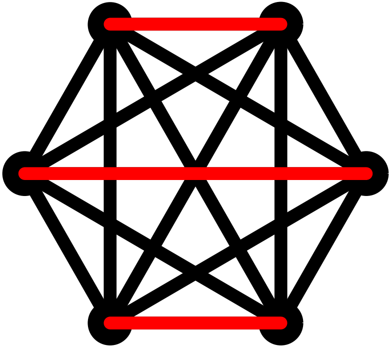

So, $c(K_n) = 1$. But $c_0(K_n) = \lceil \frac{n}{2} \rceil$ – that is, $\frac{c_0(X)}{c(X)}$ can be arbitrarily large.

(Tošić 1985, Tang 2004)

Differences with edge-searching

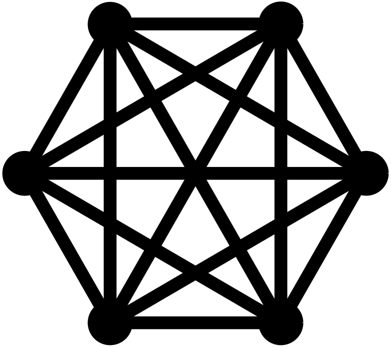

Differences with edge-searching

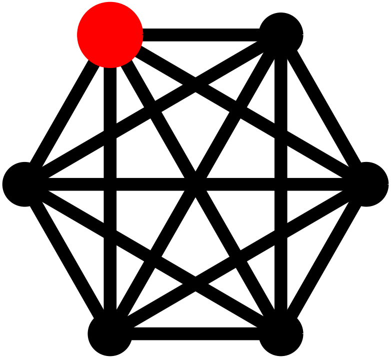

Differences with edge-searching

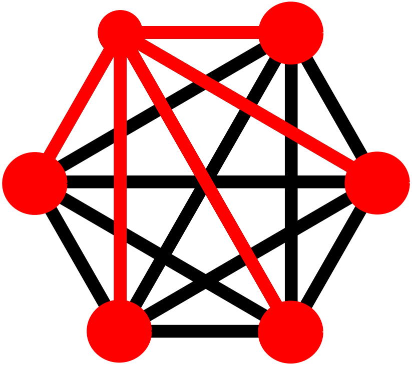

Differences with edge-searching

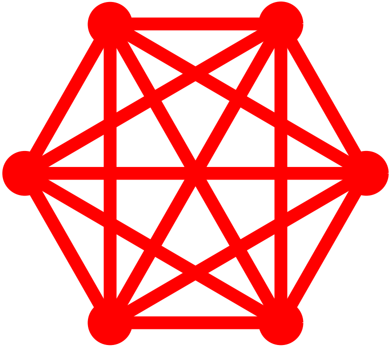

We see $c_0(K_n) = \left\lceil \frac{n}{2} \right\rceil$ and $s(K_n) = n$.

Time constraints and zero visibility

Recall that a minimum edge cover of a graph X is a set $E' \subseteq E(X)$ with the fewest edges for which every vertex of X is an end of at least one edge. We denote size of such a set as $\beta'(X)$.

Theorem (ADHY 2008)

In the zero-visibility cops and robber model, if X is a reflexive multigraph with no isolated vertices, then the minimum number of cops needed to guarantee capture of the robber in a single move is $\beta'(X)$.

The ℓ-**visibility** cop and robber model, $\ell \geq 0$

The cops...

- ...have complete knowledge of the graph.
- ...move slowly, from vertex to vertex.
- ...can see the robber when the distance between the robber and any cop is at most ℓ.
- ...can all simultaneously move.
- ...can remain in their position.

The robber...

- ...has complete knowledge of the graph.
- ...moves slowly, from vertex to vertex.
- ...can see the cops.
- ...can remain in its position.

On a graph X, the minimum number of cops needed to guarantee capture of the robber in a finite number of turns is the ℓ-visibility cop number $c_\ell(X)$.

Trees

A fundamental question:

Is it hard to catch a robber on trees?

Not for classic cops and robber. For edge-searching:

Theorem (Parsons 1978)

Let $k \geq 1$, and T be a tree. Then $s(t) \geq k+1$ if and only if T has a vertex v at which there are three branches T_1, T_2, T_3, satisfying $s(T_j) \geq k$ for $j = 1, 2, 3$.

After creating families of trees $\mathcal{T}_k$, for $k \geq 1$ for which all $T \in \mathcal{T}_k$ have $s(T) = k$, Parsons goes on to prove the following.

Theorem (Parsons 1978)

If $k \geq 2$ and T is a tree, then $s(T) = k$ if and only if T contains a minor from $\mathcal{T}_k$ and none from $\mathcal{T}_{k+1}$.

Trees and low visibility

Define $\mathcal{T}_{k,\ell}$, $\ell \geq 0$, $k \geq 1$, as follows:

1. $\mathcal{T}_{1,\ell} = \{K_1\}$;
2. $\mathcal{T}_{k,\ell}$, $k \geq 2$, is the set of trees, T, that can be formed as follows: let $T_1, T_2, T_3 \in \mathcal{T}_{k-1,\ell}$. Let r_1, r_2, r_3 be vertices of T_1, T_2, T_3 respectively. Then T is formed from the disjoint union of T_1, T_2, T_3, together with paths of length $2\ell + 2$ from each of r_1, r_2, r_3, to a common endpoint, q.

Lemma (Dereniowski, Dyer, Tifenbach, Yang 2015; Cox, Clarke, Duffy, Dyer, Fitzpatrick, Messinger 2018+)

If $T \in \mathcal{T}_{k,\ell}$, then $c_\ell(T) = k$.

Theorem (DDTY 2015; CCDDFM 2018+)

If T is a tree, then $c_\ell(T) = k$ if and only if T contains a minor from $\mathcal{T}_{k,\ell}$ and none from $\mathcal{T}_{k+1,\ell}$.

Monotonicity

Another fundamental question:

Does allowing the robber to return to "cleared" territory ever help?

Not very interesting for the classic cops and robber problem.
Solution is well known for edge-searching.

Theorem (LaPaugh 1993/Bienstock&Seymour 1991)

Every graph X that can be searched with k cops can be monotonically searched with k cops.

Is the ℓ-visibility cop and robber model monotonic?

Obviously not.

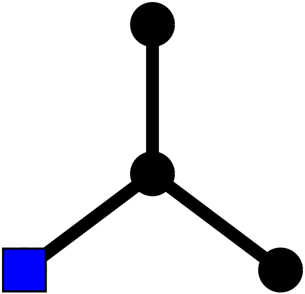

Obviously not.

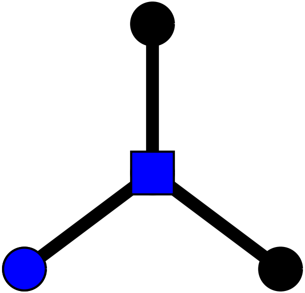

Obviously not.

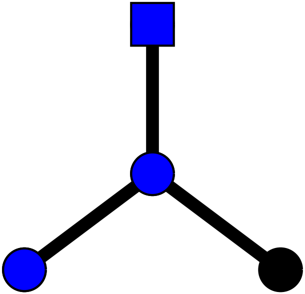

Obviously not.

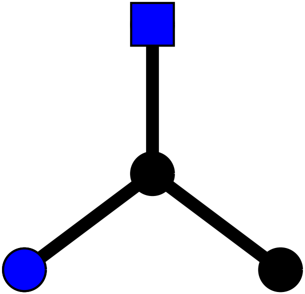

Obviously not.

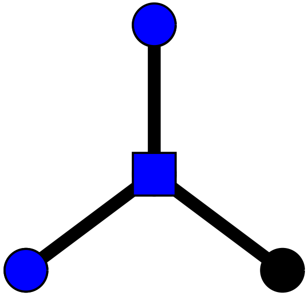

Obviously not.

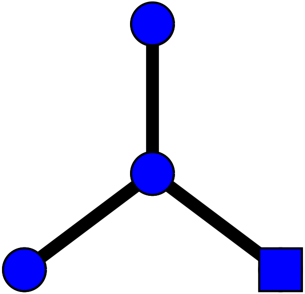

Let's get specific

1. Initially, every vertex is marked as *dirty*.
2. A dirty vertex is *cleaned* if a cop piece occupies it.
3. In between each of the cop's turns, every cleaned vertex that is unoccupied and adjacent to a dirty vertex becomes dirty.

Let X be a graph and let $\mathcal{L}$ be a strategy of length T. For each nonnegative integer $t \leq T$,

1. let $\mathcal{L}_t$ be the set of vertices occupied by cops at the end of the t-th turn by the cops;
2. let $\mathcal{R}_t$ be the set of vertices that are dirty immediately *before* the cop's t-th turn; and
3. let $\mathcal{S}_t$ be the set of vertices that are dirty immediately *after* the cop's t-th turn.

Montonicity, again

A strategy $\mathcal{L}$ is monotonic when

$$\mathcal{R}_0 \supseteq \mathcal{S}_0 \supseteq \mathcal{R}_1 \supseteq \mathcal{S}_1 \supseteq \ldots \supseteq \mathcal{R}_T \supseteq \mathcal{S}_T.$$

This is very restrictive. But can we capture the idea of edge-searching's monotonicity for zero-visibility cops and robber?

Montonicity, again

A strategy $\mathcal{L}$ is monotonic when

$$\mathcal{R}_0 \supseteq \mathcal{S}_0 \supseteq \mathcal{R}_1 \supseteq \mathcal{S}_1 \supseteq \ldots \supseteq \mathcal{R}_T \supseteq \mathcal{S}_T.$$

This is very restrictive. But can we capture the idea of edge-searching's monotonicity for zero-visibility cops and robber?

Weakly monotonic

A strategy of length T is weakly monotonic if for all $t \leq T-1$, we have $\mathcal{S}_{t+1} \subseteq \mathcal{S}_t$.

On a graph X, the minimum number of cops needed to guarantee capture of the robber with a weakly monotonic strategy is the monotonic zero visibility cop number $mc_0(X)$.

$c_0(T)$ versus $mc_0(T)$

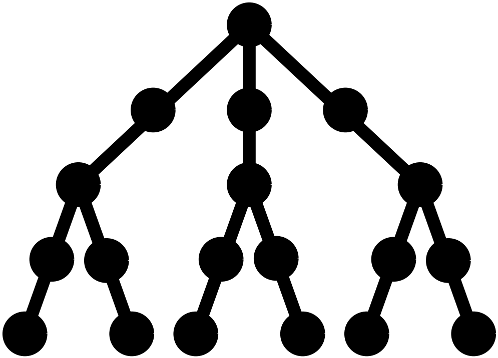

$c_0(T)$ versus $mc_0(T)$

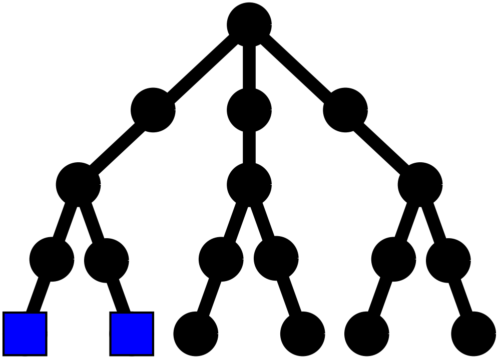

$c_0(T)$ versus $mc_0(T)$

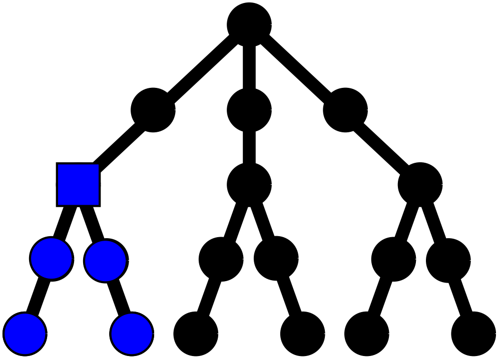

$c_0(T)$ versus $mc_0(T)$

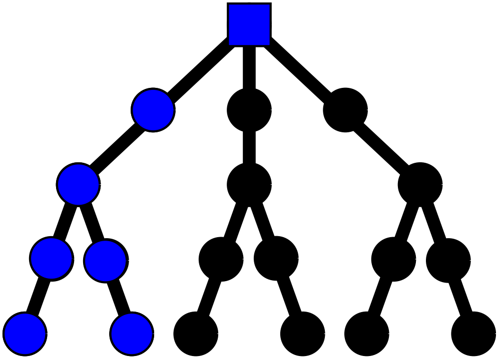

$c_0(T)$ versus $mc_0(T)$

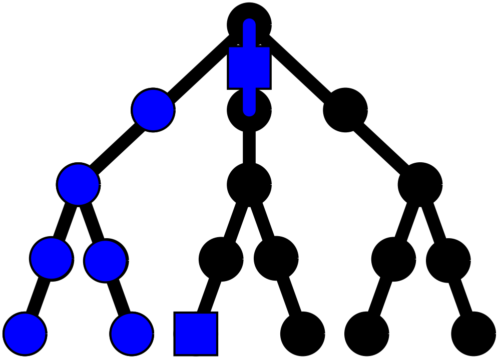

$c_0(T)$ versus $mc_0(T)$

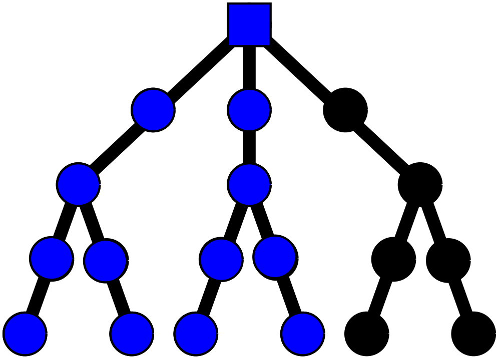

$c_0(T)$ versus $mc_0(T)$

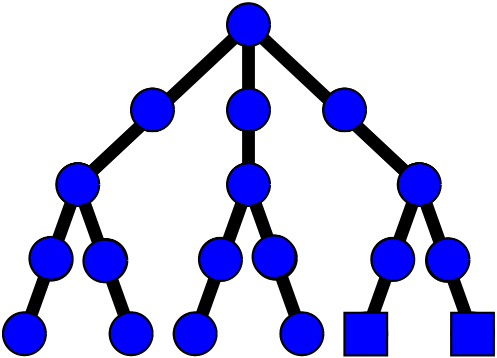

We see that $c_0(T) = 2$, but $mc_0(T) = 3$. (Similarly for $\ell > 0$.)

Pathwidth

Let X be a graph with vertex set V_X. A path decomposition of X is a finite sequence $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2, \ldots, \mathcal{B}_n)$ of sets $\mathcal{B}_i \subseteq V_X$, called bags, such that

1. $\bigcup_{i=1}^{n} \mathcal{B}_i = V_X$;
2. if $x \sim y$, then there is $i \in \{1, \ldots, n\}$ such that $\{x, y\} \subseteq \mathcal{B}_i$; and
3. if $1 \leq i < j < k \leq n$, then $\mathcal{B}_i \cap \mathcal{B}_k \subseteq \mathcal{B}_j$.

Let X be a graph and let $\mathcal{B} = (\mathcal{B}_i)$ be a path decomposition of X. We define the pathwidth of X to be

$$pw(X) = \min \max \{|\mathcal{B}_i| - 1 \mid i \in \{1, \ldots, n\}\},$$

over all possible path decompositions $\mathcal{B}$.

Relations between $c_0(X)$, $mc_0(X)$ and $pw(X)$

Theorem (Dereniowski, Dyer, Tifenbach, Yang 2014)

Let X be a graph. The following are equivalent:

1. *We have $c_0(X) = 1$, $mc_0(X) = 1$ or $pw(X) = 1$.*
2. *We have $c_0(X) = mc_0(X) = pw(X) = 1$.*
3. *The graph X is a caterpillar.*

Theorem (DDTY 2014)

Let X be a connected graph on two or more vertices. Then,

$$c_0(X) \leq pw(X) \leq 2mc_0(X) - 1 \leq 4pw(X) + 1.$$

Differences between $c_0(X)$, $pw(X)$, and $mc_0(X)$

Theorem (DDTY 2014)

For any positive integer k, there is a graph X with $c_0(X) = 2$ and $pw(X) \geq k$.

Proof. Given a graph X, we form the graph X^* by adding a universal vertex to X; a single new vertex is added, together with edges joining this new vertex and every other vertex already present in X.

Let X be a tree on two or more vertices. We will sketch a proof that for some subdivision of H of X, $c_0(H^*) = 2$. Such a graph will have a pathwidth at least that of X.

We proceed by strong induction.

(Proof cont'd)

Let X be a rooted tree with root r. We show that there is a subdivision H of X and a successful zero-visibility strategy on H^* utilising two cops such that

1. a cop visits the universal vertex u at least every second turn throughout the game; and
2. once the root r has been visited by a cop for the first time, either the game is finished or this cop vibrates on the edge ru for the remainder of the game.

Let s be a child of r and let X_2 be the subtree of X consisting of s and its descendants. Let X_1 be the subtree of X consisting of the remaining vertices.

Let H_1 and H_2 be subdivisions of X_1 and X_2 such that H_1^* and H_2^* can be cleaned using two cops subject to the above conditions.

(Proof cont'd)

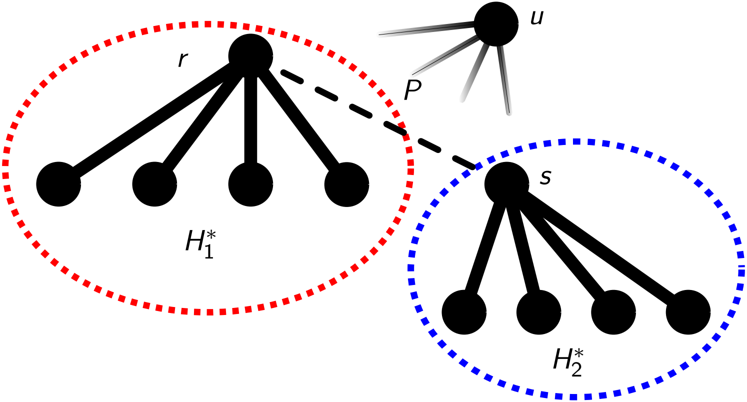

If it takes T moves to clear H_1^*, subdivide P to obtain a path of length $T + 3$.

Open questions

1. For what graphs is $c_0(X) = mc_0(X)$?
2. Are there characterisations of c_0 and mc_0 for trees, unicyclic graphs, planar graphs, series parallel graphs, *etc.*?
3. Can we characterize the graphs with $c_0(X) = 2$? (Already begun by Clarke and Jeliazkova.)
4. For ℓ-visibility cops and robber, what is the difference between 'seeing' or locating the robber, and capturing the robber? (Seeing implies capture on chordal graphs.)

Open questions

1. For what graphs is $c_0(X) = mc_0(X)$?
2. Are there characterisations of c_0 and mc_0 for trees, unicyclic graphs, planar graphs, series parallel graphs, *etc.*?
3. Can we characterize the graphs with $c_0(X) = 2$? (Already begun by Clarke and Jeliazkova.)
4. For ℓ-visibility cops and robber, what is the difference between 'seeing' or locating the robber, and capturing the robber? (Seeing implies capture on chordal graphs.)

Thank you!