Simultaneously moving cops and robbers

G. Konstantinidis, Ath. Kehagias [*]

A B S T R A C T

In this paper we study the *concurrent cops and robber (CCCR)* game. CCCR follows the same rules as the classical, turn-based game, except for the fact that the players move *simultaneously*. The cops' goal is to capture the robber and the *concurrent cop number* of a graph is defined as the minimum number of cops which guarantees capture. For the variant in which it is required to capture the robber in the *shortest possible time*, we let time to capture be the *payoff function* of CCCR; the (game theoretic) *value* of CCCR is the optimal capture time and (cop and robber) *time optimal strategies* are the ones which achieve the value of the game. In this paper we prove the following.

1. For every graph G, the concurrent cop number is equal to the "classical" cop number.
2. For every graph G, CCCR has a value, the cops have an optimal strategy and, for every $\varepsilon > 0$, the robber has an ε-optimal strategy.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

In this paper we study the *concurrent cops and robber* game. In the classical CR game [15,16] each player observes the other player's move before he performs his own. On the other hand, in concurrent CR the players move *simultaneously*. In all other aspects, the concurrent game (henceforth CCCR) follows the same rules as the classical, turn-based game (henceforth TBCR).

The CCCR game (similarly to TBCR) can be considered as either a *game of kind* (the cops' goal is to capture the robber) or a *game of degree* (the cops' goal is to capture the robber *in the shortest possible time*).[1]

This paper is organized as follows. In Section 2 we define preliminary concepts and notation and use these to define the CCCR game rigorously. In Section 3 we concentrate on the "game of kind" aspect: we define the *concurrent cop number* $\widetilde{c}(G)$ and prove that, for every graph G, it is equal to the "classical" cop number $c(G)$. In Section 4 we concentrate on the "game of degree" aspect: we equip CCCR with a *payoff function* (namely the time required to capture the robber) and prove that (a) CCCR has a game theoretic *value*, (b) the cops have an optimal strategy and (c) for every $\varepsilon > 0$ the robber has an ε-optimal strategy; in addition we provide an algorithm for the computation of the value and the optimal strategies. In Section 5 we discuss related work. Finally, in Section 6 we present our conclusions and future research directions.

* Corresponding author.
  E-mail address: kehagiat@gmail.com (Ath. Kehagias).

[1] This terminology is due to Isaacs [8].

2. Preliminaries

In this section, as well as in the rest of the paper, we will mainly concern ourselves with the case of a single cop; this is reflected in the following definitions and notation. In case $K > 1$ cops are considered, this will be stated explicitly; the extension of definitions and notation is straightforward.

2.1. Definition of the CCCR game

Both CCCR and TBCR are played on an undirected, simple and connected graph $G = (V, E)$ by two players called C and R. Player C, controlling K cops (with $K \geq 1$) pursues a single robber controlled by player R (we will sometimes call both the cops and robber *tokens*). We assume the reader is familiar with the rules of TBCR and proceed to present the rules of CCCR for the case of $K = 1$ (a single cop).

1. The game starts from given initial positions: the cop is located at $x_0 \in V$ and the robber at $y_0 \in V$.
2. At the t-th round ($t \in \mathbb{N}$) C moves the cop to $x_t \in N[x_{t-1}]$ and simultaneously R moves the robber to $y_t \in N[y_{t-1}]$.[2]
3. At every round both players know the current cop and robber location (and remember all past locations).
4. A *capture* occurs at the smallest $t \in \mathbb{N}$ for which either of the following conditions holds:
    (a) The cop is located at x_t, the robber is located at y_t, and $x_t = y_t$. This capture condition is the same as in TBCR.
    (b) The cop is located at x_{t-1} and moves to y_{t-1}, while the robber is located at y_{t-1} and moves to x_{t-1}. We will call this "*en passant*" capture; it does not have an analog in TBCR.
5. C wins if capture takes place for some $t \in \mathbb{N}$. Otherwise, R wins. The game analysis becomes easier if we assume that the game always lasts an infinite number of rounds; if a capture occurs at t_c, then we will have $x_t = y_t = x_{t_c}$ for all $t \geq t_c$.

We will denote the above defined game, played on graph $G = (V, E)$ and starting from initial position $(x, y) \in V^2$ by $\Gamma^G_{(x,y)}$. In case the game is played with K cops, it will be denoted by $\Gamma^{G,K}_{(x,y)}$ (in this case $x \in V^K$).

2.2. Nomenclature and notation

The following quantities will be used in the subsequent analysis (once again, we present definitions for the case of $K = 1$). Some of them require two separate definitions: one for TBCR and another for CCCR.

Definition 2.1. A *position* in TBCR is a triple (x, y, P) where $x \in V$ is the cop location, $y \in V$ is the robber location and $P \in \{C, R\}$ is the player whose turn it is to move. We also have $|V| + 1$ additional positions:

1. the position $(\emptyset, \emptyset, C)$ corresponds to the beginning of the game, before either player has placed his token;
2. the positions $(x, \emptyset, R)$, $x \in V$, correspond to the phase of the game in which C has placed the cop but R has not placed the robber.

The set of all TBCR positions is denoted by $S = V \times V \times \{C, R\}$.

Definition 2.2. A *position* in CCCR is a pair $(\widetilde{x}, \widetilde{y})$ where $\widetilde{x} \in V$ is the cop location and $\widetilde{y} \in V$ is the robber location. The set of all CCCR positions is denoted by $\widetilde{S} = V \times V$.

Definition 2.3. A *history* is a position sequence of finite or infinite length. The set of all game histories of *any* finite length is denoted by S^* for TBCR and $\widetilde{S}^*$ for CCCR. The set of all infinite game histories is denoted by S^∞ for TBCR and $\widetilde{S}^\infty$ for CCCR.

In both TBCR and CCCR, the players' *moves* are graph nodes, e.g., $x, y \in V$. Given the current position and the next move (in TBCR) or moves (in CCCR), the next game position is determined by the *transition function*, which encodes the rules of the respective game.

Definition 2.4. In TBCR, the *transition function* $Q : S \times V \to S$ is defined as follows:

when $x = y$: $Q\left((x, y, C), x'\right) = (x, x, R)$

when $x \neq y$ and $x' \in N[x]$: $Q\left((x, y, C), x'\right) = \left(x', y, R\right)$

---

[2] $N[u]$ denotes the *closed neighborhood* of node u, i.e., the set containing u itself and all nodes connected to u by an edge.

when $x \neq y$ and $x' \notin N[x]$: $Q\left((x, y, C), x'\right) = (x, y, R)$

when $x = y$: $Q\left((x, y, R), y'\right) = (x, x, C)$

when $x \neq y$ and $y' \in N[y]$: $Q\left((x, y, R), y'\right) = \left(x, y', C\right)$

when $x \neq y$ and $y' \notin N[y]$: $Q\left((x, y, R), y'\right) = (x, y, C)$.

Definition 2.5. In CCCR, the *transition function* $\widetilde{Q} : \widetilde{S} \times V \times V \to \widetilde{S}$ is defined as follows:

when $x = y$: $\widetilde{Q}\left((x, y), x', y'\right) = (x, x)$

when $x \neq y$ and $x' \in N[x]$ and $y' \in N[y]$

       if $x' = y$ and $y' = x$: $\widetilde{Q}\left((x, y), x', y'\right) = \left(x', x'\right)$

       otherwise: $\widetilde{Q}\left((x, y), x', y'\right) = \left(x', y'\right)$

when $x \neq y$ and $x' \notin N[x]$ and $y' \in N[y]$: $\widetilde{Q}\left((x, y), x', y'\right) = \left(x, y'\right)$

when $x \neq y$ and $x' \in N[x]$ and $y' \notin N[y]$: $\widetilde{Q}\left((x, y), x', y'\right) = \left(x', y\right)$

when $x \neq y$ and $x' \notin N[x]$ and $y' \notin N[y]$: $\widetilde{Q}\left((x, y), x', y'\right) = (x, y)$.

The above rules have the following consequences (which will facilitate our subsequent analysis).

1. The CR game continues for an infinite number of rounds; but if a capture occurs at some time t_c, the cop and robber locations remain fixed for all subsequent times.
2. The transition function accepts "illegal" moves (e.g., $x' \notin N[x]$) as input but "ignores" them, in the sense that they have no influence on the location of tokens.

Roughly speaking, a *strategy* is a rule which, given a finite game history, prescribes a player's next move. In CCCR the players gain an advantage by using *randomized* or *mixed strategies*.

Definition 2.6. A *randomized* or *mixed strategy* is a function $\widetilde{\pi} : \widetilde{S}^* \times V \to [0, 1]$, which satisfies

$$\forall \left((\widetilde{x}_0, \widetilde{y}_0), (\widetilde{x}_1, \widetilde{y}_1), \ldots, (\widetilde{x}_t, \widetilde{y}_t)\right) \in \widetilde{S}^* : \sum_{\widetilde{z} \in V} \widetilde{\pi}\left(\widetilde{z} \mid (\widetilde{x}_0, \widetilde{y}_0), (\widetilde{x}_1, \widetilde{y}_1), \ldots, (\widetilde{x}_t, \widetilde{y}_t)\right) = 1$$

and gives the probability that at time t the player moves into node $\widetilde{z}$, given that the game has started at position $(\widetilde{x}_0, \widetilde{y}_0)$ and progressed through positions $(\widetilde{x}_1, \widetilde{y}_1), \ldots, (\widetilde{x}_t, \widetilde{y}_t)$.

Two classes of strategies will be of special interest to us.

Definition 2.7. A strategy $\widetilde{\pi}$ is called *memoryless* iff

$$\forall \left((\widetilde{x}_0, \widetilde{y}_0), \ldots, (\widetilde{x}_t, \widetilde{y}_t)\right) \in \widetilde{S}^*, \forall \widetilde{z} \in V : \widetilde{\pi}\left(\widetilde{z} \mid (\widetilde{x}_0, \widetilde{y}_0), \ldots, (\widetilde{x}_t, \widetilde{y}_t)\right) = \widetilde{\pi}\left(\widetilde{z} \mid (\widetilde{x}_t, \widetilde{y}_t)\right),$$

i.e., the player's move depends only on the *current* game position.

Definition 2.8. A strategy $\widetilde{\pi}$ is called *deterministic* iff

$$\forall \left((\widetilde{x}_0, \widetilde{y}_0), \ldots, (\widetilde{x}_t, \widetilde{y}_t)\right) \in \widetilde{S} : \exists \widetilde{z} : \widetilde{\pi}\left(\widetilde{z} \mid (\widetilde{x}_0, \widetilde{y}_0), \ldots, (\widetilde{x}_t, \widetilde{y}_t)\right) = 1,$$

i.e., for every game history $(\widetilde{x}_0, \widetilde{y}_0), \ldots, (\widetilde{x}_t, \widetilde{y}_t)$, there is a position $\widetilde{z}$ to which the player will move with certainty.

If $\widetilde{\pi}$ is deterministic, it can be equivalently described by a function $\widetilde{\sigma} : \widetilde{S}^* \to V$ which is determined by $\widetilde{\pi}$ as follows:

$$\widetilde{\sigma}\left((\widetilde{x}_0, \widetilde{y}_0), \ldots, (\widetilde{x}_t, \widetilde{y}_t)\right) = \widetilde{z} \text{ iff } \widetilde{\pi}\left(\widetilde{z} \mid (\widetilde{x}_0, \widetilde{y}_0), \ldots, (\widetilde{x}_t, \widetilde{y}_t)\right) = 1.$$

Similarly, if $\widetilde{\pi}$ is memoryless and deterministic, it can be equivalently described by a function $\widetilde{\sigma} : \widetilde{S} \to V$ which is determined by $\widetilde{\pi}$ as follows:

$$\widetilde{\sigma}\left(\widetilde{x}_t, \widetilde{y}_t\right) = \widetilde{z} \text{ iff } \widetilde{\pi}\left(\widetilde{z} \mid (\widetilde{x}_t, \widetilde{y}_t)\right) = 1.$$

The above definitions and remarks concern CCCR strategies. Regarding TBCR strategies, it is well known [7] that both players lose nothing by restricting themselves to memoryless deterministic strategies of the form $\sigma : S \to V$. In other words,

if player P uses the strategy σ and the current game position is (x, y, P) (which means that it is P's turn to move) P moves his token into node $\sigma(x, y, P)$. Obviously P will only use σ when it is his turn to move; hence we can use the notation $\sigma_C(x, y)$ when talking about *cop strategies* and $\sigma_R(x, y)$ when talking about *robber strategies*. Note that a cop strategy is also defined for the initial position $(\emptyset, \emptyset, C)$ and a robber strategy is also defined for the initial positions $(x, \emptyset, R)$ (for every $x \in V$).

3. Cop numbers

In the "classical" TBCR game we have the following.

Definition 3.1. The cop number $c(G)$ of a graph G is the minimum number of cops sufficient to capture the robber when TBCR is played (optimally by both players) on G.

Note that in the above definition optimal play includes optimal initial placement (in the 0-th round) of the cops and robber on G. On the other hand, in the CCCR game $\Gamma_{(x_0, y_0)}^{G,K}$ the initial cops and robber positions are given (rather than chosen by the players). A reasonable definition of cop number should account for all possible initial positions. Hence we have the following.

Definition 3.2. The concurrent cop number $\widetilde{c}(G)$ of graph G is the minimum number of cops sufficient to ensure capture with probability one for every initial position $(\widetilde{x}_0, \widetilde{y}_0)$, when CCCR is played (optimally by both players) on G.

Note also the expression "capture with probability one" in Definition 3.2. This is different from "certain capture" in the sense that there may exist infinite game histories in which capture does not occur, but the probability of any such infinite history materializing is zero.[3]

In what follows, whenever we mention an arbitrary robber (or cop) move sequence $y_0, y_1, y_2, \ldots$ we assume that it is a *legal* move sequence, i.e., for all t we have $y_{t+1} \in N[y_t]$. Also, if capture occurs at time t_c, the robber's (and cop's) location remains fixed at $y_t = y_{t_c}$, irrespective of the moves $y_{t_c+1}, y_{t_c+2}, \ldots$.

Lemma 3.3. $c(G) = 1 \Rightarrow \widetilde{c}(G) = 1$.

Proof. We select an arbitrary graph G with $c(G) = 1$ and fix it for the rest of the proof. Both TBCR and CCCR will be played on this G. We let $n = |V|$, i.e., n is the number of nodes of G.

We will prove the proposition by constructing a (deterministic and memoryless) cop strategy $\widetilde{\pi}_C^{\#}$ which guarantees, for every starting position, CCCR capture with probability 1.

An essential component of $\widetilde{\pi}_C^{\#}$ is a deterministic cop strategy $\widetilde{\sigma}_C^*$, constructed from another deterministic, memoryless cop strategy σ_C^* which guarantees capture in the TBCR game. Since $c(G) = 1$ we know [7] that such a σ_C^* exists and guarantees capture in at most T rounds, where T depends only on G. Furthermore recall that we have defined TBCR so that after capture takes place both C and R stay in place. The rest of the proof will be divided in two parts.

Part 1. Consider the CCCR game and assume that, for every time t, C knows R's next move (this assumption will be removed in Part 2). Take an arbitrary starting position $\widetilde{s}_0 = (\widetilde{x}_0, \widetilde{y}_0)$ and suppose that at time t, when the position is $(\widetilde{x}_t, \widetilde{y}_t)$, C (knowing that R's next move will be $\widetilde{y}_{t+1}$) plays $\widetilde{x}_{t+1} = \widetilde{\sigma}_C^*(\widetilde{x}_t, \widetilde{y}_{t+1}) = \sigma_C^*(\widetilde{x}_t, \widetilde{y}_{t+1})$.[4] Then, for any robber moves $\widetilde{y}_1, \widetilde{y}_2, \ldots$ in rounds $t = 1, 2, \ldots$ the sequence of game positions will be:

$$(\widetilde{x}_0, \widetilde{y}_0), (\widetilde{x}_1 = \sigma_C^*(\widetilde{x}_0, \widetilde{y}_1), \widetilde{y}_1), (\widetilde{x}_2 = \sigma_C^*(\widetilde{x}_1, \widetilde{y}_2), \widetilde{y}_2), \ldots, (\widetilde{x}_t = \sigma_C^*(\widetilde{x}_{t-1}, \widetilde{y}_t), \widetilde{y}_t), \ldots.$$

We will prove that $\widetilde{x}_T = \widetilde{y}_T$, i.e., capture results in at most T rounds and this will be true with certainty for any starting position $\widetilde{s}_0 = (\widetilde{x}_0, \widetilde{y}_0)$ and robber moves $\widetilde{y}_1, \widetilde{y}_2, \ldots$ thereafter.

To show this, consider a TBCR game in which, at the end of the initial round ($t = 0$) the position is $(x_0, y_0, C) = (\widetilde{x}_0, \widetilde{y}_1, C)$. Further suppose that C uses σ_C^* and R plays the moves $y_1, y_2, \ldots$ with $y_t = \widetilde{y}_{t+1}$ (for $t = 0, 1, \ldots$). Note that, given $y_0 = \widetilde{y}_1$, and also that $\widetilde{y}_1, \widetilde{y}_2, \ldots$ are legal robber moves in CCCR, the resulting robber moves $y_1, y_2, \ldots$ in TBCR are also legal. Moreover recall that, for any given starting position (x_0, y_0), when the robber moves are $y_1, y_2, \ldots$ and C uses σ_C^*, we get a sequence of cop and robber locations of the following form:

$$x_0, y_0, x_1 = \sigma_C^*(x_0, y_0), y_1, x_2 = \sigma_C^*(x_1, y_1), y_2, \ldots, x_t = \sigma_C^*(x_{t-1}, y_{t-1}), y_t, \ldots.$$

---

[3] This point is further discussed in Section 6.

[4] Note that $\widetilde{\sigma}_C^*$ is deterministic and only uses two inputs: one is $\widetilde{x}_t$ (from the previous round) and the other is $\widetilde{y}_{t+1}$ (from the current round). Hence $\widetilde{\sigma}_C^*$ is memoryless in the sense that it only requires knowledge of the immediate past position, but it is also *prescient* in the sense that it requires knowledge of the current robber move.

Given $y_t = \widetilde{y}_{t+1}$ for $t = 0, 1, \ldots, x_0 = \widetilde{x}_0$ and that C uses σ_C^*, we get $x_1 = \sigma_C^*(x_0, y_0) = \sigma_C^*(\widetilde{x}_0, \widetilde{y}_1) = \widetilde{x}_1$, $x_2 = \sigma_C^*(x_1, y_1) = \sigma_C^*(\widetilde{x}_1, \widetilde{y}_2) = \widetilde{x}_2, \ldots, x_t = \sigma_C^*(x_{t-1}, y_{t-1}) = \sigma_C^*(\widetilde{x}_{t-1}, \widetilde{y}_t) = \widetilde{x}_t, \ldots$. Thus the resulting sequence of cop and robber locations in TBCR is:

$$\widetilde{x}_0, \widetilde{y}_1, \widetilde{x}_1, \widetilde{y}_2, \widetilde{x}_2, \widetilde{y}_3, \ldots, \widetilde{x}_t, \widetilde{y}_{t+1}, \ldots. \tag{1}$$

Since σ_C^* guarantees capture by time T in TBCR, we have $x_T = y_T$, irrespective of the moves $y_1, y_2, \ldots$. In fact we will have $x_T = y_{T-1}$, i.e., at most in the first (i.e. cop) phase of round T, C captures R, or else (i.e., if $x_T \neq y_{T-1}$) R can stay put in this round and then $x_T \neq y_T$, which is a contradiction. Since $x_T = \widetilde{x}_T$ and $y_{T-1} = \widetilde{y}_T$ from $x_T = y_{T-1}$ we have $\widetilde{x}_T = \widetilde{y}_T$.

We conclude that also in the CCCR game, for any starting position $\widetilde{s}_0 = (\widetilde{x}_0, \widetilde{y}_0)$ and subsequent robber moves $\widetilde{y}_1, \widetilde{y}_2, \ldots$, capture takes place by the T-th round at the latest. We repeat that this holds under the assumption that: in each round t, C knows R's next move $\widetilde{y}_{t+1}$.

Part 2. In the actual CCCR game C will not know R's next move $\widetilde{y}_{t+1}$; however he can always *guess* $\widetilde{y}_{t+1}$ to be v. Suppose that, when R is at $\widetilde{y}_t$, C guesses with uniform probability $\frac{1}{|N[\widetilde{y}_t]|}$ that R will move to $v \in N[\widetilde{y}_t]$. Let $\widetilde{y}_{t+1}$ be R's actual move at $t+1$ and $\widehat{y}_{t+1}$ be C's guess of that move. We have

$$\Pr(\widehat{y}_{t+1} = v | \widetilde{y}_{t+1} = v) = \frac{1}{|N[\widetilde{y}_t]|} \geq \frac{1}{n}$$

and

$$\Pr(C \text{ guesses R's move correctly})$$
$$= \Pr(\widehat{y}_{t+1} = \widetilde{y}_{t+1}) = \sum_{v \in N[\widetilde{y}_t]} \Pr(\widehat{y}_{t+1} = v | \widetilde{y}_{t+1} = v) \Pr(\widetilde{y}_{t+1} = v) \geq \frac{1}{n} \sum_{v \in N[\widetilde{y}_t]} \Pr(\widetilde{y}_{t+1} = v) = \frac{1}{n}.$$

In other words, C guesses correctly R's next move with probability at least $\frac{1}{n}$. It follows that C guesses correctly R's next T moves (and captures R) with probability at least $\left(\frac{1}{n}\right)^T$.

Now we define the following set of CCCR infinite game histories:

$$\forall k \in \mathbb{N} : A_k = \left\{ \mathbf{s} : \mathbf{s} \in \widetilde{S}^\infty \text{ and R is still free after the first } k \cdot T \text{ rounds} \right\},$$
$$A = \limsup A_k = \cap_{m=1}^\infty \cup_{k=m}^\infty A_k.$$

Since $A_{k+1} \subseteq A_k$ (for all $k \in \mathbb{N}$) we have

$$A = \cap_{m=1}^\infty \cup_{k=m}^\infty A_k = \cap_{m=1}^\infty A_m = \left\{ \mathbf{s} : \mathbf{s} \in \widetilde{S}^\infty \text{ and } \forall m \in \mathbb{N} : \text{R is still free after the first } m \cdot T \text{ rounds} \right\}.$$

In other words, A is the set of all CCCR infinite game histories in which R is never captured. Since

$$\sum_{k=1}^\infty \Pr(A_k) \leq \sum_{k=1}^\infty \left(1 - \left(\frac{1}{n}\right)^T \right)^k < \infty$$

we have (from the first Borel–Cantelli lemma [3]) that $\Pr(A) = 0$.

To sum up: using the deterministic memoryless strategy σ_C^* in conjunction with uniform guessing, we obtain a randomized memoryless strategy $\widetilde{\pi}_C^\#$ which guarantees capture with probability one in the CCCR game; thus $\widetilde{c}(G) = 1$ as claimed. $\square$

Remark 3.4. Another winning strategy for the cop is to always play with uniform probability in his closed neighborhood.

In case of a graph G with $c(G) = K$, it is straightforward to extend the previous argument using K cops. Note however that, at the end of the proof we will know that: if K cops are *required* to capture the robber in TBCR, then K cops *suffice* to capture (with probability one) the robber in CCCR. Hence we have the following.

Lemma 3.5. $c(G) = K \Rightarrow \widetilde{c}(G) \leq K$.

Next we show the "reverse" of Lemma 3.3.

Lemma 3.6. $\widetilde{c}(G) = 1 \Rightarrow c(G) = 1$.

Proof. We select an arbitrary graph G with $\widetilde{c}(G) = 1$ and fix it for the rest of the proof. Both TBCR and CCCR will be played on this G. The Lemma can be stated equivalently as:

$$c(G) > 1 \Rightarrow \widetilde{c}(G) > 1$$

and this is what we will prove.

If $c(G) > 1$ then there exists a (memoryless and deterministic) winning robber strategy σ_R^* for TBCR with one cop on G. More specifically, σ_R^* guarantees that, for every cop starting position x_0, the robber will never be captured.

Choose any $\widetilde{x}_0 \in V$ and let $\widetilde{y}_0 = \sigma_R^*(\widetilde{x}_0, \emptyset)$. Using σ_R^*, we will construct a CCCR robber strategy $\widetilde{\sigma}_R^*$ such that: when CCCR (played on G with a single cop) starts from position $(\widetilde{x}_0, \widetilde{y}_0)$ and R uses $\widetilde{\sigma}_R^*$, the capture probability is zero. This, clearly, implies that $\widetilde{c}(G) > 1$.

It suffices to define $\widetilde{\sigma}_R^*$ only for the case when CCCR starts from $(\widetilde{x}_0, \widetilde{y}_0)$, as follows.

1. In round $t = 1$: $\widetilde{y}_1 = \widetilde{y}_0$ (R stays put);
2. In rounds $t = 2, 3, \ldots$, R plays according to σ_R^*. In other words, if $\widetilde{x}_{t-1} = u$ and $\widetilde{y}_{t-1} = v$, then $\widetilde{y}_t = \widetilde{\sigma}_R^*(u, v) = \sigma_R^*(u, v)$.

Clearly, $\widetilde{\sigma}_R^*$ is not strictly memoryless. The move $\widetilde{y}_1 = \widetilde{y}_0$ depends not only on the game position $(\widetilde{x}_0, \widetilde{y}_0)$ but also on the fact that this is the first round. However, the part of $\widetilde{\sigma}_R^*$ used in rounds $t \geq 2$ is memoryless.

Suppose that in CCCR (starting from $(\widetilde{x}_0, \widetilde{y}_0)$) R plays the strategy $\widetilde{\sigma}_R^*$ while C plays any move sequence $\widetilde{x}_1, \widetilde{x}_2, \ldots$. To prove that capture will never occur, consider a TBCR in which R plays the strategy σ_R^* and C plays the same move sequence $\widetilde{x}_0, \widetilde{x}_1, \ldots$ as in CCCR. Since σ_R^* is winning, capture will never take place in TBCR; as will be shown, this implies capture will never occur in CCCR either and, since this holds for *any* $\widetilde{x}_1, \widetilde{x}_2, \ldots$, we will conclude that $\widetilde{c}(G) > 1$.

Let $y_0, y_1, y_2 \ldots$ be the robber moves occurring in TBCR, given that R plays σ_R^* and C plays $\widetilde{x}_0, \widetilde{x}_1, \ldots$. Let us use $d(u, v)$ to denote the distance of nodes u, v in G, i.e., the length of shortest path between u and v. Obviously we have

$$\forall t \geq 0 : d(\widetilde{x}_{t+1}, y_t) \geq 1 \tag{2}$$

(if we had $d(\widetilde{x}_{t+1}, y_t) = 0$ then σ_R^* would not be a winning strategy). Furthermore

$$\forall t \geq 0 : \widetilde{y}_{t+1} = y_t. \tag{3}$$

Indeed, $\widetilde{y}_1 = y_0$ by construction and if, for some n, we have $\widetilde{y}_n = y_{n-1}$, then

$$\widetilde{y}_{n+1} = \widetilde{\sigma}_R^*(\widetilde{x}_n, \widetilde{y}_n) = \sigma_R^*(\widetilde{x}_n, y_{n-1}) = y_n.$$

From (2) and (3) follows that

$$\forall t : 1 \leq d(\widetilde{x}_t, \widetilde{y}_t). \tag{4}$$

This almost completes the proof that capture never occurs in CCCR. However, we must also consider the possibility of an "en passant" capture, i.e., the case $\widetilde{x}_{t+1} = \widetilde{y}_t$ and $\widetilde{y}_{t+1} = \widetilde{x}_t$. But this would mean

$$d(\widetilde{x}_t, y_t) = d(\widetilde{x}_t, \widetilde{y}_{t+1}) = d(\widetilde{x}_t, \widetilde{x}_t) = 0;$$

in other words, we would have capture in TBCR which contradicts the assumption that σ_R^* is a winning robber strategy. Hence "en passant" capture is also impossible in CCCR. The proof is complete. $\square$

It is straightforward to extend the above for the case of $\widetilde{c}(G) = K$ and obtain the following.

Lemma 3.7. $\widetilde{c}(G) = K \Rightarrow c(G) \leq K$.

Now we can prove our main result.

Theorem 3.8. $c(G) = K \Leftrightarrow \widetilde{c}(G) = K$.

Proof. Assume that $c(G) = K$. By Lemma 3.5 we have $c(G) = K \Rightarrow \widetilde{c}(G) \leq K$; if $\widetilde{c}(G) = K' < K$, then by Lemma 3.7 we have $c(G) \leq K' < K = c(G)$, which is a contradiction. Thus $c(G) = K \Rightarrow \widetilde{c}(G) = K$.

Conversely, assume that $\widetilde{c}(G) = K$. By Lemma 3.7 we have $\widetilde{c}(G) = K \Rightarrow c(G) \leq K$; if $c(G) = K' < K$, then by Lemma 3.5 we have $\widetilde{c}(G) = K' < K = \widetilde{c}(G)$, which is a contradiction. Thus $\widetilde{c}(G) = K \Rightarrow c(G) = K$. $\square$

4. Time optimality

Recall that $\Gamma_{(x_0, y_0)}^G$ denotes the CCCR game played on graph G by a single cop starting at location x_0 and a single robber starting at location y_0. In this section we will only consider $\Gamma_{(x_0, y_0)}^G$; our results can be easily generalized to any game $\Gamma_{(x_0, y_0)}^{G, K}$ (i.e., using $K \geq 1$ cops).

4.1. Existence of value and optimal strategies

To equip $\Gamma_{(x_0,y_0)}^G$ with a *payoff function*, we first define the auxiliary function

$$\forall\, (x,y) \in V^2 : r(x,y) = \begin{cases} 1 & \text{iff } x \neq y \\ 0 & \text{iff } x = y \end{cases}$$

where x and y are cop and robber locations, respectively. Suppose that for every round of $\Gamma_{(x_0,y_0)}^G$ in which the robber remains uncaptured, C pays R one unit of utility and denote by $v_{(x_0,y_0)}^G(\widetilde{\pi}_C, \widetilde{\pi}_R)$ the expected total amount collected by R (obviously it depends on the strategies $\widetilde{\pi}_C, \widetilde{\pi}_R$). Then the *expected robber payoff* in the game $\Gamma_{(x_0,y_0)}^G$ is

$$v_{(x_0,y_0)}^G(\widetilde{\pi}_C, \widetilde{\pi}_R) = \mathbb{E}\left(\sum_{t=0}^{\infty} r(x_t, y_t) \right)$$

where $\mathbb{E}(\cdot)$ denotes expected value and, for notational brevity, the dependence of x_t, y_t on $\widetilde{\pi}_C, \widetilde{\pi}_R$ has been suppressed. The expected payoff to C is $-v_{(x_0,y_0)}^G(\widetilde{\pi}_C, \widetilde{\pi}_R)$.

Following the terminology of [5], we recognize that CCCR equipped with the above payoff is a *positive stochastic game*, R is Player 1 or the *Maximizer* and C is Player 2 or the *Minimizer*. These terms reflect the fact that R (resp. C) chooses $\widetilde{\pi}_R$ (resp. $\widetilde{\pi}_C$) to maximize (resp. to minimize) $v_{(x,y)}^G(\widetilde{\pi}_C, \widetilde{\pi}_R)$. We always have

$$\sup_{\widetilde{\pi}_R} \inf_{\widetilde{\pi}_C} v_{(x,y)}^G(\widetilde{\pi}_C, \widetilde{\pi}_R) \leq \inf_{\widetilde{\pi}_C} \sup_{\widetilde{\pi}_R} v_{(x,y)}^G(\widetilde{\pi}_C, \widetilde{\pi}_R). \tag{5}$$

The following is standard game theoretic terminology [5].

Definition 4.1. If we have

$$\sup_{\widetilde{\pi}_R} \inf_{\widetilde{\pi}_C} v_{(x,y)}^G(\widetilde{\pi}_C, \widetilde{\pi}_R) = \inf_{\widetilde{\pi}_C} \sup_{\widetilde{\pi}_R} v_{(x,y)}^G(\widetilde{\pi}_C, \widetilde{\pi}_R) \tag{6}$$

then we denote the common quantity of (6) by $\widehat{v}_{(x,y)}^G$ and call it the *value* of $\Gamma_{(x,y)}^G$.

Definition 4.2. We denote the *capture time of G* by $CT(G)$ and define it by

$$CT(G) = \max_{(x,y) \in V^2} \widehat{v}_{(x,y)}^G.$$

What is the connection between the games $\Gamma_{(x,y)}^G$ and the values[5] $\widehat{v}_{(x,y)}^G$ for various $(x,y) \in V^2$? It is natural to assume that: if at some stage of $\Gamma_{(x,y)}^G$ we reach the position (x',y') then we can play the remaining portion of $\Gamma_{(x,y)}^G$ as if *we are just starting the game* $\Gamma_{(x',y')}^G$. This plausible assumption can be proved rigorously (see [18] and [5, pp. 89–91]) and has the important consequence that, for a given G and a given position (x,y), $\widehat{v}_{(x,y)}^G$ *exists and is the same for every game* $\Gamma_{(x',y')}^G$ (and hence it is correct to omit mention of a specific game in the notation $\widehat{v}_{(x,y)}^G$). An additional important consequence is the existence of *memoryless optimal strategies which are the same for all* $\Gamma_{(x,y)}^G$ *games*, as will be seen in Theorem 4.4. Before stating and proving this theorem we need some additional definitions.

Definition 4.3. Given $\varepsilon \geq 0$, we say that the cop strategy $\widetilde{\pi}_C^{\varepsilon}$ is *ε-optimal* (for the game $\Gamma_{(x,y)}^G$) iff

$$\left| \widehat{v}_{(x,y)}^G - \sup_{\widetilde{\pi}_R} v_{(x,y)}^G(\widetilde{\pi}_C^{\varepsilon}, \widetilde{\pi}_R) \right| \leq \varepsilon.$$

Similarly, we say that the robber strategy $\widetilde{\pi}_R^{\varepsilon}$ is ε-optimal (for the game $\Gamma_{(x,y)}^G$) iff

$$\left| \widehat{v}_{(x,y)}^G - \inf_{\widetilde{\pi}_C} v_{(x,y)}^G(\widetilde{\pi}_C, \widetilde{\pi}_R^{\varepsilon}) \right| \leq \varepsilon.$$

A 0-optimal (cop or robber) strategy is simply called *optimal*.

---

[5] If they exist.

If both $\widetilde{\pi}_C^*$ and $\widetilde{\pi}_R^*$ are optimal, then we have $\widehat{v}_{(x,y)}^G = v_{(x,y)}^G\left(\widetilde{\pi}_C^*, \widetilde{\pi}_R^*\right)$. The main facts about the $\Gamma_{(x,y)}^G$ games are summarized in the following.

Theorem 4.4. *For every graph $G = (V, E)$ the following hold.*

1. *For every $(x, y) \in V^2$, the game $\Gamma_{(x,y)}^G$ has the value $\widehat{v}_{(x,y)}^G$.*
2. *V^2 can be partitioned into the sets V_1 and V_2 defined by*

$$V_1 = \left\{(x, y) : \widehat{v}_{(x,y)}^G < \infty\right\}, \quad V_2 = \left\{(x, y) : \widehat{v}_{(x,y)}^G = \infty\right\}.$$

If $\widetilde{c}(G) = 1$, then $V_1 = V^2$, i.e., $\widehat{v}_{(x,y)}^G < \infty$ for every $(x, y) \in V^2$.
3. *There exists a memoryless cop strategy $\widetilde{\pi}_C^*$ which is optimal for every game $\Gamma_{(x,y)}^G$.*
4. *If $\widetilde{c}(G) = 1$, then for every $\varepsilon > 0$, there exists a memoryless robber strategy $\widetilde{\pi}_R^\varepsilon$ which is ε-optimal for every game $\Gamma_{(x,y)}^G$.*

Proof. Part 1 follows immediately from Theorems 1 and 2 of [13].

In Part 2, the partition of V^2 into V_1 and V_2, is just a definition. Now suppose that $\widetilde{c}(G) = 1$. To complete the proof of Part 2, it suffices to show the existence of a cop strategy $\widetilde{\pi}_C^{\#}$ and a constant M_G such that

$$\forall \widetilde{\pi}_R, x, y : v_{(x,y)}^G\left(\widetilde{\pi}_C^{\#}, \widetilde{\pi}_R\right) \leq M_G < \infty; \tag{7}$$

in other words, $\widetilde{\pi}_C^{\#}$ guarantees finite (not necessarily optimal) capture time for every robber strategy and every starting position. The required $\widetilde{\pi}_C^{\#}$ is the strategy used in the proof of Lemma 3.3. Indeed, recall that

$$A_k = \left\{\mathbf{s} : \mathbf{s} \in \widetilde{S}^\infty \text{ and } R \text{ is still free after the first } k \cdot T \text{ rounds}\right\}$$

and when C uses $\widetilde{\pi}_C^{\#}$ and R uses any $\widetilde{\pi}_R$ we have

$$\sum_{k=1}^\infty \Pr\left(A_k\right) \leq \sum_{k=1}^\infty \left(1 - \left(\frac{1}{n}\right)^T\right)^k.$$

It follows that

$$\forall \widetilde{\pi}_R : v_{x_0, y_0}^G\left(\widetilde{\pi}_C^{\#}, \widetilde{\pi}_R\right) = \mathbb{E}\left(\sum_{t=0}^\infty r\left(x_t, y_t\right)\right)$$

$$\leq \sum_{k=1}^\infty k \cdot T \cdot \Pr\left(A_k\right) \leq \sum_{k=1}^\infty k \cdot T \cdot \left(1 - \left(\frac{1}{n}\right)^T\right)^k = \widetilde{M}_{G,x,y} < \infty.$$

Letting $M_G = \max_{(x,y) \in V^2} \widetilde{M}_{G,x,y} < \infty$, where M_G depends only on G, we see that $\widetilde{\pi}_C^{\#}$ satisfies (7). Hence we also have:

$$\forall x, y : \widehat{v}_{(x,y)}^G \leq \sup_{\widetilde{\pi}_R} v_{(x,y)}^G\left(\widetilde{\pi}_C^{\#}, \widetilde{\pi}_R\right) \leq M_G < \infty; \tag{8}$$

which shows that $V_1 = V^2$.

Part 3 follows from Theorem 2 of [13].

For Part 4, we have seen that $\widetilde{c}(G) = 1$ implies that $\widehat{v}_{(x,y)}^G < \infty$ for every (x, y); hence the required result follows from Theorem 3.(i) of [13]. □

Remark 4.5. The theorem can be extended to the game $\Gamma_{(x,y)}^{G,K}$ for any graph G (with any $\widetilde{c}(G)$), any number of cops K and any initial position (x, y) (we will now have $x \in V^K$). If $K \geq \widetilde{c}(G)$, then $V_1 = V^{K+1}$. Note that the set V_1 will never be empty; for example, when $K = 1$, (x, x) belongs to V_1 for any $\widetilde{c}(G) \in \mathbb{N}$ (since $v_{(x,x)}^G(\widetilde{\pi}_C, \widetilde{\pi}_R) = 0$ for any $G, x, \widetilde{\pi}_C, \widetilde{\pi}_R$).

Remark 4.6. Parts 1, 2 and 3 of the theorem can also be proved immediately using the results of either [5] or [13].

4.2. Computation of value and optimal strategies

The value and optimal strategies of $\Gamma_{(x,y)}^G$ can be computed by *value iteration*, as shown by Theorem 4.7. Before presenting the theorem and its proof let us give its intuitive justification.

Suppose at time t the game position is (x, y). As already mentioned, we can assume that the "*remainder game*" is $\Gamma_{(x,y)}^G$, i.e., it can be played as a new CCCR game starting at (x, y); the remainder game has value $\widehat{v}_{(x,y)}^G$. Suppose further that C uses the move u and R uses the move v. The new game position is

$$\left(x', y'\right) = \widetilde{Q}\left(\left(x', y'\right), u, v\right),$$

R receives

$$r\left(x', y'\right) = r\left(\widetilde{Q}\left((x, y), u, v\right)\right)$$

units from C and, invoking memorylessness again, the remainder-game is $\Gamma_{(x',y')}^G = \Gamma_{\widetilde{Q}((x,y),u,v)}^G$ and has value $\widehat{v}_{(x',y')}^G = \widehat{v}_{\widetilde{Q}((x,y),u,v)}^G$. To describe the relationship between $\widehat{v}_{(x,y)}^G$ and $\widehat{v}_{\widetilde{Q}((x,y),u,v)}^G$ we need some new notation.

Recall that a finite *two-person zero-sum game in normal form* can be specified by a single $M \times N$ matrix A [17]. The game is played in a single round as follows: *simultaneously* the maximizing Player 1 chooses the row index m and the minimizing Player 2 chooses the column index n; then Player 2 pays to Player 1 the amount A_{mn}. It is well known that every such game has a value and many algorithms are available to compute it. We denote the game matrix A by the notation $\{A_{mn}\}_{m=1,\ldots,M}^{n=1,\ldots,N}$ and its value by **Val** $\left[\{A_{mn}\}_{m=1,\ldots,M}^{n=1,\ldots,N}\right]$.

It seems reasonable (and can be rigorously justified) that $\Gamma_{(x,y)}^G$ can be considered as a single-round finite two-person zero-sum game as follows: when C chooses move u and R chooses move v the expected payoff to R is

$$r\left(\widetilde{Q}\left((x, y), u, v\right)\right) + \widehat{v}_{\widetilde{Q}((x,y),u,v)}^G. \tag{9}$$

In other words, R receives $r\left(\widetilde{Q}\left((x, y), u, v\right)\right)$ units as the payoff of the current round and $\widehat{v}_{\widetilde{Q}((x,y),u,v)}^G$ units as the value of the "remainder-game" $\Gamma_{\widetilde{Q}((x,y),u,v)}^G$ (which is assumed to be played optimally by both players). Hence the game matrix of $\Gamma_{(x,y)}^G$ is $\left\{r\left(\widetilde{Q}\left((x, y), u, v\right)\right) + \widehat{v}_{\widetilde{Q}((x,y),u,v)}^G\right\}_{u \in V}^{v \in V}$ and has value

$$\widehat{v}_{(x,y)}^G = \mathbf{Val}\left[\left\{r\left(\widetilde{Q}\left((x, y), u, v\right)\right) + \widehat{v}_{\widetilde{Q}((x,y),u,v)}^G\right\}_{u \in V}^{v \in V}\right]. \tag{10}$$

Note that (10) implies especially for $x = y$ that $\widehat{v}_{(x,y)}^G = 0$.

The above is an informal argument for the connection between the values $\widehat{v}_{(x,y)}^G$. The following theorem shows that the argument can be made rigorous; furthermore, the theorem provides a method for computing the values, as well as the optimal strategies.

Theorem 4.7. *For every graph* $G = (V, E)$ *with* $\widetilde{c}(G) = 1$ *and for every* $(x, y) \in V^2$ *the values* $\left\{\widehat{v}_{(x,y)}^G\right\}_{(x,y) \in V^2}$ *are the smallest (componentwise) positive solution of the system of* optimality equations*:*

$$\widehat{v}_{(x,y)}^G = \mathbf{Val}\left[\left\{r\left(\widetilde{Q}\left((x, y), u, v\right)\right) + \widehat{v}_{\widetilde{Q}((x,y),u,v)}^G\right\}_{u \in V}^{v \in V}\right] \quad \text{when } x \neq y, \tag{11}$$

$$\widehat{v}_{(x,y)}^G = 0 \quad \text{when } x = y. \tag{12}$$

Furthermore, for $n = 0, 1, 2, \ldots$, *define the initial conditions*

$$\widehat{v}_{(x,y)}^G(0) \geq 0 \quad \text{when } x \neq y \quad \text{and} \quad \widehat{v}_{(x,y)}^G(0) = 0 \quad \text{when } x = y$$

and, for $n \in \mathbb{N}$, *the recursion (*value iteration*)*

$$v_{(x,y)}^G(n+1) = \mathbf{Val}\left[\left\{r\left(\widetilde{Q}\left((x, y), u, v\right)\right) + v_{\widetilde{Q}((x,y),u,v)}^G(n)\right\}_{u \in V}^{v \in V}\right] \quad \text{when } x \neq y, \tag{13}$$

$$v_{(x,y)}^G(n+1) = 0 \quad \text{when } x = y. \tag{14}$$

Then

$$\forall (x, y) \in V^2 : \lim_{n \to \infty} v_{(x,y)}^G(n) = \widehat{v}_{(x,y)}^G.$$

Proof. This is essentially the combination of Theorems 4.4.3 and 4.4.4 from [5], but the following modifications are required.
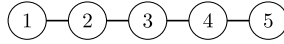
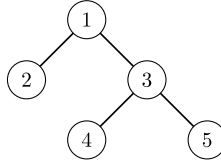**Fig. 1.** A five node path.

Fig. 2. A tree.

In [5] the optimality equations (11)–(12) and the recursion (13)–(14) are given in terms of transition probabilities which in our notation would be written as $P\left((x', y') \mid (x, y), u, v\right)$; this is the probability that the new position is (x', y') given the old position is (x, y) and the player moves are u and v. However, in CCCR transitions are deterministic, i.e.,

$$P\left((x', y') \mid (x, y), u, v\right) = \begin{cases} 1 & \text{when } (x', y') = \widetilde{Q}\left((x, y), u, v\right) \\ 0 & \text{otherwise.} \end{cases}$$

Furthermore, once the game reaches a capture position (x, x), it will always stay in this position, which has value $\widehat{v}^G_{(x,x)} = 0$.

Taking the above in account, the optimality equations and the recursion of [5] reduce to (11)–(14). □

Remark 4.8. The modification of the theorem for the game $\Gamma^{G,K}_{(x,y)}$, with $K > 1$, is obvious.

We conclude this section with some examples. We apply the value iteration (13)–(14) to several graphs and discuss the results.

Example 4.9. In the first example G is a path of five nodes, as illustrated in Fig. 1. Let the (x, y) element of matrix $\widehat{V}^G$ be equal to $\widehat{v}^G_{(x,y)}$, the value of $\Gamma^G_{(x,y)}$ when the cop is at node x and the robber at node y. Value iteration yields

$$\widehat{V}^G = \begin{bmatrix} 0 & 4 & 4 & 4 & 4 \\ 1 & 0 & 3 & 3 & 3 \\ 2 & 2 & 0 & 2 & 2 \\ 3 & 3 & 3 & 0 & 1 \\ 4 & 4 & 4 & 4 & 0 \end{bmatrix}.$$

Cop and robber optimal strategies can be described quite easily: the cop should always move towards the robber and the robber should always move away from the robber.[6] Clearly in this graph the CCCR time-optimal strategies are the same as those for the TBCR game.

Example 4.10. Not surprisingly, for the tree G illustrated in Fig. 2, the optimal cop and robber strategies are again the same for the CCCR and TBCR games.

The value-iteration algorithm yields

$$\widehat{V}^G = \begin{bmatrix} 0 & 1 & 2 & 2 & 2 \\ 3 & 0 & 3 & 3 & 3 \\ 2 & 2 & 0 & 1 & 1 \\ 3 & 3 & 3 & 0 & 2 \\ 3 & 3 & 3 & 2 & 0 \end{bmatrix}.$$

Example 4.11. The next example involves the clique of three nodes, as illustrated in Fig. 3.

After eight iterations, the algorithm yields $\overline{V}^G$ which is (componentwise) within 10^{-2} of the true solution

$$\widehat{V}^G = \begin{bmatrix} 0 & 2 & 2 \\ 2 & 0 & 2 \\ 2 & 2 & 0 \end{bmatrix}.$$

---

[6] Actually the robber has several other optimal strategies; he can also stay in place if the cop is at a distance greater than one. These are pure (i.e., deterministic) strategies; the robber also has mixed optimal strategies.
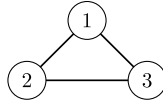
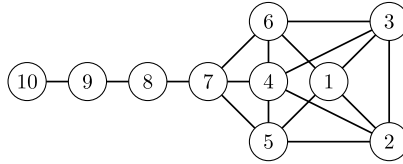**Fig. 3.** A three node clique.

Fig. 4. A Gavenciak graph.

The algorithm also yields the optimal strategies, which are symmetrical with respect to the cop and robber positions (x_0, y_0). For example, when $(x_0, y_0) = (3, 1)$ we have

$$\pi_C^* = \left[\frac{1}{2}, \frac{1}{2}, 0\right] \text{ and } \pi_R^* = \left[0, \frac{1}{2}, \frac{1}{2}\right].$$

In other words, under these strategies, both cop and robber always move with equal probability to one of the two nodes they don't currently occupy. It can be verified analytically that these strategies yield the previously displayed value matrix $\widehat{V}^G$. Because of symmetry, many other optimal strategies exist for both cop and robber.

Example 4.12. The final example involves a Gavenciak graph [6], as illustrated in Fig. 4.

From the results of [6] we know that the TBCR capture time of this graph is 7 (this is the minimax of the capture time over all initial positions). The cop is able to achieve this capture time by first maneuvering himself to node 7 and forcing the robber into the path subgraph, and then chasing the robber all the way to node 10. In the CCCR game the results are similar but they require the use of randomized strategies. We do not present the entire $\widehat{V}^G$ (because of space limitations) but let us give some indicative results. For example, when the initial positions are $(x_0, y_0) = (2, 1)$, the cop cannot be certain of capturing the robber in one move (since the moves are simultaneous). It turns out that, by the application of randomized strategies, the optimal expected capture time is $\widehat{v}_{2,1}^G \cong 18.82 \ldots$ However, the part of the strategies which concerns the path subgraph is, as in TBCR, deterministic. For instance once the cop reaches node 8 (with the robber in either node 9 or 10) he should deterministically perform the transitions $8 \to 9 \to 10$. Let us also note that for this ten-nodes graph, the value iteration algorithm required 90 iterations to get (componentwise) within 10^{-2} of the true solution.

5. Related work

While the assumption of simultaneous moves is a natural one (and is better than turn-based movement as a model of real world pursuit / evasion problems) it appears that CCCR has not been studied in the cops and robber literature. However, our analysis of time optimal CCCR strategies follows closely the corresponding study of time optimal TBCR strategies presented in [7] (and expanded in [4]). Both Hahn's algorithm in [7] and the recursion (13) of Theorem 4.7 are value iteration algorithms. The main difference between the two is this: while in Hahn's algorithm updating the value in every iteration only requires taking a minimum or a maximum, *every value iteration* of (13) requires solving a one-round, zero-sum game (this is indicated by the **Val**[·] operator in (13)). Consequently, (13) is computationally more intensive than Hahn's algorithm.

As already mentioned, simultaneous moves have not been explored in the CR literature. On the other hand, an interesting analog can be found in the literature of *reachability games* [2,14]. As we have pointed out in [11,12], TBCR is a special case of a "classical" (i.e., turn-based) reachability game. Similarly, CCCR is a special case of a *concurrent* reachability game.[7] The literature on concurrent reachability games [1] can furnish useful insights for the analysis of CCCR.

All the above problems can be considered as special cases of the general *stochastic game*. The book [5] is an excellent, comprehensive and relatively recent study of the topic; it also contains many references to important earlier work.

6. Concluding remarks

We conclude this paper by presenting questions which, in our opinion, merit further study.

---

[7] This is the source of our term "concurrent CR game".

One group of questions concern the definition of cop number, which in turn depends on the definition of the properties of the *capture event*. To understand the issue, we must turn back to concurrent reachability games. Let A be the set of all infinite histories of a reachability game and $B \subseteq A$ the set of all infinite histories in which the target state is reached (in CCCR, B would be the set of all infinite histories $\{(x_t, y_t)\}_{t=0}^{\infty}$ for which there is some t_c such that $x_{t_c} = y_{t_c}$). As pointed out in [1], the target state can be reached in at least three different senses (and each of these implies the next one in the list).

1. *Sure reachability*: $B = A$.
2. *Almost sure reachability*: $\Pr(B) = 1$.
3. *Limit sure reachability*: For every real ε, player 1 has a strategy such that for all strategies of player 2, the target state is reached with probability greater than $1 - \varepsilon$.

The above carry over to CCCR and can be used to define corresponding cop numbers: $c_{sure}(G)$, $c_{almostsure}(G)$ and $c_{limitsure}(G)$. In this paper we have worked exclusively with $\widetilde{c}(G) = c_{almostsure}(G)$. Obviously

$$c_{sure}(G) \geq c_{almostsure}(G) \geq c_{limitsure}(G)$$

but several additional questions can be asked. For example can the ratios $\frac{c_{sure}(G)}{c_{almostsure}(G)}$ and $\frac{c_{sure}(G)}{c_{limit\text{-}sure}(G)}$ be bounded by a constant? By a number depending on the size of G? How about the differences $c_{sure}(G) - c_{almostsure}(G)$ and $c_{almostsure}(G) - c_{limitsure}(G)$?

Another group of questions concerns the CCCR variants obtained by modifying the cop's and/or the robber's behavior.

1. For example, what is the *cost of drunkenness*? In other words, what is the ratio of expected capture times for the previously described CCCR game and a variant in which the robber performs a random walk on the nodes of the graph? The same question has been studied for the TBCR case in [9,10].
2. Similarly, what is the *cost of visibility*? In this case we study the ratio of expected capture times for the previously described CCCR game and a variant in which the robber is invisible to the cop. For the TBCR case, this has been studied in [10,11].

References

[1] L. de Alfaro, T.A. Henzinger, O. Kupferman, Concurrent reachability games, Theoret. Comput. Sci. 386 (2007) 188–217.
[2] D. Berwanger, Graph games with perfect information, unpublished manuscript.
[3] P. Billingsley, Probability and Measure, John Wiley & Sons, 2008.
[4] A. Bonato, G. MacGillivray, A general framework for discrete-time pursuit games, unpublished manuscript.
[5] J. Filar, K. Vrieze, Competitive Markov Decision Processes, Springer Science & Business Media, 2012.
[6] T. Gavenčiak, Cop-win graphs with maximum capture-time, Discrete Math. 310 (2010) 1557–1563.
[7] G. Hahn, G. MacGillivray, A note on k-cop, l-robber games on graphs, Discrete Math. 306 (2006) 2492–2497.
[8] R. Isaacs, Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization, Courier Corporation, 1999.
[9] Ath. Kehagias, P. Prałat, Some remarks on cops and drunk robbers, Theoret. Comput. Sci. 463 (2012) 133–147.
[10] Ath. Kehagias, D. Mitsche, P. Prałat, Cops and invisible robbers: the cost of drunkenness, Theoret. Comput. Sci. 481 (2013) 100–120.
[11] Ath. Kehagias, D. Mitsche, P. Prałat, The role of visibility in pursuit/evasion games, Robotics 3 (2014) 371–399.
[12] Ath. Kehagias, G. Konstantinidis, Cops and robbers, game theory and Zermelo's early results, arXiv preprint arXiv:1407.1647, 2014.
[13] P.R. Kumar, T.H. Shiau, Existence of value and randomized strategies in zero-sum discrete-time stochastic dynamic games, SIAM J. Control Optim. 19 (1981) 617–634.
[14] R. Mazala, Infinite games, in: Automata Logics, and Infinite Games, Springer Berlin Heidelberg, 2002, pp. 23–38.
[15] R. Nowakowski, P. Winkler, Vertex to vertex pursuit in a graph, Discrete Math. 43 (1983) 230–239.
[16] A. Quilliot, Jeux et pointes fixes sur les graphes, Ph.D. Dissertation, Universite de Paris VI, 1978.
[17] M.J. Osborne, A. Rubinstein, A Course in Game Theory, MIT press, 1994.
[18] L.S. Shapley, Stochastic games, Proc. Natl. Acad. Sci. USA 39 (1953) 1095–1100.