

Connected Surveillance Game¹

F. Giroire^{a,b}, I. Lamprouc, D. Mazauric^d, N. Nisse^{b,a}, S. Pérennes^{a,b}, R. Soares^{b,a,e}

^a*Univ. Nice Sophia Antipolis, CNRS, I3S, UMR 7271, Sophia Antipolis, France.*

^b*Inria, France.*

^c*Department of Informatics & Telecommunications, University of Athens, Greece.*

^d*Aix-Marseille Université, CNRS, LIF UMR 7279, 13000, Marseille, France.*

^e*ParGO Research Group, UFC, Fortaleza, Brazil.*

Abstract

The *surveillance game* [Fomin *et al.*, 2012] models the problem of web-page prefetching as a pursuit evasion game played on a graph. This two-player game is played turn-by-turn. The first player, called the *observer*, can mark a fixed amount of vertices at each turn. The second one controls a *surfer* that stands at vertices of the graph and can slide along edges. The surfer starts at some initially marked vertex of the graph, its objective is to reach an unmarked node before all nodes of the graph are marked. The *surveillance number* $\text{sn}(G)$ of a graph G is the minimum amount of nodes that the observer has to mark at each turn ensuring it wins against any surfer in G . Fomin *et al.* also defined the *connected surveillance game* where the observer must ensure that marked nodes always induce a connected subgraph. They ask what is the cost of connectivity, i.e., is there a constant $c > 0$ such that the ratio between the *connected surveillance number* $\text{csn}(G)$ and $\text{sn}(G)$ is at most c for any graph G . It is straightforward to show that $\text{csn}(G) \leq \Delta \text{sn}(G)$ for any graph G with maximum degree Δ . Moreover, it has been shown that there are graphs G for which $\text{csn}(G) = \text{sn}(G) + 1$. In this paper, we investigate the question of the cost of the connectivity.

We first provide new non-trivial upper and lower bounds for the cost of connectivity in the surveillance game. More precisely, we present a family of graphs G such that $\text{csn}(G) > \text{sn}(G) + 1$. Moreover, we prove that $\text{csn}(G) \leq \sqrt{\text{sn}(G)n}$ for any n -node graph G . While the gap between these bounds remains huge, it seems difficult to reduce it. We then define the *online surveillance game* where the observer has no *a priori* knowledge of

¹Funded by STINT, Associated Team AIDyNet, project ECOS-Sud Chile

the graph topology and discovers it little-by-little. This variant, which fits better the prefetching motivation, is a restriction of the connected variant. Unfortunately, we show that no algorithm for solving the online surveillance game has competitive ratio better than $\Omega(\Delta)$. That is, while interesting, this variant does not help to obtain better upper bounds for the connected variant. We finally answer an open question [Fomin *et al.*, 2012] by proving that deciding if the surveillance number of a digraph with maximum degree 6 is at most 2 is NP-hard.

Key words: Surveillance game; Cops and robber games; Cost of connectivity; Online strategy; Competitive ratio; Prefetching.

1. Introduction

In this paper, we study two variants of the *surveillance game* introduced in [4]. This two-player game involves one Player moving a mobile agent, called *surfer*, along the edges of a graph, while a second Player, called *observer*, marks the vertices of the graph. The surfer wins if it manages to reach an unmarked vertex. The observer wins otherwise.

Surveillance game. More formally, let $G = (V, E)$ be an undirected simple n -node graph, $v_0 \in V$, and $k \in \mathbb{N}^*$. Initially, the surfer stands at v_0 which is marked and all other nodes are not marked. Then, turn-by-turn, the observer first marks k unmarked vertices and then the surfer may move to a neighbour of its current position. Once a node has been marked, it remains marked until the end of the game. The surfer wins if, at some step, it reaches an unmarked vertex; and the observer wins otherwise. Note that the game lasts at most $\lceil \frac{n}{k} \rceil$ turns. When the game is played on a directed graph, the surfer has to follow arcs when it moves [4]. A k -strategy for the observer from v_0 , or simply a k -strategy from v_0 , is a function $\sigma : V \times 2^V \rightarrow 2^V$ that assigns the set $\sigma(v, M) \subseteq V$ of vertices, $|\sigma(v, M)| \leq k$, that the observer should mark in the *configuration* (v, M) , where $M \subseteq V$, $v_0 \in M$, is the set of already marked vertices and $v \in M$ is the current position of the surfer. We emphasize that σ depends implicitly on the graph G , i.e., it is based on the full knowledge of G . A k -strategy from v_0 is *winning* if it allows the observer to win whatever be the sequence of moves of the surfer starting in v_0 . The *surveillance number* of a graph G with initial node v_0 , denoted by $\text{sn}(G, v_0)$, is the smallest k such that there exists a winning k -strategy starting from v_0 .

Let us define some notations used in the paper. Let Δ be the maximum degree of the nodes in G and, for any $v \in V$, let $N(v)$ be the set of neighbours

of v . More generally, the neighbourhood $N(F)$ of a set $F \subseteq V$ is the subset of vertices of $V \setminus F$ which have a neighbour in F . Moreover, we define the closed neighbourhood of a set F as $N[F] = N(F) \cup F$.

As an example, let us consider the following *basic strategy*: let σ_B be the strategy defined by $\sigma_B(v, M) = N(v) \setminus M$ for any $M \subseteq V$, $v_0 \in M$, and $v \in M$. Intuitively, the basic strategy σ_B asks the observer to mark all unmarked neighbours of the current position of the surfer. It is straightforward, and it was already shown in [4], that σ_B is a winning strategy for any $v_0 \in V$ and it easily implies that $\text{sn}(G, v_0) \leq \max\{|N(v_0)|, \Delta - 1\}$.

Web-page prefetching, connected and online variants. The surveillance game has been introduced because it models the web-page prefetching problem. This problem can be stated as follows. A web-surfer is following the hyperlinks in the digraph of the web. The web-browser aims at downloading the web-pages before the web-surfer accesses it. The number of web-pages that the browser may download before the web-surfer accesses another web-page is limited due to bandwidth constraints. Therefore, designing efficient strategies for the surveillance game would allow to preserve bandwidth while, at the same time, avoiding the waiting time for the download of the web-page the web-surfer wants to access.

By nature of the web-page prefetching problem, in particular because of the huge size of the web digraph, it is not realistic to assume that a strategy may mark any node of the network, even nodes that are “far” from the current position of the surfer. For this reason, [4] defines the *connected* variant of the surveillance game. A strategy σ is said *connected* if $\sigma(v, M) \cup M$ induces a connected subgraph of G for any M , $v_0 \in M \subseteq V(G)$. Note that the basic strategy σ_B is connected. The *connected surveillance number* of a graph G with initial node v_0 , denoted by $\text{csn}(G, v_0)$, is the smallest k such that there exists a winning connected k -strategy starting from v_0 . By definition, $\text{csn}(G, v_0) \geq \text{sn}(G, v_0)$ for any graph G and $v_0 \in V(G)$. In [4], it is shown that there are graphs G and $v_0 \in V(G)$ such that $\text{csn}(G, v_0) = \text{sn}(G, v_0) + 1$. Only the trivial upper bound $\text{csn}(G, v_0) \leq \Delta \text{sn}(G, v_0)$ is known and a natural question is how big the gap between $\text{csn}(G, v_0)$ and $\text{sn}(G, v_0)$ may be [4]. This paper provides a partial answer to this question.

Still the connected surveillance game seems unrealistic since the web-browser cannot be asked to have the full knowledge of the web digraph. For this reason, we define the *online surveillance game*. In this game, the observer discovers the considered graph while marking its nodes. That is, initially, the observer only knows the starting node v_0 and its neighbours.

After the observer has marked the subset M of nodes, it knows M and the vertices that have a neighbour in M and the next set of vertices to be marked depends only on this knowledge, i.e., the nodes at distance at least two from M are unknown. In other words, an *online strategy* is based on the current position of the surfer, the set of already marked nodes and knowing only the subgraph H of the marked nodes and their neighbours (a more formal definition is postponed to Section 3). By definition, the next nodes marked by such a strategy must be known, i.e., adjacent to an already marked vertex. Therefore, an online strategy is connected. We are interested in the competitive ratio of winning online strategies. The competitive ratio $\rho(\mathcal{S})$ of a winning online strategy \mathcal{S} is defined as $\rho(\mathcal{S}) = \max_{G, v_0 \in V(G)} \frac{\mathcal{S}(G, v_0)}{\text{sn}(G, v_0)}$, where $\mathcal{S}(G, v_0)$ denotes the maximum number of vertices marked by \mathcal{S} in G at each turn, when the surfer starts in v_0 . Note that, because any online winning strategy \mathcal{S} is connected, $\text{csn}(G, v_0) \leq \rho(\mathcal{S}) \text{sn}(G, v_0)$ for any graph G and $v_0 \in V(G)$.

1.1. Related work

The surveillance game has mainly been studied in the computational complexity point of view. It is shown that the problem of computing the surveillance number is NP-hard in split graphs [4, 5]. Moreover, deciding whether the surveillance number is at most 2 is NP-hard in chordal graphs and deciding whether the surveillance number is at most 4 is PSPACE-complete. Polynomial-time algorithms that compute the surveillance number in trees and interval graphs are designed in [4]. All previous results also hold for the connected surveillance number. Finally, it is shown that, for any graph G and $v_0 \in V(G)$, $\max_{S \subseteq V(G)} \lceil \frac{|N[S]|-1}{|S|} \rceil \leq \text{sn}(G, v_0) \leq \text{csn}(G, v_0)$ where the maximum is taken over every subset $S \subseteq V(G)$ inducing a connected subgraph with $v_0 \in S$. Moreover, both previous inequalities turn into an equality in case of trees. [4] asks for an example where the inequalities are strict.

In the literature, there are mainly three types of prefetching: server based hints prefetching [2, 1, 8], local prefetching [10] and proxy based prefetching [3]. In local prefetching, the client has no aid from the server when deciding which documents to prefetch. In the server based hints prefetching, the server can aid the client to decide which pages to prefetch. Lastly, in the proxy based prefetching, a proxy that connects its clients with the server decides which pages to prefetch. Moreover, some studies consider that the prefetching mechanism has perfect knowledge of the web-surfer's behaviour [9, 7]. In these studies, the objective is to minimize the waiting

time of the web-surfer with a given bandwidth, by designing good prediction strategies for which pages to prefetch. In the context of prefetching web-pages, the surveillance game is a model to study a local prefetching scheme to guarantee that a websurfer never has to wait a web-page to be downloaded, whilst minimizing the bandwidth necessary to achieve this.

1.2. Our results

In this paper, we study both the connected and online variants of the surveillance game. First, we try to evaluate the gap between non-connected and connected surveillance number of graphs. We give a new upper bound, independent from the maximum degree, for the ratio csn / sn . More precisely, we show that, for any n -node graph G and any $v_0 \in V(G)$, $\text{csn}(G, v_0) \leq \sqrt{\text{sn}(G, v_0)n}$. Then, we describe a family of graphs G such that $\text{csn}(G, v_0) = \text{sn}(G, v_0) + 2$. Note that, contrary to the simple example that shows that connected and not connected surveillance number may differ by one, a larger difference seems much more difficult to obtain.

As mentioned above, the online variant of the surveillance game is a more constrained version of the connected game. We prove that any online strategy has competitive ratio at least $\Omega(\Delta)$. More formally, we describe a family of trees with constant surveillance number such that, for any online winning strategy, there is a step when the strategy has to mark at least $\frac{\Delta}{4}$ vertices. Unfortunately, this shows that the best (up to constant ratio) online strategy is the basic one.

We finish this paper answering an open question in [4] by proving that deciding if the “classical” surveillance number of a graph is at most two is NP-complete even when the graph has maximum degree at most 6.

2. Cost of connectedness

In this section, we investigate the cost of the connectivity constraint. We first prove the first non-trivial upper bound for the ratio csn / sn . More precisely, we show that for any n -node graph G , $\text{csn}(G, v_0) \leq \sqrt{\text{sn}(G, v_0)n}$. Then, we improve the lower bound of [4]. That is, we show a family of graphs where $\text{csn}(G, v_0) > \text{sn}(G, v_0) + 1$. Finally, we disprove a conjecture in [4].

2.1. Upper bound

In this section, we give the first non-trivial upper bound (independent from the degree) of the cost of the connectivity in the surveillance game.

Theorem 1. *Let G be any connected n -node graph and $v_0 \in V(G)$, then*

$$\text{csn}(G, v_0) \leq \sqrt{\text{sn}(G, v_0) \cdot n}.$$

Proof. $\text{sn}(G, v_0) = 1$ if and only if G is a path with v_0 as one of the extremities. In this case, $\text{csn}(G, v_0) = \text{sn}(G, v_0)$ and the result holds.

Assume that $k = \text{sn}(G, v_0) > 1$ and that $n \geq 2$. We describe a connected strategy σ marking at most \sqrt{kn} nodes per turn. Moreover, we prove by induction on the number of turns that σ is connected. Let $M^0 = \{v_0\}$ and let M^t be the set of vertices marked after $t \geq 1$ turns. By the induction hypothesis, let us assume that M^t induces a connected graph of G containing v_0 (it is clearly true for $t = 0$). Let v_t be the vertex occupied by the surfer after turn t . The set $\sigma(v_t, M^t)$ of nodes marked by the observer at step $t + 1$ is defined as follows. If $|V(G) \setminus M^t| \leq \sqrt{kn}$, then let $\sigma(v_t, M^t) = V(G) \setminus M^t$. Otherwise, let $H \subseteq V(G) \setminus M^t$ be such that $|H| = \sqrt{kn}$, $H \cup M^t$ induces a connected subgraph and $|H \cap N(v_t)|$ is maximum. Then, $\sigma(v_t, M^t) = H$, i.e., the strategy marks \sqrt{kn} new nodes in a connected way and, moreover, it marks as many unmarked nodes as possible among the neighbours of v_t . In particular, if $|N(v_t) \setminus M^t| \leq \sqrt{kn}$, then all neighbours of v_t are marked after turn $t + 1$. Moreover, the set $M^{t+1} = M^t \cup \sigma(v_t, M^t)$ is connected in both cases.

Hence, σ is connected and marks at most \sqrt{kn} nodes per turn. We need to show that σ is winning.

For purpose of contradiction, let us assume that the surfer wins against σ by following the path $P = (v_0, \dots, v_t, v_{t+1})$. At its $(t + 1)^{\text{th}}$ turn, the surfer moves from a marked vertex v_t to an unmarked vertex v_{t+1} .

Therefore, $n > t\sqrt{kn}$, otherwise the observer marking \sqrt{kn} nodes at each turn would have already marked every vertex on the graph by the end of turn t . Moreover, by definition of sigma, $|N(v_t) \setminus M^t| > \sqrt{kn}$.

Since, $\text{sn}(G, v_0) = k$, let \mathcal{S} be any k -winning (non necessarily connected) strategy for the observer. Assume that the observer follows \mathcal{S} against the surfer following $P \setminus \{v_{t+1}\}$. Since, \mathcal{S} is winning, all vertices of $N(v_t)$ must be marked after turn t , otherwise the surfer would win by moving to an unmarked neighbour of v_t . Therefore, since \mathcal{S} can mark at most k vertices each turn, $|N(v_t)| \leq kt$.

Taking both inequalities, we have that $\sqrt{kn} < |N(v_t)| \leq kt$. Hence, $\sqrt{n} < t\sqrt{k}$. Since $n > t\sqrt{kn}$ and $\sqrt{n} < t\sqrt{k}$, we have that $t^2k < n < t^2k$, a contradiction. \square

2.2. Lower Bound

This section is devoted to proving the following theorem.

Theorem 2. *There exists a family of graphs such that, for any graph \mathcal{G} of this family, there exists $v_0 \in V(\mathcal{G})$ such that*

$$\text{csn}(\mathcal{G}, v_0) > \text{sn}(\mathcal{G}, v_0) + 1.$$

For this purpose, we define a particular graph, denoted by \mathcal{G} . We first prove that there exists a k -strategy for the observer that is not connected for some k depending on \mathcal{G} (Lemma 3). Then, using two technical results (Claim 4 and Lemma 5), we prove that there are no connected $(k + 1)$ -strategies (Lemma 6). Finally, for completeness, we prove that there exists a connected $(k + 2)$ -strategy for the observer in \mathcal{G} (Lemma 7).

We use the following result proved in [4]. For any graph $G = (V, E)$ and any vertex $v_0 \in V$, a k -strategy for G with initial vertex v_0 is winning if and only if it is winning against a surfer that is constrained to follow induced paths on G . In other words, the walk of the surfer is constrained to be an induced path.

In the following theorem, by *adding a path $P = (v_1, \dots, v_r)$ between two vertices u and v of G* , we mean that the induced path P is added as an induced subgraph of G and the edges $\{u, v_1\}$ and $\{v_r, v\}$ are added.

Let x, α, β and γ be four strictly positive integers satisfying the following.

$$\begin{aligned} \max\{\beta, \frac{\beta}{2} + \gamma + 1\} < \alpha < \min\{\beta + \gamma + 1, 2\gamma + 2\} \text{ and } \beta < 2\gamma + 2 \text{ and} \\ 3x \geq \alpha + \beta + 2\gamma + 12 \text{ and } x > \frac{4}{5}(\alpha + \beta + \gamma) + 10 \text{ and } 2\alpha \geq 73 + \beta + 2\gamma. \end{aligned}$$

For instance, $x = 250, \alpha = 146, \beta = \gamma = 73$ satisfy all the above inequalities.

For proving the main theorem in this section we mainly rely in the family of graphs built in the following procedure described below.

Let $\mathcal{G} = (V, E)$ be a graph with 10 isolated vertices $\{v_0, w_0, w_1, w_2, w'_0, w'_1, w'_2, s_0, s_1, s_2\}$. Then, for all $i \in \{0, 1, 2\}$ do the following:

1. $4x - 9$ vertices of degree one are added and made adjacent to s_i ;
2. $3x - 2$ vertices of degree one are added and made adjacent to w_i , respectively $3x - 2$ neighbours of degree one are added to w'_i ;
3. two disjoint paths $A^i = (a_1^i, \dots, a_\alpha^i)$ and $A'^i = (a_1'^i, \dots, a_\alpha'^i)$ are added between v_0 and s_i ;
4. a path $B^i = (b_1^i, \dots, b_\beta^i)$ is added between v_0 and w_i , and a path $B'^i = (b_1'^i, \dots, b_\beta'^i)$ is added between v_0 and w'_i ;




Figure 1: Graph Family Scheme. Here we show only one “layer” of the graph.

5. for any $j \in \{i, i+1 \bmod 3\}$ a path $C^{i,j} = (c_1^{i,j}, \dots, c_\gamma^{i,j})$ is added between s_j and w_i , and a path $C^{i,j} = (c_1^{i,j}, \dots, c_\gamma^{i,j})$ is added between s_j and w'_i ;
6. for any $1 \leq j \leq \alpha$, $3x-1$ vertices of degree one are added and made adjacent to a_j^i , respectively $3x-1$ neighbours of degree one are added to a_j^i ;
7. for any $1 \leq j \leq \beta$, $3x-1$ vertices of degree one are added and made adjacent to b_j^i , respectively $3x-1$ neighbours of degree one are added to b_j^i ;
8. for any $1 \leq j \leq \gamma$, $\ell \in \{i, i+1 \bmod 3\}$, $3x-1$ vertices of degree one are added and made adjacent to $c_j^{i,\ell}$, respectively $3x-1$ neighbours of degree one are added to $c_j^{i,\ell}$.

The shape of \mathcal{G} is depicted in Figure 1. \mathcal{G} has $(30+18(\alpha+\beta)+36\gamma)x-29$ vertices. For any $i \in \{0, 1, 2\}$, the node s_i has $4x-3$ neighbours, v_0 has 12 neighbours, and any other non-leaf node has degree $3x+1$.

Lemma 3. $\text{sn}(\mathcal{G}, v_0) \leq 3x$.

Proof. To show that $\text{sn}(\mathcal{G}, v_0) \leq 3x$, consider the following strategy for the observer. For any $i \in \{0, 1, 2\}$, in the first step, it marks $x-4$ one-degree neighbours of s_i and the 12 neighbours of v_0 . Then, at any subsequent step, marks all unmarked neighbours of the current position of the surfer.

It is easy to see, by induction on the number of steps that, each time that the surfer arrives at a new node, this node is marked and has at most $3x$ unmarked neighbours. \square

In next claim, we describe some particular Steiner-trees of \mathcal{G} . Intuitively, these Steiner-trees will be used in Lemma 6 to characterize connected strategies.

Claim 4. *If $\max\{\beta, \frac{\beta}{2} + \gamma + 1\} < \alpha < \min\{\beta + \gamma + 1, 2\gamma + 2\}$ and $\beta < 2\gamma + 2$, the unique (up to symmetries) minimum Steiner-tree for $S = N[v_0] \cup \{s_0, s_1, s_2\}$ in \mathcal{G} has $15 + \alpha + \beta + 2\gamma$ vertices and consists of the vertices of the paths $A^0, B^1, C^{1,1}, C^{1,2}$ and the vertices in $S \cup \{w_1\}$.*

Proof. The subgraph induced by the vertices of the paths $A^0, B^1, C^{1,1}, C^{1,2}$ and the vertices in $S \cup \{w_1\}$ is a subtree spanning S and with $15 + \alpha + \beta + 2\gamma$ vertices. Let us enumerate all the possible (up to symmetries) Steiner-trees for S . Consider the subgraph induced by the vertices of:

- A^0, A^1, A^2 and S . The number of vertices in this subgraph is $3\alpha + 13$.
- $A^0, A^1, C^{1,1}, C^{1,2}$ and $S \cup \{w_1\}$. The number of vertices in this subgraph is $2\alpha + 2\gamma + 15$.
- $A^0, A^1, B^1, C^{1,2}$ and $S \cup \{w_1\}$. The number of vertices in this subgraph is $2\alpha + \beta + \gamma + 14$.
- $A^0, C^{0,0}, C^{0,1}, C^{2,0}, C^{2,2}$ and $S \cup \{w_0, w_2\}$. The number of vertices in this subgraph is $\alpha + 4\gamma + 17$.
- $B^0, B^1, C^{0,0}, C^{1,1}, C^{1,2}$ and $S \cup \{w_0, w_1\}$. The number of vertices in this subgraph is $2\beta + 3\gamma + 16$.
- $B^1, C^{1,1}, C^{1,2}, C^{2,2}, C^{2,0}$ and $S \cup \{w_1, w_2\}$. The number of vertices in this subgraph is $\beta + 4\gamma + 17$.

If the subgraph induced by the vertices of the paths $A^0, B^1, C^{1,1}, C^{1,2}$ and the vertices in $S \cup \{w_1\}$, is the unique (up to symmetries) minimum Steiner-

tree for $S = N[v_0] \cup \{s_0, s_1, s_2\}$ in G , then we get the following inequalities:

$$\begin{cases} \alpha > \frac{\beta}{2} + \gamma + 1 \\ \alpha > \beta \\ \alpha > \gamma + 1 \\ \beta < 2\gamma + 2 \\ \alpha < \beta + \gamma + 1 \\ \alpha < 2\gamma + 2. \end{cases}$$

Thus $\max\{\beta, \frac{\beta}{2} + \gamma + 1\} < \alpha < \min\{\beta + \gamma + 1, 2\gamma + 2\}$ and $\beta < 2\gamma + 2$. \square

In Fig. 1, the scheme of a minimum Steiner-tree for $S = N[v_0] \cup \{s_0, s_1, s_2\}$ is depicted with dashed lines.

For any $i \in \{0, 1, 2\}$, let $\mathcal{A}_i = N[v_0] \cup N[A^i] \cup N[s_i]$ (resp., $\mathcal{A}'_i = N[v_0] \cup N[A'^i] \cup N[s_i]$). Note that $|\mathcal{A}_i| = |\mathcal{A}'_i| = (3\alpha + 4)x + 9$ and that the \mathcal{A}_i and \mathcal{A}_j , $i \neq j$, pairwise intersect only in $N[v_0]$.

For any $i \in \{0, 1, 2\}$, let $\mathcal{B}_i = N[v_0] \cup N[B^i] \cup N[w_i] \cup N[C^{i,i}] \cup N[C^{i,i+1 \bmod 3}] \cup N[s_i] \cup N[s_{i+1 \bmod 3}]$ and \mathcal{B}'_i is defined similarly. $|\mathcal{B}_i| = |\mathcal{B}'_i| = (3\beta + 6\gamma + 11)x + 5$. Finally, for any $i \in \{0, 1, 2\}$ and $j \in \{i, i+1 \bmod 3\}$, let $\mathcal{B}_{i,j} = N[v_0] \cup N[B^i] \cup N[w_i] \cup N[C^{i,j}] \cup N[s_j]$ and $\mathcal{B}'_{i,j} = N[v_0] \cup N[B'^i] \cup N[w'_i] \cup N[C^{i,j}] \cup N[s_j]$.

In the next lemma, we characterize the first step of some strategies for \mathcal{G} , showing that several subsets of nodes have to be marked.

Lemma 5. *For any $i \in \{0, 1, 2\}$ and $j \in \{i, i+1 \bmod 3\}$, during its first step, any winning $(3x + y)$ -strategy for \mathcal{G} must mark at least*

- $x + 8 - y(\alpha + 1)$ nodes in \mathcal{A}_i (resp., in \mathcal{A}'_i), and
- $x + 8 - y(\beta + \gamma + 2)$ nodes in $\mathcal{B}_{i,j}$ (resp., in $\mathcal{B}'_{i,j}$), and
- $2x + 4 - y(\beta + 2\gamma + 3)$ nodes in \mathcal{B}_i (resp., in \mathcal{B}'_i).

Proof. Let \mathcal{S} be any winning $(3x + y)$ -strategy and F be the set of nodes that \mathcal{S} marks during its first step.

Let $M = F \cap \mathcal{A}_0$. The surfer goes to a_1^0 . We may assume that \mathcal{S} had marked it since the strategy fails otherwise. Now, the surfer first goes to s_0 through A^0 unless, at some turn, its position has an unmarked neighbour. In the latter case, the surfer goes to this unmarked node and wins. During these $(\alpha + 1)$ steps, the strategy \mathcal{S} can mark at most $(\alpha + 1)(3x + y)$ extra nodes in \mathcal{A}_0 . Hence, in total, at most $|M| + (\alpha + 1)(3x + y)$ nodes have been marked

in \mathcal{A}_0 when the surfer is at s_0 and it is its turn. Because \mathcal{S} is a winning strategy, all nodes in \mathcal{A}_0 must have been marked since otherwise the surfer would have won. Therefore, $|M| + (\alpha + 1)(3x + y) \geq |\mathcal{A}_0 \setminus \{v_0\}| = (3\alpha + 4)x + 8$ and $|M| \geq x + 8 - y(\alpha + 1)$.

The proof is similar for $\mathcal{B}_{i,j}$.

Now, let $M = F \cap \mathcal{B}_0$ and let $M' = F \cap (N[v_0] \cup N[B^0] \cup N[w_0]) \subseteq M$. The surfer goes to b_1^0 . We may assume that \mathcal{S} had marked it since the strategy fails otherwise. Now, the surfer first goes to w_0 through B^0 unless, at some turn, its position has an unmarked neighbour. In the latter case, the surfer goes to this unmarked node and wins. At the turn of the surfer when it is in w_0 , the strategy has marked $|M| + (\beta + 1)(3x + y)$ and all nodes in $N[v_0] \cup N[B^0] \cup N[w_0]$ must have been marked. Therefore, at most $|M| + (\beta + 1)(3x + y) - (12 + 3(\beta + 1)x) = |M| + y(\beta + 1) - 12$ nodes of $\mathcal{B}'_0 \setminus (N[v_0] \cup N[B^0] \cup N[w_0])$ are marked. Therefore, w.l.o.g., there are at most $\lfloor \frac{|M| + y(\beta + 1) - 12}{2} \rfloor$ nodes that are marked in $(N[C^{0,0}] \cup N[s_0]) \setminus N[w_0]$. The surfer now goes from w_0 to s_0 . During these steps, at most $(\gamma + 1)(3x + y)$ new vertices are marked. Because \mathcal{S} is a winning strategy, all nodes in $(N[C^{0,0}] \cup N[s_0]) \setminus N[w_0]$ must have been marked since otherwise the surfer would have won. Therefore, $\lfloor \frac{|M| + y(\beta + 1) - 12}{2} \rfloor + (\gamma + 1)(3x + y) \geq |(N[C^{0,0}] \cup N[s_0]) \setminus N[w_0]| = 3\gamma x + 4x - 4$. Hence, $|M| \geq 2x + 4 - y(\beta + 2\gamma + 3)$. \square

Lemma 6. $\text{csn}(\mathcal{G}, v_0) > 3x + 1$.

Proof. For purpose of contradiction, let us assume that there is a winning connected $(3x + 1)$ -strategy. Let F be the set of vertices marked by this strategy during the first step. Clearly, $N(v_0) \subseteq F$ and $|F| \leq 3x + 1$.

For any $0 \leq i \leq 2$, let $f_i = |F \cap N[s_i]|$ and let $f_{\min} = \min_i f_i$. Without loss of generality, $f_{\min} = f_0$. We first show that $f_{\min} > 3$.

By Lemma 5, for any $i \in \{0, 1, 2\}$, $|F \cap (\mathcal{A}_i \setminus N[v_0])| \geq x - 5 - \alpha$ and, for any $i \in \{0, 2\}$, $|F \cap (\mathcal{B}_{i,0} \setminus N[v_0])| \geq x - 6 - (\beta + \gamma)$ and $|F \cap (\mathcal{B}'_{i,0} \setminus N[v_0])| \geq x - 6 - (\beta + \gamma)$. Therefore,

$$\begin{aligned} 3x + 1 &\geq |F \cap (\mathcal{A}_0 \cup \mathcal{A}'_0 \cup \mathcal{A}_1 \cup \mathcal{A}_2 \cup \mathcal{B}_{0,0} \cup \mathcal{B}_{2,0} \cup \mathcal{B}'_{0,0} \cup \mathcal{B}'_{2,0})| \\ &\geq 12 + 4(x - 5 - \alpha) + 4(x - 6 - (\beta + \gamma)) - 5|F \cap N[s_0]| \\ &\geq 8x - 4(\alpha + \beta + \gamma) - 32 - 5f_{\min} \end{aligned}$$

Hence, $5f_{\min} \geq 5x - 4(\alpha + \beta + \gamma) - 33$, and $f_{\min} \geq x - \frac{4}{5}(\alpha + \beta + \gamma) - 7 > 3$.

Therefore, by definition of f_{\min} , $|F \cap N[s_i]| \geq 4$ for any $i \in \{0, 1, 2\}$. By connectivity of the strategy, $s_i \in F \cap N[s_i]$ for any $i \in \{0, 1, 2\}$. Hence, F must contain a subset of vertices inducing a subtree spanning $S = N[v_0] \cup$

$\{s_0, s_1, s_2\}$. Let T be an inclusion-minimal subset of F that induces a subtree spanning S . By Claim 4, $|T| \geq \alpha + \beta + 2\gamma + 15$. Let $T' = T \setminus (N[v_0] \cup \bigcup_{0 \leq i \leq 2} N[s_i])$. Then, $|T'| \geq \alpha + \beta + 2\gamma - 4$. Moreover, because of the symmetries, we may assume w.l.o.g., that $T' \subseteq \bigcup_{0 \leq i \leq 2} (\mathcal{A}_i \cup \mathcal{B}_i)$.

By Lemma 5 and because $N(v_0) \subseteq F$, for any $0 \leq i \leq 2$, $|F \cap (\mathcal{A}'_i \cup \mathcal{B}'_{i+1 \bmod 3})| \geq x + 8 - (\alpha + 1) + 2x + 4 - (\beta + 2\gamma + 3) - 12 = 3x - (\alpha + \beta + 2\gamma) - 4$. Hence, $|T'| + |F \cap (\mathcal{A}'_i \cup \mathcal{B}'_{i+1 \bmod 3})| \geq 3x - 8$. Let $W_i = F \setminus (\mathcal{A}'_i \cup \mathcal{B}'_{i+1 \bmod 3} \cup T')$. Since $|F| \leq 3x + 1$, it follows that $|W_i| \leq 9$.

Let $f_{\max} = \max_i f_i$ and assume w.l.o.g. that $f_{\max} = f_2$. Since $\sum_{0 \leq i \leq 2} f_i \leq |F \setminus T'|$, we get that $f_0 + f_1 \leq \lfloor \frac{2}{3}(5 + 3x - (\alpha + \beta + 2\gamma)) \rfloor$.

To conclude, $|F \cap \mathcal{B}'_0| = |N(v_0)| + f_0 + f_1 + |W_0| \leq 21 + \lfloor \frac{2}{3}(5 + 3x - (\alpha + \beta + 2\gamma)) \rfloor$. On the other hand, Lemma 5 implies that $|F \cap \mathcal{B}'_0| \geq 2x + 1 - (\beta + 2\gamma)$. Therefore, $22 + \frac{2}{3}(5 + 3x - (\alpha + \beta + 2\gamma)) > 2x + 1 - (\beta + 2\gamma)$ and it follows $73 > 2\alpha - \beta - 2\gamma$. This contradicts the inequalities. \square

Lemmas 3 and 6 are sufficient to prove Theorem 2. More precisely, it shows that there exists a family of graphs G and $v_0 \in V(G)$ such that $\text{csn}(G, v_0) \geq \text{sn}(G, v_0) + 2$. However, Lemma 7 shows that the family of graphs we described does not allow to increase further the cost of connectivity.

Lemma 7. $\text{csn}(\mathcal{G}, v_0) \leq 3x + 2$.

Proof. Consider the following strategy. At the first step, the observer marks the 12 neighbours of v_0 , all nodes of the paths $A^0, B^1, C^{1,1}$ and $C^{1,2}$, the vertices w_1, s_0, s_1 and s_2 and finally $Z = \lfloor (3x - \alpha - \beta - 2\gamma - 12)/3 \rfloor$ one-degree neighbours of each s_i . Note that $Z \geq 0$ by the fact that $3x \geq \alpha + \beta + 2\gamma + 12$.

Then, the strategy goes on as follows. Let $i \in \{0, 1, 2\}$. When the surfer arrives at some node a_j^i (resp., a_j^i), $1 \leq j \leq \alpha$, the observer marks the at most $3x$ unmarked neighbours of a_j^i and marks at least 2 unmarked neighbours of s_i . When the surfer arrives at some node b_j^i (resp., b_j^i), $1 \leq j \leq \beta$, or at w_i , the observer marks the at most $3x$ unmarked neighbours of this node and marks at least 1 unmarked neighbour of s_i and at least 1 unmarked neighbour of $s_{i+1 \bmod 3}$. When the surfer arrives at some node $c_j^{i,\ell}$ (resp., $c_j^{i,\ell}$), $1 \leq j \leq \gamma$, $\ell \in \{i, i+1 \bmod 3\}$, the observer marks the at most $3x$ unmarked neighbours of $c_j^{i,\ell}$ and marks at least 2 unmarked neighbours of s_ℓ (if any) and, if all neighbours of s_ℓ are already marked, the observer marks at least 2 unmarked neighbours of s_k where $\{k\} = \{i, i+1 \bmod 3\} \setminus \{\ell\}$. Finally, when the surfer arrives at s_i , the observer marks $3x + 2$ unmarked neighbours of it.

To prove the validity of this strategy, it is sufficient to show that the surfer will lose for the following three different trajectories. This is sufficient, because the surfer is only able to win when moving from s_0 , s_1 or s_2 and because $\alpha < 2\gamma$, i.e., the amount of steps it takes for the surfer to move from s_i to s_j , with $j \neq i$ is bigger than the amount of steps it takes it to move from v_0 to s_j . Meaning that, if the surfer wins it wins the first time it moves out of one of these three vertices.

First, let us assume that the surfer goes from v_0 to s_i through A^i ($i \in \{0, 1, 2\}$). Clearly, at each step before reaching s_i , all neighbours of the current position of the surfer are marked. Now, when the surfer arrives at s_i , there are at least $2(\alpha + 1) + Z$ neighbours of s_i that are already marked. To show that the observer wins, it is sufficient to note that $|N(s_i)| - (2(\alpha + 1) + Z) = 4x - 3 - 2\alpha - 2 - \lfloor (3x - \alpha - \beta - 2\gamma - 12)/3 \rfloor \leq 3x - 2\alpha - 5 + (\alpha + \beta + 2\gamma + 12)/3 = 3x - 1 + (\beta + 2\gamma - 5\alpha)/3 \leq 3x + 2$ because $2\alpha > \beta + 2\gamma + 1$.

Second, let us assume that the surfer goes from v_0 to s_i through B^i, w_i and $C^{i,i}$ ($i \in \{0, 1, 2\}$). When the surfer arrives at s_i , there are at least $\beta + 1 + 2\gamma + Z$ neighbours of s_i that are already marked. To show that the observer wins, it is sufficient to note that $|N(s_i)| - (\beta + 1 + 2\gamma + Z) = 4x - 3 - \beta - 1 - 2\gamma - \lfloor (3x - \alpha - \beta - 2\gamma - 12)/3 \rfloor \leq 3x - \beta - 4 - 2\gamma + (\alpha + \beta + 2\gamma + 12)/3 \leq 3x + (\alpha - 2\beta - 4\gamma)/3 \leq 3x + 2$ because $\alpha < \beta + \gamma + 1$.

Finally, let us assume that the surfer goes from s_i (all neighbours of which are already marked) to $s_{i+1 \bmod 3}$ through $C^{i,i}, w_i$ and $C^{i,i+1 \bmod 3}$ ($i \in \{0, 1, 2\}$). When the surfer arrives at $s_{i+1 \bmod 3}$, there are at least $4\gamma + 2 + Z$ neighbours of $s_{i+1 \bmod 3}$ that are already marked. To show that the observer wins, it is sufficient to note that $|N(s_{i+1 \bmod 3})| - (4\gamma + 2 + Z) = 4x - 3 - 4\gamma - 2 - \lfloor (3x - \alpha - \beta - 2\gamma - 12)/3 \rfloor \leq 3x - 5\gamma - 4 + (\alpha + \beta + 2\gamma + 12)/3 \leq 3x - 1 + (\alpha + \beta - 10\gamma)/3 \leq 3x + 2$ because $\beta < \alpha < 2\gamma + 1$. \square

2.3. Relationship with another graph parameter

It is shown that, for any graph G and $v_0 \in V(G)$, $\max_{S \subseteq V(G)} \lceil \frac{|N[S]| - 1}{|S|} \rceil \leq \text{sn}(G, v_0) \leq \text{csn}(G, v_0)$ where the maximum is taken over every subset $S \subseteq V(G)$ inducing a connected subgraph with $v_0 \in S$ [4]. Moreover, both previous inequalities turn into an equality in case of trees. The authors of [4] ask whether the first inequality may be strict.

First, let us notice that such an equality is unlikely to hold since it would imply that the problem of computing the surveillance number of a graph is in co-NP while this problem is known to be PSPACE-complete in DAGs. We actually show that there are graphs where the inequality is strict.

Let us build a graph as follows. Starting from the vertex set $V = \{a, b, c, ab, ac, bc, s\}$ and $E = \{(s, a), (s, b), (s, c), (a, ab), (a, ac), (b, ab), (b, bc), (c, ac),$

$(c, bc)\}$. Then, we add $\frac{11k-21-2x}{6}$ leaves to each vertex ab, ac and bc , moreover, add 3 leaves to each vertex a, b and c , and, finally, add x leaves to s . A scheme of this family can be found in Figure 2.




Figure 2: Scheme of the graph family described in the proof of Theorem 8.

We moreover assume that $k - 5 \equiv 0 \pmod{2}$, $k - x - 3 \equiv 0 \pmod{3}$, $11k - 21 - 2x \equiv 0 \pmod{6}$, $x \leq k - 36$ and $k \geq 34$. For instance, $k = 105$ and $x = 42$ are possible.

Let \mathcal{G} be the graph obtained by the above construction and where parameters satisfy the above constraints.

Theorem 8. $\text{sn}(\mathcal{G}, s) = k$ and $\max_{S \subseteq V(\mathcal{G})} \lceil \frac{|N[S]|-1}{|S|} \rceil < k$.

Proof. Throughout this proof, let $M \subseteq V$ denote the set of (currently) marked vertices in \mathcal{G} .

We show a strategy for the surfer that wins against an observer that can mark at most $k - 1$ vertices per turn. Let $S_a = (N[a] \cup N[ab] \cup N[ac]) \setminus \{s, a, b, c\}$, $S_b = (N[b] \cup N[ab] \cup N[bc]) \setminus \{s, a, b, c\}$, $S_c = (N[c] \cup N[bc] \cup N[ac]) \setminus \{s, a, b, c\}$.

In the first step and after the observer has used its marks, the surfer chooses to move to i where $i = \arg \min_{i=\{a,b,c\}} |S_i \cap M|$. Since the observer must mark the vertices in $N(s)$ (including a, b, c) we have that $|S_i \cap M| \leq \frac{2}{3}(k - 1 - x - 3)$. Without loss of generality assume that $i = a$. In the second step, all neighbours of a must have been marked, otherwise the surfer wins by moving to an unmarked leaf of a . Let $S_{ab} = N[ab] \setminus \{a, b, ab\}$ and

$S_{ac} = N[ac] \setminus \{a, c, ac\}$, therefore, after all marks are spent in the second step, $\min_{j=\{ab, ac\}} |S_j \cap M| \leq \frac{k-1-5+\frac{2}{3}(k-1-x-3)}{2}$. The surfer then chooses to move to $\arg \min_{i=\{ab, ac\}} |S_i \cap M|$, w.l.o.g. assume that it is the vertex ab . In the third step, the observer might use all its available marks onto the leaves of ab , hence, after spending all the marks, $|S_{ab} \cap M| \leq k-1 + \frac{k-1-5+\frac{2}{3}(k-1-x-3)}{2} = \frac{11k-32-2x}{6}$ which is less than $|S_{ab}|$, hence there is an unmarked leaf of ab that the surfer can reach.

We consider now a winning strategy for the observer that marks k vertices per step. At the first step, the observer marks all vertices in $N[s]$, with the remaining marks, $k-x-3$, being spread evenly among vertices in the sets $N[ab] \setminus \{a, b, ab\}$, $N[ac] \setminus \{a, c, ac\}$ and $N[bc] \setminus \{b, c, bc\}$. Hence, there are at least $\lfloor \frac{k-x-3}{3} \rfloor = \frac{k-x-3}{3}$ vertices marked in each of those sets. Without loss of generality assume that the surfer moves towards a . Then, the observer marks the vertices in $N(a)$ and, with the remaining marks, proceeds to distribute them evenly among the vertices of the sets $N(ab)$ and $N(ac)$. When the surfer is about to move there are at least $\lfloor \frac{k-5}{2} \rfloor + \frac{k-x-3}{3} = \frac{k-5}{2} + \frac{k-x-3}{3}$ vertices in $(N(ab) \setminus \{a, b\}) \cap M$ and in $(N(ac) \setminus \{a, c\}) \cap M$. Without loss of generality assume that the surfer moves towards ab . Then the observer uses all its available marks on the unmarked vertices in $N(ab) \setminus \{a, b\}$. Therefore, after all marks are spent, there are $k + \frac{k-5}{2} + \frac{x-3}{3}$ marked vertices in $N(ab) \setminus \{a, b\}$. It remains to show that $k + \frac{k-5}{2} + \frac{x-3}{3} \geq \frac{11k-21-2x}{6}$.

$$\begin{aligned} k + \frac{k-5}{2} + \frac{x-3}{3} &\geq \frac{6k}{6} + \frac{3k-15}{6} + \frac{2x-6}{6} - 2 = \frac{9k-21+2x}{6} - 2 \\ \frac{9k-21+2x}{6} - 2 &= \frac{9k-33+4x-2x}{6} \geq \frac{11k-21-2x}{6} \end{aligned}$$

Now we show that for all connected sets S such that $s \in S$ we have that $\lceil \frac{|N[S]|-1}{|S|} \rceil < k$.

Claim 9. *For all connected sets S such that $s \in S$, then $\lceil \frac{|N[S]|-1}{|S|} \rceil \leq k-1$.*

Proof. First we prove that if S contains a vertex $v \in V$ with degree 1, then $\lceil \frac{|N[S]|-1}{|S|} \rceil \leq \lceil \frac{|N[S \setminus \{v\}]|-1}{|S \setminus \{v\}|} \rceil$. Since S contains s and induces a connected subgraph, then $N(v) \subset S$ because $|N(v)| = 1$. Thus $N[S \setminus \{v\}]$ contains v and so $N[S \setminus \{v\}] = N[S]$.

In the rest of the proof, we consider sets S that do not contain a node with degree 1. Let $L_{ab} = N(ab) \setminus \{a, b\}$, $L_{ac} = N(ac) \setminus \{a, c\}$, and $L_{bc} = N(bc) \setminus \{b, c\}$. By the previous assumption, if a node $v \in L_{ab}$ is such that

$v \in N[S]$, then all nodes in L_{ab} are in $N[S]$. By symmetry, we have the similar property for L_{ac} and L_{bc} . Note that $(N(s) \setminus \{a, b, c\}) \subset N[S]$ because $s \in S$ by definition.

We have four different cases:

- Consider S such that $N[S] \cap (L_{ab} \cup L_{ac} \cup L_{bc}) = \emptyset$. We get that $|S| \geq 1$ and $N[S] \leq x+16$. Thus $\lceil \frac{|N[S]|-1}{|S|} \rceil \leq x+15 \leq k-1$ because $x \leq k-36$.
- Consider S such that $N[S] \cap (L_{ac} \cup L_{bc}) = \emptyset$ and $L_{ab} \subset N[S]$. We get that $|S| \geq 3$ and $N[S] \leq x+16 + \frac{11k-21-2x}{6}$. Thus $\lceil \frac{|N[S]|-1}{|S|} \rceil \leq \lceil \frac{11k+4x+69}{18} \rceil \leq k-1$ because $x \leq k-36$ and $k \geq 34$. The case $N[S] \cap (L_{ab} \cup L_{bc}) = \emptyset$ and $L_{ac} \subset N[S]$ is similar and the case $N[S] \cap (L_{ab} \cup L_{ac}) = \emptyset$ and $L_{bc} \subset N[S]$ is also similar.
- Consider S such that $N[S] \cap L_{bc} = \emptyset$ and $L_{ab} \cup L_{ac} \subset N[S]$. We get that $|S| \geq 4$ and $N[S] \leq x+16 + \frac{11k-21-2x}{3}$. Thus $\lceil \frac{|N[S]|-1}{|S|} \rceil \leq \lceil \frac{11k+x+24}{12} \rceil \leq k-1$ because $x \leq k-36$ and $k \geq 34$. The case $N[S] \cap L_{ac} = \emptyset$ and $L_{ab} \cup L_{bc} \subset N[S]$ is similar and the case $N[S] \cap L_{ab} = \emptyset$ and $L_{ac} \cup L_{bc} \subset N[S]$ is also similar.
- Consider S such that $L_{ab} \cup L_{ac} \cup L_{bc} \subset N[S]$. We get that $|S| \geq 6$ and $N[S] \leq x+16 + \frac{11k-21-2x}{2}$. Thus $\lceil \frac{|N[S]|-1}{|S|} \rceil \leq \lceil \frac{11k+9}{12} \rceil \leq k-1$ because $k \geq 34$.

□

This concludes the proof of Claim 9 and therefore, the proof of Theorem 8 because we have proved that $\text{sn}(\mathcal{G}, s) = k$. □

3. Online Surveillance Number

In this section, we study the online variant of the surveillance game motivated by the web-page prefetching problem where the observer (the web-browser) discovers new nodes through hyperlinks in already marked nodes. In this variant, the observer does not know *a priori* the graph in which it is playing. That is, initially, the observer only knows v_0 , its degree and the identifiers of its neighbours. Then, when a new node is marked, the observer discovers all its neighbours that are not yet marked. Note that the degree of a node is not known before it is marked.

Another property of an online strategy that must be defined concerns the moment when the observer discovers the unmarked neighbours of a node

that it has decided to mark. There are two natural models. Assume that the set M of nodes have been marked and this is the turn of the observer, and let $N(M)$ be the set of nodes with a neighbour in M . Either it first chooses the k nodes that will be marked among the set $N(M) \setminus M$ of the unmarked neighbours of the nodes that were already marked and then the observer marks each of these k nodes and discovers their unknown neighbours simultaneously. Or, the observer first chooses one node x in $N(M) \setminus M$, marks it and discovers its unmarked neighbours, then it chooses a new node to be marked in $N(M \cup \{x\}) \setminus (M \cup \{x\})$ and so on until the observer finishes its turn after marking k nodes. Note that the second model is less restricted since the observer has more power. However, we show that, even in this model, the basic strategy is the best one with respect to the competitive ratio.

Formal definition of online strategy. Now we are ready to formally define an online strategy. Let $k \geq 1$, let $G = (V, E)$ be a graph, $v_0 \in V$, and let \mathcal{G} be the set of subgraphs of G .

Let $M \subseteq V$ be a subset of nodes inducing a connected subgraph containing v_0 in G . Let $G_M \in \mathcal{G}$ be the subgraph of G known by the observer when M is the set of marked nodes. That is, $G_M = (M \cup N(M), E_M)$ where $E_M = \{(u, v) \in E \mid u \in M\}$. For any $u, v \in N(M) \setminus M$, let us set $u \sim_M v$ if and only if $N(u) \cap M = N(v) \cap M$. Let χ_M be the set of equivalent classes, called *modules*, of $N(M) \setminus M$ with respect to \sim_M . The intuition is that two nodes in the same module of χ_M are known by the observer but cannot be distinguished. For instance, $\chi_{\{v_0\}} = \{N(v_0)\}$.

A *k-online strategy for the observer starting from v_0* is a function $\sigma : \mathcal{G} \times V \times 2^V \times \{1, \dots, k\} \rightarrow 2^V$ such that, for any subset $M \subseteq V$ of nodes inducing a connected subgraph containing v_0 in G , for any $v \in M$, and for any $1 \leq i \leq k$, then $\sigma(G_M, v, M, i) \in \chi_M$. This means that, if M is the set of nodes already marked and thus the observer only knows the subgraph G_M , if v is the position of the surfer and it remains $k - i + 1$ nodes to be marked by the observer before the surfer moves, then the observer will mark one node in $\sigma(G_M, v, M, i)$.

More precisely, we say that the observer *follows* the *k-online strategy* σ if the game proceeds as follows. Let $M = M^0$ be the set of marked nodes just after the surfer has moved to $v \in M$. Initially, $M^0 = \{v_0\}$ and $v = v_0$. Then, the strategy proceeds sequentially in k steps for $i = 1$ to k . First, the observer marks an arbitrary node $x_1 \in \sigma(G_{M^0}, v, M^0, 1)$. Let $M^1 = M^0 \cup \{x_1\}$. Sequentially, after having marked $1 < i < k$ nodes at this turn, the observer marks one arbitrary node $x_{i+1} \in \sigma(G_{M^i}, v, M^i, i+1)$.




Figure 3: Tree T_k described in the proof of Theorem 10.

and $M^{i+1} = M^i \cup \{x_{i+1}\}$. When the observer has marked k nodes, that is after choosing $x_k \in \sigma(G_{M^{k-1}}, v, M^{k-1}, k)$, it is the turn of the surfer, when it may move to a node adjacent to its current position and then a new turn for the observer starts. Note that because we are interested in the worst case for the observer, each marked node $x_i \in \sigma(G_{M^{i-1}}, v, M^{i-1}, i)$ is chosen by an adversary.

The *online surveillance number* of a graph G with initial node v_0 , denoted by $\text{osn}(G, v_0)$, is the smallest k such that there exists a winning k -online strategy starting from v_0 . In other words, there is a winning k -online strategy σ starting from v_0 such that an observer following σ wins whatever be the trajectory of the surfer and the choices done by the adversary at each step. Note that, since we consider the worst scenario for the observer, we may assume that the surfer has full knowledge of G .

Theorem 10. *There exists an infinite family of rooted trees such that, for any T with root $v_0 \in V(T)$ in this family, $\text{sn}(T, v_0) = 2$ and $\text{osn}(T, v_0) = \Omega(\Delta)$ where Δ is the maximum degree of T .*

Proof. We first define the family $(T_k)_{k \geq 1}$ of rooted trees as follows.

Let $k \geq 4$ be a power of two and let $i = 2^k$ and $d = \frac{2^k}{k}$.

Let us consider a path $P = (v_0, v_1, \dots, v_{i-1})$ with i nodes. Let B be a complete binary tree of height $h = 3k+1$ and rooted at some vertex v_i , i.e., B has $2^{h+1} - 1$ vertices. Let w_0 be any leaf of B . Finally, let $Q = (w_1, \dots, w_k)$ be a path on k nodes. Note that, P, B and Q depend on k .

The tree T_k is obtained from P, B and Q by adding an edge between v_{i-1} and v_i , an edge between w_0 and w_1 . Finally, for any $1 \leq j \leq k$, let us add an independent set, S_j , with d vertices and an edge between each vertex of S_j and w_j (i.e., each node in S_j is a leaf). T_k is then rooted in v_0 .

Let Q^+ denote the union of vertices of Q and $\bigcup_{j=1}^k S_j$. The maximum degree Δ of T_k is reached by any node w_j , $1 \leq j < k$, and $\Delta = d+2 = \frac{2^k}{k} + 2$.

Clearly, $\text{sn}(T_k, v_0) > 1$. We show that $\text{sn}(T_k, v_0) = 2$.

Consider the following (offline) strategy for the observer. At each turn $j \leq i$, the observer marks the vertex v_j and one unmarked vertex of Q^+ that is closest to the surfer. For each turn $j > i$ and while the surfer does not occupy a node in $Q^+ \cup \{w_0\}$, the observer marks the neighbours of the current position of the surfer if they are not already marked. Finally, if the surfer occupies a node in $Q^+ \cup \{w_0\}$, the observer marks two unmarked nodes of Q^+ that are closest to the surfer. It is easy to see, by induction on the number of steps that, each time that the surfer arrives at a new node, this node is marked and has at most 2 unmarked neighbours. Hence, $\text{sn}(T_k, v_0) = 2$.

Now it remains to show that $\text{osn}(T_k, v_0) = \Omega(\Delta)$. Let γ be any online strategy for T_k and marking at most $\frac{d}{4} = \frac{2^{k-2}}{k}$ nodes per turn. We show that γ fails.

For this purpose, we model the fact that the observer does not know the graph by “building” the tree during the game. More precisely, each time the observer marks a node v , then the adversary may add new nodes adjacent to v or decide that v is a leaf. Of course, the adversary must satisfy the constraint that eventually the graph is T_k . Initially, the observer only knows v_0 that has one neighbour v_1 . Now, for any $1 \leq j < i$, when the observer marks the node v_j of P , then the adversary “adds” a new node v_{j+1} adjacent to v_j , i.e., the observer discovers its single unmarked neighbour v_{j+1} . Now, let v be any node of B . Recall that h is the height of B . When the observer marks v , there are three cases to be considered: if v is at distance at most $h - 1$ from v_i , then the adversary adds two new nodes adjacent to v ; if v is at distance h from v_i and not all nodes of B have been marked then the adversary decides that v is a leaf; finally, if all nodes of B have been marked (v is the last marked node of B , i.e., B is a complete binary tree of height h), the adversary decides that $v = w_0$ and adds one new neighbour w_1 adjacent

to it. Note that we can ensure that the last node of B to be marked is at distance h of v_i by connectivity of any online strategy.

Now, let us consider the following execution of the game. During the first i steps, the surfer goes from v_0 to v_i . Just after the surfer arrives in v_i , the observer has marked at most $(di)/4$ nodes and all nodes of $P \cup \{v_i\}$ must be marked since otherwise the surfer would have won. Therefore, at most $i(d/4 - 1) + 1 = 2^{2k-2}/k - 2^k + 1$ nodes of B are marked when it is the turn of the surfer at v_i . Since B has $2^{h+1} - 1 = 2^{3k+2} - 1$ nodes, at least one node of B is not marked.

From v_i , the surfer always goes toward w_0 . Note that the observer may guess this strategy but it does not know where is w_0 while all nodes of B have not been marked.

Then let $0 \leq t \leq h$ and let $v'_t \in V(B)$ be the position of the surfer at step $i + t$ and B^t the subtree of B rooted at v'_t . Note that, at step i , $v'_0 = v_i$ and $B^0 = B$. Let B_l^t and B_r^t be the subtrees of B rooted at the children of v'_t . W.l.o.g., let us assume that the number of marked nodes in B_l^t is at most the number of marked nodes in B_r^t , when it is the turn of the surfer standing at v'_t . Then, the surfer moves to the root of B_l^t . That is, v'_{t+1} is the child of v_t whose subtree contains the minimum number of marked nodes.

Let m_t be the number of marks in the subtree of B rooted at v'_t when it is the turn of the surfer at v'_t . Since, at the beginning of step i there are at most $2^{2k-2}/k - 2^k + 1$ nodes of B that are marked and $k \geq 4$, then $m_0 \leq 2^{2k-2}/k - 2^k + 1 \leq 2^{2k-2}/k$. Note that, for any $t > 0$, $m_t \leq (m_{t-1} - 1 + \frac{d}{4})/2 \leq (m_{t-1} + \frac{d}{4})/2$. Simply expanding this expression we get that, for any $t > 0$,

$$m_t \leq \frac{m_0}{2^t} + \frac{2^k}{k} \sum_{j=3}^{t+2} 2^{-j} \leq \frac{2^{2k-(t+2)}}{k} + \frac{2^k}{k} \sum_{j=3}^{t+2} 2^{-j}.$$

Therefore, for any $t \geq 2k$:

$$m_t \leq \frac{1}{k} + \frac{2^k}{k} \sum_{j=3}^{t+2} 2^{-j} \leq \frac{2^k + 1}{k}.$$

In particular, at step $i + 2k$ (when it is the turn of the surfer), the surfer is at v'_{2k} which is at distance $k + 1$ from w_0 . Hence, $|B^{2k}| \geq 2^{k+1} - 1$ and at most $\frac{2^k + 1}{k} < 2^{k+1} - 1$ of its nodes are marked. Hence, neither w_0 nor nodes in Q^+ are marked.

From this step, the surfer directly goes to w_k unless it meets an unmarked node, in which case, it goes to it and wins. When the surfer is at w_k and it is its turn, the observer may have marked at most $(2k+2)\frac{d}{4} \leq \frac{kd}{2} + \frac{d}{2} \leq 2^{k-1} + \frac{2^{k-1}}{k}$ nodes in Q^+ . Since $|Q^+| = (d+1)k = 2^k + k$ and $k \geq 4$, at least one neighbour of w_k is not marked yet and the surfer wins. \square

Theorem 10 implies that, for any online strategy \mathcal{S} , $\rho(\mathcal{S}) = \Omega(\Delta)$. Recall that the basic strategy \mathcal{B} , that marks all unmarked neighbours of the surfer at each step, is an online strategy. \mathcal{B} has trivially competitive ratio $\rho(\mathcal{B}) = O(\Delta)$. Hence, no online winning strategy has better competitive ratio than the basic strategy up to a constant factor. In other words:

Corollary 1. The best competitive ratio of online winning strategies is $\Theta(\Delta)$, with Δ the maximum degree.

As mentioned in the introduction, any online strategy is connected and therefore, for any graph G and $v_0 \in V(G)$, $\text{csn}(G, v_0) \leq \text{osn}(G, v_0)$. Moreover, we recall that, for any tree T and for any $v_0 \in V(T)$, $\text{csn}(T, v_0) = \text{sn}(T, v_0)$ [4]. Hence, the previous theorem shows that there might be an arbitrary gap between $\text{csn}(G, v_0)$ and $\text{osn}(G, v_0)$.

4. Bounded Degree Hardness

The question of the complexity of computing the surveillance number in the class of bounded degree graphs was left open in [4]. Notice that computing $\text{sn}(G, v_0)$ is trivial for graphs of maximum degree 3 [4].

In this section, we show that the problem is difficult in the class of DAGs with maximum degree 6. More precisely, we prove the following theorem.

Theorem 11. *Deciding whether $\text{sn}(G, v_0) \leq 2$, for a directed acyclic graph G of maximum degree 6 and a starting vertex $v_0 \in V(G)$, is NP-hard.*

A graph is called *cubic*, if every node has degree exactly 3. To prove NP-hardness, a reduction from a special case of the well-known *Vertex Cover* problem is employed. Given a cubic graph G and a constant k , the *Cubic Vertex Cover* problem consists in deciding whether there exists a set $C \subseteq V(G)$, $|C| \leq k$ and such that for any $\{v_i, v_j\} \in E(G)$, $|\{v_i, v_j\} \cap C| \geq 1$. NP-hardness for the above problem was proved in [6]. From now on, we shall refer to the problem shortly as VC-3.

Let (G, k) be any instance of VC-3 and set $V(G) = \{x_1, \dots, x_n\}$ and $E(G) = \{e_1, \dots, e_m\}$. We build an instance (D, v_0) of the surveillance game

problem from the instance (G, k) as follows. We start with a directed path $(v_0, v_1, \dots, v_{k+m-2}, c_1, c_2, \dots, c_m)$. For any $1 \leq i \leq m$, let us add three new nodes $c_i^{left}, c_i^{right}, c_i^{mid}$ and the following arcs $(c_i, c_i^{left}), (c_i, c_i^{right}), (c_i^{left}, c_i^{mid}), (c_i^{right}, c_i^{mid})$. Finally, let us add n nodes u_1, \dots, u_n to D , and, for any edge $e_i = \{u_j, u_\ell\}$ of G , let us add the corresponding arcs $(c_i^{left}, u_j), (c_i^{left}, u_\ell), (c_i^{right}, u_j)$ and (c_i^{right}, u_ℓ) in D . The digraph D is depicted in Figure 4. Notice that, since G is cubic, the sum of in-degree and out-degree of each node of D is at most 6. In particular, u_i has in-degree 6 for any $i \leq n$. Furthermore, $|V(D)| = n + 5m + k - 1$ and $|E(D)| = 10m + k - 2$.




Figure 4: Scheme representing the digraph D constructed from G .

Lemma 12. *If G has a vertex cover of size at most k , then $\text{sn}(D, v_0) \leq 2$.*

Proof. If there exists a vertex cover of size at most k for G , then we show how the observer can win in D by using 2 marks per round. Let $X \subseteq V(G)$ be a vertex cover of size at most k and let $X' = \{u_i \mid x_i \in X, 1 \leq i \leq n\} \subseteq V(D)$. That is, X' is the set of nodes of D that *corresponds* to X . Initially, v_0 is marked. When the surfer moves to any node $v_i \in \{v_0, v_1, \dots, v_{m-1}\}$, the observer's strategy is to mark v_{i+1} and c_{i+1} . When the surfer moves to any node $v_i \in \{v_m, \dots, v_{m+k-3}\}$, the observer's strategy is to mark v_{i+1} and an unmarked node in X' . When the surfer moves to node v_{m+k-2} , the observer's strategy is to mark the last two unmarked nodes in X' . At this step, the set of marked nodes is $\{v_0, \dots, v_{m+k-2}\} \cup \{c_1, \dots, c_m\} \cup X'$.

Now, while the surfer moves to c_i , $i \leq m$, the observer marks both unmarked neighbours c_i^{left} and c_i^{right} of c_i . Then, at some step, the surfer goes from c_i to $w \in \{c_i^{left}, c_i^{right}\}$. Let $e_i = \{x_j, x_\ell\}$ be the edge of G corresponding to c_i . Since X is a vertex cover of G , either x_j or x_ℓ , say x_j , belongs to it. Hence, $u_j \in X'$ has been already marked. Thus w has at most two unmarked neighbours: c_i^{mid} and x_ℓ . The observer marks them and the surfer cannot access any unmarked node.

Hence, $sn(D, v_0) \leq 2$. \square

Lemma 13. *If G does not admit a vertex cover of size k , then $sn(D, v_0) > 2$.*

Proof. If G does not admit a vertex cover of size k , we provide a winning strategy for the surfer against an observer that marks 2 nodes per step in D .

Let \mathcal{S} be any strategy for the observer marking at most 2 nodes per steps. We will show that this strategy is not winning. First, we show that we may assume that \mathcal{S} marks a node c_i^{mid} ($i \leq m$) only if the surfer occupies an in-neighbour of it. Indeed, let us assume that there is a step such that \mathcal{S} marks c_i^{mid} while the surfer is neither at c_i^{left} nor at c_i^{right} . Let $e_i = \{x_j, x_\ell\}$ be the edge of G corresponding to c_i . We modify \mathcal{S} such that, instead of marking c_i^{mid} , it marks one unmarked node in $\{u_j, u_\ell\}$. Clearly, if \mathcal{S} was a winning strategy, then the modified strategy is still winning since, when the surfer arrives in c_i^{left} or c_i^{right} , at most 2 out-neighbours are not marked. Hence, we may assume that \mathcal{S} satisfies the desired property.

Now, let us consider the following strategy for the surfer. First, the surfer follows the path from v_0 to c_1 . Just before the surfer leaves c_1 (i.e., after the observer has marked at most $2m + 2k$ nodes), all nodes in $\{v_0, \dots, v_{k+m-2}, c_1\}$ have been marked by \mathcal{S} since otherwise the surfer would have already won. Moreover, c_1^{left} , c_1^{right} and c_2 must also be marked since otherwise the surfer would win during its next move. Hence, at this step, at most $m+k-1$ nodes are marked in $V(D) \setminus \{v_0, \dots, v_{k+m-2}, c_1, c_2, c_1^{left}, c_1^{right}\}$.

Let $1 \leq i < m$. Assume that the surfer has followed the path from c_1 to c_i and, when it is about to leave c_i : all nodes in $M_i = \{v_0, \dots, v_{k+m-2}, c_1, \dots, c_{i+1}, c_1^{left}, \dots, c_i^{left}, c_1^{right}, \dots, c_i^{right}\}$ have been marked, and at most $m+k-i$ nodes are marked in $V(D) \setminus M_i$. Note that, in the above paragraph, we reached this configuration for $i = 1$.

Note that, by the property of \mathcal{S} , c_i^{mid} is not marked yet. Let u_j and u_ℓ be the two out-neighbours, distinct from c_i^{mid} , of c_i^{left} . If both u_j and u_ℓ are not marked yet, then the surfer goes to c_i^{left} that has then 3 unmarked out-neighbours. Therefore, the surfer will win during its next move. Hence,

we may assume that either u_j or u_ℓ is marked. In such a case, the surfer goes to c_{i+1} . Now, the observer marks at most 2 nodes. Note that, c_{i+1}^{left} , c_{i+1}^{right} and (if $i + 1 < m$) c_{i+2} must have been marked since otherwise the surfer would win during its next move. Therefore, the surfer is about to leave c_{i+1} and all nodes in M_{i+1} have been marked, and at most $m + k - (i + 1)$ nodes are marked in $V(D) \setminus M_{i+1}$.

Therefore, either the surfer wins at some step or it eventually reached c_m . Let X' be the set of marked nodes in $\{u_1, \dots, u_n\}$ when the surfer is about to leave c_m . By the above reasoning, $|X'| \leq k$. Moreover, for any $i < m$, there is at least one node in X' that is adjacent to c_i^{left} . Hence, by definition of D , the set $X = \{x_i \mid u_i \in X'\}$ covers all edges in $\{e_1, \dots, e_{m-1}\} \subset E(G)$. Since no vertex cover of G has size at most k , this implies that $e_m = \{x_j, x_\ell\}$ is not covered by X . Therefore, none of the two neighbours u_j and u_ℓ of c_m^{left} belongs to X' . That is, when the surfer is about to leave c_m , neither u_j nor u_ℓ are marked. Moreover, by the property of \mathcal{S} , c_m^{mid} is not marked yet. To conclude, the surfer goes to c_m^{left} that has 3 unmarked out-neighbours. The surfer will win during its next move. \square

The proof of Theorem 11 follows directly from Lemma 12 and Lemma 13.

5. Conclusion

Despite our results, the main question remains open. Can the difference or the ratio between the connected surveillance number of a graph and its surveillance number be bounded by some constant?

References

- [1] Albrecht, D., Zukerman, I., Nicholson, A., 1999. Pre-sending documents on the www: a comparative study. In: 16th Int. Joint Conf. on Artificial Intelligence. pp. 1274–1279.
- [2] Aumann, Y., Etzioni, O., Feldman, R., Perkowitz, M., 1998. Predicting event sequences: Data mining for prefetching web-pages.
- [3] Fan, L., Cao, P., Lin, W., Jacobson, Q., 1999. Web prefetching between low-bandwidth clients and proxies: potential and performance. In: ACM SIGMETRICS Int. Conf. on Measurement and Modeling of Computer Systems. pp. 178–187.

- [4] Fomin, F. V., Giroire, F., Jean-Marie, A., Mazauric, D., Nisse, N., 2012. To satisfy impatient web surfers is hard. In: FUN. pp. 166–176.
- [5] Fomin, F. V., Giroire, F., Jean-Marie, A., Mazauric, D., Nisse, N., 2014. To satisfy impatient web surfers is hard. Theoretical Computer Science 526 (0), 1 – 17.
- [6] Garey, M., Johnson, D., Stockmeyer, L., 1976. Some simplified np-complete graph problems. Theoretical Computer Science 1 (3), 237 – 267.
- [7] Kroeger, T. M., Long, D. D. E., Mogul, J. C., 1997. Exploring the bounds of web latency reduction from caching and prefetching. In: USENIX Symposium on Internet Technologies and Systems. pp. 2–2.
- [8] Mogul, J. C., 1996. Hinted caching in the web. In: 7th workshop on ACM SIGOPS European workshop: Systems support for worldwide applications. pp. 103–108.
- [9] Padmanabhan, V. N., Mogul, J. C., 1996. Using predictive prefetching to improve world wide web latency. SIGCOMM Comput. Commun. Rev. 26 (3), 22–36.
- [10] Wang, Z., Lin, F. X., Zhong, L., Chishtie, M., 2012. How far can client-only solutions go for mobile browser speed? In: 21st Int. Conf. on World Wide Web. pp. 31–40.