Randomized Pursuit-Evasion in Graphs

Micah Adler[1], Harald Räcke[2 ⋆], Naveen Sivadasan[3 ⋆], Christian Sohler[2 ⋆], and Berthold Vöcking[3 ⋆]

[1] Department of Computer Science University of Massachusetts, Amherst, micah@cs.umass.edu
[2] Heinz Nixdorf Institute and Department of Mathematics and Computer Science, Paderborn University, Germany, {harry,csohler}@upb.de
[3] Max-Planck-Institut für Informatik, Saarbrücken, Germany, {ns,voecking}@mpi-sb.mpg.de

Abstract. We analyze a randomized pursuit-evasion game on graphs. This game is played by two players, a *hunter* and a *rabbit*. Let G be any connected, undirected graph with n nodes. The game is played in rounds and in each round both the hunter and the rabbit are located at a node of the graph. Between rounds both the hunter and the rabbit can stay at the current node or move to another node. The hunter is assumed to be *restricted* to the graph G: in every round, the hunter can move using at most one edge. For the rabbit we investigate two models: in one model the rabbit is restricted to the same graph as the hunter, and in the other model the rabbit is *unrestricted*, i.e., it can jump to an arbitrary node in every round.

We say that the rabbit is *caught* as soon as hunter and rabbit are located at the same node in a round. The goal of the hunter is to catch the rabbit in as few rounds as possible, whereas the rabbit aims to maximize the number of rounds until it is caught. Given a randomized hunter strategy for G, the *escape length* for that strategy is the worst case expected number of rounds it takes the hunter to catch the rabbit, where the worst case is with regards to all (possibly randomized) rabbit strategies. Our main result is a hunter strategy for general graphs with an escape length of only $\mathcal{O}(n \log(\operatorname{diam}(G)))$ against restricted as well as unrestricted rabbits. This bound is close to optimal since $\Omega(n)$ is a trivial lower bound on the escape length in both models. Furthermore, we prove that our upper bound is optimal up to constant factors against unrestricted rabbits.

1 Introduction

In this paper we introduce a pursuit evasion game called the *Hunter vs. Rabbit* game. In this round-based game, a pursuer (the *hunter*) tries to catch an evader (the *rabbit*) while they both travel from vertex to vertex of a connected, undirected graph G. The hunter catches the rabbit when in some round the hunter and the rabbit are both located on the same vertex of the graph. We assume that both players know the graph in advance but they cannot see each other until the rabbit gets caught. Both players may use a randomized (also called *mixed*) strategy, where each player has a secure source of randomness

⋆ Partially supported by the IST Programme of the EU under contract number IST-1999-14186 (ALCOM-FT)

which cannot be observed by the other player. In this setting we study upper bounds (i.e., good hunter strategies) as well as lower bounds (i.e., good rabbit strategies) on the expected number of rounds until the hunter catches the rabbit.

The problem we address is motivated by the question of how long it takes a single pursuer to find an evader on a given graph that, for example, corresponds to a computer network or to a map of a terrain in which the evader is hiding. A natural assumption is that both the pursuer and the evader have to follow the edges of the graph. In some cases however it might be that the evader has more advanced possibilities than the pursuer in the terrain where he is hiding. Therefore we additionally consider a stronger adversarial model in which the evader is allowed to jump arbitrarily between vertices of the graph. Such a jump between vertices corresponds to a short-cut between two places which is only known to the evader (like a rabbit using rabbit holes). Obviously, a strategy that is efficient against an evader that can jump is efficient as well against an evader who may only move along the edges of the graph.

One approach to use for a hunter strategy would be to perform a random walk on the graph G. Unfortunately, the hitting time of a random walk can be as large as $\Omega(n^3)$ with n denoting the number of nodes. Thus it would require at least $\Omega(n^3)$ rounds to find a rabbit even if the rabbit does not move at all. We show that one can do significantly better. In particular, we prove that for any graph G with n vertices there is a hunter strategy such that the expected number of rounds until a rabbit that is not necessarily restricted to the graph is caught is $\mathcal{O}(n \log n)$ rounds. Furthermore we show that this result cannot be improved in general as there is a graph with n nodes and an unrestricted rabbit strategy such that the expected number of rounds required to catch this rabbit is $\Omega(n \log n)$ for any hunter strategy.

1.1 Preliminaries

Definition of the game. In this section we introduce the basic notations and definitions used in the remainder of the paper. The Hunter vs. Rabbit game is a round-based game that is played on an undirected connected graph $G = (V, E)$ without self loops and multiple edges. In this game there are two players - the hunter and the rabbit - moving on the vertices of G. The hunter tries to catch the rabbit, i.e., he tries to move to the same vertex as the rabbit, and the rabbit tries not to be caught.

During the game both players cannot “see” each other, i.e., a player has no information about the movement decisions made by his opponent and thus does not know his position in the graph. The only interaction between both players occurs when the game ends because the hunter and the rabbit move to the same vertex in G and the rabbit is caught. Therefore the movement decisions of both players do not depend on each other. We want to find good strategies for both hunter and rabbit. We say that a hunter strategy has *expected (worst case) escape length* k, if for any rabbit strategy the expected number of rounds until the hunter catches the rabbit is k. Analogously a rabbit strategy is said to have *expected (worst case) escape length* k, if for any hunter strategy the expected number of rounds until the rabbit is caught is k. We can define a strategy for a player in the following way.

Definition 1. *A pure strategy for a player in the Hunter vs. Rabbit game on a graph $G = (V, E)$ is a sequence $\mathcal{S} = \mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2, \ldots$, where $S_t \in V$ denotes the position of* *the player in round $t \in \mathbb{N}_0$ of the game. A mixed strategy or strategy $\mathbb{S}$ for a player is a probability distribution over the set of pure strategies.*

Note that both players may use mixed strategies, i.e., we assume that they both have a source of random bits for randomizing their movements on the graph.

As mentioned in the previous section we assume that the hunter cannot change his position arbitrarily between two consecutive rounds but has to follow the edges of G. To model this we call a pure strategy $\mathcal{S}$ *restricted* (to G) if either $(\mathcal{S}_t, \mathcal{S}_{t+1}) \in E$ or $\mathcal{S}_t = \mathcal{S}_{t+1}$ holds for any $t \in \mathbb{N}_0$. A (mixed) strategy is called restricted if it is a probability distribution over the set of restricted pure strategies. For the analysis we will consider only restricted strategies for the hunter and both restricted as well as unrestricted strategies for the rabbit.

Notice that in our definition, the hunter may start his walk on the graph at an arbitrary vertex. However, we want to point out that defining a fixed starting position for the hunter would not asymptotically effect the results of the paper.

1.2 Previous Work

A first study of the Hunter vs. Rabbit game can be found in [1]. The presented hunter strategy is based on a random walk on the graph and it is shown that the hunter catches an unrestricted rabbit within $O(nm^2)$ rounds, where n and m denote the number of nodes and edges, respectively. In fact, the authors place some additional restrictions on the space requirements for the hunter strategy, which is an aspect that we do not consider in this paper.

In the area of mobile ad-hoc networks related models are used to design communication protocols (see e.g. [2, 3]). In this scenario, some mobile users (the "hunters") aid in transmitting messages to the receivers (the "rabbits"). The expected number of rounds needed to catch the rabbit in our model corresponds directly to the expected time needed to deliver a message. We improve the deliver time of known protocols, which are based on random walks.

Deterministic pursuit-evasion games in graphs are well-studied. In the early work by Parsons [14, 15] the graph was considered to be a system of tunnels in which the fugitive is hiding. Parsons introduced the concept of the *search number* of a graph which is, informally speaking, the minimum number of guards needed to capture a fugitive who can move with arbitrary speed. LaPaugh [9] showed that there is always a search strategy (a sequence of placing, removing, or moving a pebble along an edge) such that no edge that is cleared at a point of time can be recontaminated again, i.e., if the fugitive is known not to be in edge e then there is no chance for him to enter edge e again in the remainder of the search. Meggido et al. [11] proved that the computation of the search number of a graph is an NP-hard problem which implies its NP-completeness because of LaPaugh's result.

If an edge can be cleared without moving along it, but it suffices to 'look into' an edge from a vertex, then the minimum number of guards needed to catch the fugitive is called the *node search number* of a graph [8].

Pursuit evasion problems in the plane were introduced by Suzuki and Yamashita [16]. They gave necessary and sufficient condition for a simple polygon to be searchable by a single pursuer. Later Guibas et al. [6] presented a complete algorithm and showed that the problem of determining the minimal number of pursuers needed to clear a polygonal region with holes is NP-hard. Recently, Park et al. [13] gave 3 necessary and sufficient conditions for a polygon to be searchable and showed that there is an $\mathcal{O}(n^2)$ time algorithm for constructing a search path for an n-sided polygon.

Efrat et al. [4] gave a polynomial time algorithm for the problem of clearing a simple polygon with a chain of k pursuers when the first and last pursuer have to move on the boundary of the polygon.

1.3 New Results

We present a hunter strategy for general networks that improves significantly on the results obtained by using random walks. Let $G = (V, E)$ denote a connected graph with n vertices and diameter $\mathrm{diam}(G)$. Recall that $\Omega(n)$ is a trivial lower bound on the escape length against restricted as well as against unrestricted rabbit strategies on every graph with n vertices. Our hunter strategy achieves escape length close to this lower bound. In particular, we present a hunter strategy that has an expected escape length of only $\mathcal{O}(n \log(\mathrm{diam}(G)))$ against any unrestricted rabbit strategy. Clearly, an upper bound on the escape length against unrestricted rabbit strategies implies the same upper bound against restricted strategies.

Our general hunter strategy is based on a hunter strategy for cycles which is then simulated on general graphs. In fact, the most interesting and original parts of our analysis deal with hunter strategies for cycles. Observe that if hunter and rabbit are restricted to a cycle, then there is a simple, efficient hunter strategy with escape length $\mathcal{O}(n)$. (In every nth round, the hunter chooses a *direction* at random, either clockwise or counterclockwise, and then it follows the cycle in this direction for the next n rounds.) Against unrestricted rabbits, however, the problem of devising efficient hunter strategies becomes much more challenging. (For example, for the hunter strategy given above, there is a simple rabbit strategy that results in an escape length of $\Theta(n\sqrt{n})$.) For unrestricted rabbits on cycles of length n, we present a hunter strategy with escape length $\mathcal{O}(n \log n)$. Furthermore, we prove that this result is optimal by devising an unrestricted rabbit strategy with escape length $\Omega(n \log n)$ against any hunter strategy on the cycle.

Generalizing the lower bound for cycles, we can show that our general hunter strategy is optimal in the sense that for any positive integers n, d with $d < n$ there exists a graph G with n nodes and diameter d such that any hunter strategy on G has escape length $\Omega(n \cdot \log(\mathrm{diam}(G)))$. This gives rise to the question whether $n \cdot \log(\mathrm{diam}(G))$ is a universal lower bound on the escape length in any graph. We can answer this question negatively. In fact, in a full version of this paper, we present a hunter strategy with escape length $\mathcal{O}(n)$ for complete binary trees against unrestricted rabbits.

Finally, we investigate the Hunter vs. Rabbit game on strongly connected directed graphs. We show that there exists a directed graph for which every hunter needs $\Omega(n^2)$ rounds to catch a restricted rabbit. Furthermore, for every strongly connected directed graph, there is a hunter strategy with escape length $\mathcal{O}(n^2)$ against unrestricted rabbits. Due to space limitations, the analyses for directed graphs as well has to be moved to the full version.

1.4 Basic Concepts

The strategies will be analyzed in phases. A phase consists of m consecutive rounds, where m will be defined depending on the context. Suppose that we are given an m-round hunter strategy $\mathbb{H}$ and an m-round rabbit strategy $\mathbb{R}$ for a phase. We want to determine the probability that the rabbit is caught during the phase. Therefore we introduce the indicator random variables $hit(t), 0 \leq t < m$ for the event $\mathcal{H}_t = \mathcal{R}_t$ that the pure hunter strategy $\mathcal{H}$ and the pure rabbit strategy $\mathcal{R}$ chosen according to $\mathbb{H}$ and $\mathbb{R}$, respectively, meet in round t of the phase. Furthermore, we define indicator random variables $\mathit{fhit}(t), 0 \leq t < n$ describing first hits, i.e., $\mathit{fhit}(t) = 1$ iff $hit(t) = 1$ and $hit(t') = 0$ for every $t' \in \{0, \ldots, t-1\}$. Finally we define $hits = \sum_{t=0}^{m-1} hit(t)$.

The goal of our analysis is to derive upper and lower bounds for $\mathbf{Pr}[hits \geq 1]$, the probability that the rabbit is caught in the phase. To analyze the quality of an m-round hunter strategy we fix a pure rabbit strategy $\mathcal{R}$ and derive a lower bound on the probability $\mathbf{Pr}[hits \geq 1]$. Similarly to analyze the quality of an m-round rabbit strategy we fix a pure hunter strategy and derive an upper bound on $\mathbf{Pr}[hits \geq 1]$. Then we apply Yao's min-max principle [12] to derive the bounds for the mixed strategies.

The following two non-standard but fundamental propositions are important tools in the analysis of the upper and lower bounds.

Proposition 1. *Let $\mathbb{R}$ be an m-round rabbit strategy and let $\mathcal{H}$ be a pure m-round hunter strategy. Then*

$$\mathbf{Pr}\left[hits \geq 1\right] = \frac{\mathbf{E}\left[hits\right]}{\mathbf{E}\left[hits \mid hits \geq 1\right]} \quad .$$

Proposition 2. *Let $\mathbb{H}$ be an m-round hunter strategy and let $\mathcal{R}$ be a pure m-round rabbit strategy. Then*

$$\mathbf{Pr}\left[hits \geq 1\right] \geq \frac{\mathbf{E}\left[hits\right]^2}{\mathbf{E}\left[hits^2\right]} \quad .$$

The proofs of these propositions are interesting exercises.

2 Efficient hunter strategies

In this section we prove that for a graph G with n nodes and diameter $\operatorname{diam}(G)$, there exists a hunter strategy with expected escape length $\mathcal{O}(n \cdot \log(\operatorname{diam}(G)))$. For this general strategy we cover G with a set of small cycles and then use a subroutine for searching these cycles. We first describe this subroutine: an efficient hunter strategy for catching the rabbit on a cycle. The general strategy is described in section 2.2.

2.1 Strategies for cycles and circles

We prove that there is an $\mathcal{O}(n)$-round hunter strategy on an n-node cycle that has a probability of catching the rabbit of at least $\frac{1}{2H_n+1} = \Omega(\frac{1}{\log(n)})$, where H_n is the n^{th} harmonic number which is defined as $\sum_{i=1}^{n} \frac{1}{i}$. Notice that by repeating this strategy until the rabbit is caught we get a hunter strategy with an expected escape length of $\mathcal{O}(n \cdot \log(n))$. In order to keep the description of the strategy as simple as possible, we introduce a continuous version of the Hunter vs. Rabbit game for cycles. In this version the hunter tries to catch the rabbit on the boundary of a circle with circumference n. The rules are as follows. In every round the hunter and the rabbit reside at arbitrary, i.e., continuously chosen points on the boundary of the circle. The rabbit is allowed to jump, i.e., it can change its position arbitrarily between two consecutive rounds whereas the hunter can cover at most a distance of one. For the notion of *catching*, we partition the boundary of the circle into n distinct half open intervals of length one. The hunter catches the rabbit if and only if there is a round in which both the hunter and the rabbit reside in the same interval. Since each interval of the boundary corresponds directly to a node of the cycle and vice versa we can make the following observation.

Observation 1 *Any hunter strategy for the Hunter vs. Rabbit game on the circle with circumference n can be simulated on the n-node cycle, achieving the same expected escape length.*

The $\mathcal{O}(n)$-round hunter strategy for catching the rabbit on the circle consists of two phases that work as follows. In an *initialization phase* that lasts for $\lceil n/2 \rceil$ rounds the hunter first selects a random position on the boundary as the *starting position* of the following *main phase*. Then the hunter goes to this position. Note that $\lceil n/2 \rceil$ rounds suffice for the hunter to reach any position on the circle boundary. We will not care whether the rabbit gets caught during the initialization phase. Therefore there is no need for specifying the exact route taken by the hunter to get to the starting position.

After the first $\lceil n/2 \rceil$ rounds the *main phase* lasting for n rounds starts. The hunter selects a velocity uniformly at random between 0 and 1 and proceeds in a clockwise direction according to this velocity. This means that a hunter with starting position $s \in [0, n)$ and velocity $v \in [0, 1]$ resides at position $(s + t \cdot v) \bmod n$ in the tth round of the main phase. Obviously this so called RANDOMSPEED-strategy is an $\mathcal{O}(n)$-round hunter strategy since at most $\lceil \frac{3}{2} n \rceil$ nodes are visited. The following analysis shows that it achieves the desired probability of catching the rabbit when simulated on the n-node cycle.

Theorem 2. *On an n-node cycle a hunter using the* RANDOMSPEED-*strategy catches the rabbit with probability at least $\frac{1}{2H_n+1} = \Omega(\frac{1}{\log(n)})$.*

Proof. We prove that the bound holds for the Hunter vs. Rabbit game on the circle. The theorem then follows from Observation 1.

Since the rabbit strategy is oblivious in the sense that it does not know the random choices made by the hunter we can assume that the rabbit strategy is fixed in the beginning before the hunter starts. Thus, let $\mathcal{R} = \mathcal{R}_0, \mathcal{R}_1, \ldots, \mathcal{R}_{n-1}$ denote the rabbit strategy during the main phase, i.e., $\mathcal{R}_t$ is the interval containing the rabbit in round t of this phase.

For this rabbit strategy let $hits$ denote a random variable counting how often the hunter catches the rabbit. This means $hits$ is the number of rounds during the main phase in which the hunter and the rabbit reside in the same interval. The theorem follows by showing that for any rabbit strategy $\mathcal{R}$ the probability $\mathbf{Pr}\left[hits \geq 1\right] =$

$\mathbf{Pr}$ [hunter catches rabbit] is larger than $\frac{1}{2H_n+1}$. For this we use Proposition 2 to estimate $\mathbf{E}[hits]$ and $\mathbf{E}\left[hits^2\right]$. Let $\Omega = [0, n) \times [0, 1]$ denote the sample space of the random experiment performed by the hunter. Further let $S_i^t \subset \Omega$ denote the subset of random choices such that the hunter resides in the ith interval during the tth round of the main phase. The hunter catches the rabbit in round t iff his random choice $\omega \in \Omega$ is in the set $S_{\mathcal{R}_t}^t$. By identifying $S_{\mathcal{R}_t}^t$ with its indicator function we can write $hits(\omega) = \sum_{t=0}^{n-1} S_{\mathcal{R}_t}^t(\omega)$.

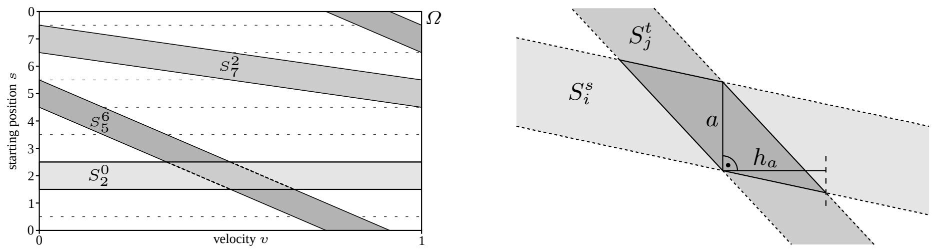

(a) (b)

Fig. 1. (a) The sample space Ω of the RANDOMSPEED strategy can be viewed as the surface of a cylinder. The sets S_i^t correspond to stripes on this surface. (b) The intersection between two stripes of slope $-s$ and $-t$, respectively.

The following interpretation of the sets S_i^t will help derive bounds for $\mathbf{E}[hits]$ and $\mathbf{E}\left[hits^2\right]$. We represent Ω as the surface of a cylinder as shown in Figure 1(a). In this representation a set S_i^t corresponds to a stripe around the cylinder that has slope $\frac{\mathrm{d}s}{\mathrm{d}v} = -t$ and area 1. To see this recall that a point $\omega = (s, v)$ belongs to the set S_i^t iff the hunter position p_t in round t resulting from the random choice ω lies in the ith interval I_i. Since $p_t = (s + t \cdot v) \bmod n$ according to the RANDOMSPEED-strategy we can write S_i^t as $\{(s, v) \mid s = (p_t - t \cdot v) \bmod n \wedge p_t \in I_i\}$ which corresponds to a stripe of slope $-t$. For the area, observe that all n stripes S_i^t of a fixed slope t together cover the whole area of the cylinder which is n. Therefore each stripe has the same area of 1. This yields the following equation.

$$\mathbf{E}[hits] = \mathbf{E}\left[\sum_{t=0}^{n-1} S_{\mathcal{R}_t}^t\right] = \sum_{t=0}^{n-1} \mathbf{E}\left[S_{\mathcal{R}_t}^t\right] = \sum_{t=0}^{n-1} \int_\Omega \frac{1}{n} S_{\mathcal{R}_t}^t(\omega)\, d\omega = 1 \tag{1}$$

Note that $\int_\Omega S_{\mathcal{R}_t}^t(\omega)\, d\omega$ is the area of a stripe and that $\frac{1}{n}$ is the density of the uniform distribution over Ω.
We now provide an upper bound on $\mathbf{E}\left[hits^2\right]$. By definition of $hits$ we have,

$$\begin{aligned}\mathbf{E}\left[hits^2\right] &= \mathbf{E}\left[\left(\sum_{t=0}^{n-1} S_{\mathcal{R}_t}^t\right)^2\right] = \mathbf{E}\left[\sum_{s=0}^{n-1}\sum_{t=0}^{n-1} S_{\mathcal{R}_s}^s \cdot S_{\mathcal{R}_t}^t\right] \\ &= \sum_{s=0}^{n-1}\sum_{t=0}^{n-1} \int_\Omega \frac{1}{n} S_{\mathcal{R}_s}^s(\omega) \cdot S_{\mathcal{R}_t}^t(\omega)\, d\omega \;.\end{aligned} \tag{2}$$

$S^s_{\mathcal{R}_s}(\omega) \cdot S^t_{\mathcal{R}_t}(\omega)$ is the indicator function of the intersection between $S^s_{\mathcal{R}_s}$ and $S^t_{\mathcal{R}_t}$. Therefore $\int_\Omega S^s_{\mathcal{R}_s}(\omega) \cdot S^t_{\mathcal{R}_t}(\omega)\, d\omega$ is the area of the intersection of two stripes and can be bounded using the following lemma.

Lemma 1. *The area of the intersection between two stripes S^s_i and S^t_j with $s, t \in \{0, \ldots, n-1\}$, is at most $\frac{1}{|t-s|}$.*

Proof. W.l.o.g. we assume $t > s$. Figure 1(b) illustrates the case where the intersection between both stripes is maximal. Note that the limitation for the slope values together with the size of the cylinder surface ensures that the intersection is contiguous. This means the stripes only "meet" once on the surface of the cylinder.

By the definition of S^s_i and S^t_j the length of the leg a in the figure corresponds to the length of an interval on the boundary of the circle. Thus $a = 1$. The length of h_a is $\frac{a}{t-s}$ and therefore the area of the intersection is $a \cdot h_a = \frac{a^2}{t-s} = \frac{1}{t-s}$. This yields the lemma. □

Using this Lemma we get

$$\sum_{t=0}^{n-1} \int_\Omega S^s_{\mathcal{R}_s}(\omega) \cdot S^t_{\mathcal{R}_t}(\omega)\, d\omega \leq \sum_{t=0}^{s-1} \frac{1}{|t-s|} + \int_\Omega S^s_{\mathcal{R}_s}(\omega) \cdot S^s_{\mathcal{R}_s}(\omega)\, d\omega + \sum_{t=s+1}^{n-1} \frac{1}{|t-s|}$$

$$= \sum_{t=1}^{s} \frac{1}{t} + \int_\Omega S^s_{\mathcal{R}_s}(\omega)\, d\omega + \sum_{t=1}^{n-s-1} \frac{1}{t} \leq 2H_n + 1 \ .$$

Plugging the above inequality into Equation 2 yields $\mathbf{E}\left[hits^2\right] \leq 2H_n + 1$. Combining this with Proposition 2 and Equation 1 we get $\mathbf{Pr}\left[\text{hunter catches rabbit}\right] \geq \frac{1}{2H_n+1}$ which yields the theorem. □

2.2 Hunter strategies for general graphs

In this section we extend the upper bound of the previous section to general graphs.

Theorem 3. *Let $G = (V, E)$ denote a graph and let $\operatorname{diam}(G)$ denote the diameter of this graph. Then there exists a hunter strategy on G that has expected escape length $\mathcal{O}(|V| \cdot \log(\operatorname{diam}(G)))$.*

Proof. (Sketch) We *cover* the graph with $r = \Theta(n/d)$ cycles $C_1, \ldots, C_r$ of length d where $d = \Theta(diam(G))$, that is, each node is contained in at least one of these cycles (in order to obtain this covering, construct a tour of length $2n - 2$ along an arbitrary spanning tree, cut the tour into subpaths of length $d/2$ and then form a cycle of length d from each of these subpaths). From now on, if hunter or rabbit resides at a node of G corresponding to several cycle nodes, then we assume they *commit* to one of these virtual nodes and the hunter catches the rabbit only if they commit to the same node. This only slows down the hunter.

The hunter strategy now is to choose one of the r cycles uniformly at random and simulate the RANDOMSPEED-strategy on this cycle. Call this a *phase*. Observe that each phase takes only $\Theta(d)$ rounds. The hunter executes phase after phase, each time choosing a new random cycle, until the rabbit is caught. In the following we will show that the success probability within each phase is $\Omega(d/nH_d)$, which implies the theorem.

Let us focus on a particular phase. Suppose the hunter has chosen cycle C_i. Recall that the hunter chooses a random node v from C_i, walks to v and then starts traversing the cycle at a random speed. In the following consider only the hits in the *main phase*, i.e., those rounds after reaching v. The probability that hunter and rabbit reside in the same cycle at the beginning of the main phase is $1/r$. For simplicity, let us assume that the rabbit does not leave this cycle during the main phase. Under this simplifying assumption, Theorem 2 yields immediately that $\mathbf{Pr}\left[hits \geq 1\right] \geq \frac{1}{r(2H_d+1)} = \Omega(d/nH_d)$. In a full version we will prove this bound rigorously without this simplifying assumption.

$\square$

3 Lower bounds and efficient rabbit strategies

We first prove that the hunter strategy for the cycle described in Section 2.1 is tight by giving an efficient rabbit strategy for the cycle. Then we provide lower bounds that match the upper bounds for general graphs given in Section 2.2.

3.1 An optimal rabbit strategy for the cycle

In this section we will prove a tight lower bound for any (mixed) hunter strategy on a cycle of length n. In particular, we describe a rabbit strategy such that every hunter needs $\Omega(n\log(n))$ expected time to catch the rabbit. We assume that the rabbit is unrestricted, i.e., can jump between arbitrary nodes, whereas the hunter is restricted to follow the edges of the cycle.

Theorem 4. *For the cycle of length n, there is a mixed, unrestricted rabbit strategy with escape length $\Omega(n\log(n))$ against any restricted hunter strategy.*

The rabbit strategy is based on a non-standard random walk. Observe that a standard random walk has the limitation that after n rounds, the rabbit is confined to a neighborhood of about $\sqrt{n}$ nodes around the starting position. Hence the rabbit is easily caught by a hunter that just sweeps across the ring (in one direction) in n steps. Also, the other extreme where the rabbit makes a jump to a node chosen uniformly at random in every round does not work, since in each round the rabbit is caught with probability exactly $1/n$, giving an escape length of $O(n)$. But the following strategy will prove to be good for the rabbit. The rabbit will change to a randomly chosen position every n rounds and then, for the next $n-1$ rounds, it performs a "heavy-tailed random walk". For this n-round strategy and any n-round hunter strategy, we will show that the hunter catches the rabbit with probability $\mathcal{O}(1/H_n)$. As a consequence, the expected escape length is $\Omega(n\log n)$, which gives the theorem.

A heavy-tailed random walk. We define a random walk on $\mathbb{Z}$ as follows. At time 0 a particle starts at position $X_0 = 0$. In a *step* $t \geq 1$, the particle makes a random jump $x_t \in \mathbb{Z}$ from position X_{t-1} to position $X_t = X_{t-1} + x_t$, where the jump length is determined by the following heavy-tailed probability distribution $\mathcal{P}$.

$$\mathbf{Pr}\left[x_t = k\right] \;=\; \mathbf{Pr}\left[x_t = -k\right] \;=\; \frac{1}{2(k+1)(k+2)} \;,$$

for every $k \geq 1$ and $\mathbf{Pr}\left[x_t = 0\right] = \frac{1}{2}$. Observe that $\mathbf{Pr}\left[|x_t| \geq k\right] = (k+1)^{-1}$, for every $k \geq 0$. The following lemma gives a property of this random walk that will be crucial for the proof of our lower bound. Due to space limitations, the proof for this lemma appears in the full version.

Lemma 2. *There is a constant $c_0 > 0$, such that, for every $t \geq 1$ and $\ell \in \{-t, \ldots, t\}$, $\mathbf{Pr}\left[X_t = \ell\right] \geq c_0/t$.*

The rabbit strategy. Our n-round rabbit strategy starts at a random position on the cycle. Starting from this position, for the next $n-1$ rounds, the rabbit simulates the heavy-tailed random walk in a wrap-around fashion on the cycle. The following lemma immediately implies Theorem 4.

Lemma 3. *The probability that the hunter catches the rabbit within n rounds is $\mathcal{O}(1/H_n)$.*

Proof. Fix any n-round hunter strategy $\mathcal{H} = \mathcal{H}_0, \mathcal{H}_1, \ldots, \mathcal{H}_{n-1}$. Because of Proposition 1 we only need to estimate $\mathbf{E}[hits]$ and $\mathbf{E}[hits \mid hits \geq 1]$. First, we observe that $\mathbf{E}[hits] = 1$. This is because the rabbit chooses its starting position uniformly at random so that $\mathbf{Pr}[hit(t) = 1] = 1/n$ for $0 \leq t < n$, and hence $\mathbf{E}[hit(t)] = \mathbf{Pr}[hit(t) = 1] = 1/n$. By linearity of expectation, we obtain $\mathbf{E}[hits] = \sum_{t=0}^{n-1} \mathbf{E}[hit(t)] = 1$. Thus, it remains only to show that $\mathbf{E}\left[hits \,|\, hits \geq 1\right] \geq c_0 H_n$. In fact, the idea behind the following proof is that we have chosen the rabbit strategy in such a way that when the rabbit is hit by the hunter in a round then it is likely that it will be hit additionally in several later rounds as well.

Claim. For every $\tau \in \{0, \ldots, \frac{n}{2} - 1\}$, $\mathbf{E}\left[hits \mid fhit(\tau) = 1\right] \geq c_1 H_n$, for a suitable constant c_1.

Proof. Assume hunter and rabbit met at time τ for the first time, i.e., $fhit(\tau) = 1$. Observe that the hunter has to stay somewhere in interval $[\mathcal{H}_\tau - (t-\tau), \mathcal{H}_\tau + (t-\tau)]$ in round $t > \tau$ as he is restricted to the cycle. The heavy-tailed random walk will also have some tendency to stay in this interval. In particular, Lemma 2 implies, for every $t > \tau$, $\mathbf{Pr}[hit(t)] \geq c_0/(t-\tau)$. Consequently, $\mathbf{E}[hits \mid fhit(\tau) \geq 1] \geq 1 + \sum_{t=\tau+1}^{n-1} c_0/(t-\tau)$, which is $\Omega(H_n)$ since $\tau < n/2$. $\square$

With this result at hand, we can now estimate the expected number of repeated hits as follows.

$$\begin{aligned}\mathbf{E}\left[hits \mid hits \geq 1\right] &= \sum_{\tau=0}^{n-1} \mathbf{E}\left[hits \mid \mathit{fhit}(\tau)=1\right] \cdot \mathbf{Pr}\left[\mathit{fhit}(\tau)=1 \mid hits \geq 1\right] \\ &\geq \sum_{\tau=0}^{n/2-1} \mathbf{E}\left[hits \mid \mathit{fhit}(\tau)=1\right] \cdot \mathbf{Pr}\left[\mathit{fhit}(\tau)=1 \mid hits \geq 1\right] \\ &\geq c_1\, H_n \sum_{\tau=0}^{n/2-1} \mathbf{Pr}\left[\mathit{fhit}(\tau)=1 \mid hits \geq 1\right] \quad .\end{aligned}$$

Finally, observe that

$$\sum_{\tau=0}^{n/2-1} \mathbf{Pr}\left[\mathit{fhit}(\tau)=1 \mid hits \geq 1\right] + \sum_{\tau=n/2}^{n-1} \mathbf{Pr}\left[\mathit{fhit}(\tau)=1 \mid hits \geq 1\right] = 1 \quad .$$

Thus, one of the two sums must be greater than or equal to $\frac{1}{2}$. If the first sum is at least $\frac{1}{2}$, then we directly obtain $\mathbf{E}\left[hits \mid hits \geq 1\right] \geq \frac{1}{2} c_1\, H_n$. In the other case, one can prove the same lower bound by going backward instead of forward in time, that is, by summing over the last hits instead of the first hits. Hence Lemma 3 is shown. □

3.2 A lower bound in terms of the diameter

In this section, we show that the upper bound of Section 2.2 is asymptotically tight for the parameters n and $\mathrm{diam}(G)$. We will use the efficient rabbit strategy for cycles as a subroutine on graphs with arbitrary diameter.

Theorem 5. *For any positive integers n, d with $d < n$ there exists a graph G with n nodes and diameter d such that any hunter strategy on G has escape length $\Omega(n \cdot \log(d))$.*

Proof. For simplicity, we assume that n is odd, $d = 3d'$ and $N = (n-1)/2$ is a multiple of d'. The graph G consists of a *center* $s \in V$ and N/d' subgraphs called loops. Each *loop* consists of a cycle of length d' and a linear array of $d'+2$ nodes such that the first node of the linear array is identified with one of the nodes on the cycle and the last node is identified with s. Thus, all loop subgraphs share the center s, otherwise the node sets are disjoint.

Every d' rounds the rabbit chooses uniformly at random one of the N/d' loops and performs the optimal d'-round cycle strategy from Section 3.1 on the cycle of this loop graph. Observe that the hunter cannot visit nodes in different cycles during a phase of length d'. Hence, the probability that the rabbit chooses a cycle visited by the hunter is at most d'/N. Provided that the rabbit chooses the cycle visited by the hunter the probability that it is caught during the next d' rounds is $\mathcal{O}(\frac{1}{H_{d'}})$ by Lemma 3. Consequently, the probability of being caught in one of the independent d'-round games is $\mathcal{O}(\frac{d'}{nH_{d'}})$. Thus, the escape length is $\Omega(nH_{d'})$ which is $\Omega(n \cdot \log(d))$. □

References

1. R. Aleliunas, R. M. Karp, R. J. Lipton, L. Lovasz, and C. Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *Proceedings of the 20th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 218–223, 1979.
2. I. Chatzigiannakis, S. Nikoletseas, N. Paspallis, P. Spirakis, and C. Zaroliagis. An experimental study of basic communication protocols in ad-hoc mobile networks. In *Proceedings of the 5thWorkshop on Algorithmic Engineering*, pages 159–171.
3. I. Chatzigiannakis, S. Nikoletseas, and P. Spirakis. Self-organizing ad-hoc mobile networks: The problem of end-to-end communication. In *Proceedings of the 20th ACM Symposium on Principles of Distributed Computing (PODC 2001)*, 2001.
4. A. Efrat, L. J. Guibas, S. Har-Peled, D. C. Lin, J. S. B. Mitchell, and T. M. Murali. Sweeping simple polygons with a chain of guards. In *Proceedings of the 11th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 927–936, 2000.
5. M. K. Franklin, Z. Galil, and M. Yung. Eavesdropping games: a graph-theoretic approach to privacy in distributed systems. *Journal of the ACM*, 47(2):225–243, 2000.
6. L. J. Guibas, J.-C. Latombe, S. M. LaValle, D. Lin, and R. Motwani. A visibility-based pursuit-evasion problem. *International Journal of Computational Geometry and Applications (IJCGA)*, 9(4):471–493, 1999.
7. K. Hatzis, G. Pentaris, P. Spirakis, and V. Tampakas. Implementation and testing eavesdropper protocols using the DSP tool. In *Proceedings of the 2nd Workshop on Algorithm Engineering (WAE)*, pages 74–85, 1998.
8. L. M. Kirousis and C. H. Papadimitriou. Searching and pebbling. *Theoretical Computer Science*, 47:205–218, 1986.
9. A. S. LaPaugh. Recontamination does not help to search a graph. *Journal of the ACM*, 40(2):224–245, 1993.
10. S. M. LaValle and J. Hinrichsen. Visibility-based pursuit-evasion: The case of curved environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1677–1682, 1999.
11. N. Megiddo, S. L. Hakimi, M. R. Garey, D. S. Johnson, and C. H. Papadimitriou. The complexity of searching a graph. *Journal of the ACM*, 35(1):18–44, 1988.
12. R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
13. S.-M. Park, J.-H. Lee, and K.-Y. Chwa. Visibility-based pursuit-evasion in a polgonal region by a searcher. In *Proceedings of the 28th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 456–468, 2001.
14. T. D. Parsons. Pursuit-evasion in a graph. In Y. Alavi and D. Lick, editors, *Theory and Applications of Graphs*, Lecture Notes in Mathematics, pages 426–441. Springer, 1976.
15. T. D. Parsons. The search number of a connected graph. In *Proceedings of the 9th Southeastern Conference on Combinatorics, Graph Theory and Computing*, pages 549–554, 1978.
16. I. Suzuki and M. Yamashita. Searching for a mobile intruder in a polygonal region. *SIAM Journal on Computing*, 21(5):863–888, 1992.