Combinatorics, Probability and Computing (2003) 12, 225–244. © 2003 Cambridge University Press
DOI: 10.1017/S0963548303005625 Printed in the United Kingdom

Randomized Pursuit-Evasion in Graphs

MICAH ADLER,[1] HARALD RÄCKE,[2†]
NAVEEN SIVADASAN,[3‡] CHRISTIAN SOHLER[2]
and BERTHOLD VÖCKING[4 §]

[1] Department of Computer Science, University of Massachusetts, Amherst, MA 01003-4610, USA
(e-mail: micah@cs.umass.edu)

[2] Heinz Nixdorf Institute
and
Faculty of Computer Science, Electrical Engineering and Mathematics,
University of Paderborn, Germany
(e-mail: {harry, csohler}@upb.de)

[3] Max-Planck-Institut für Informatik, 66123 Saarbrücken, Germany
(e-mail: ns@mpi-sb.mpg.de)

[4] Department of Computer Science, Universität Dortmund, 44221 Dortmund, Germany
(e-mail: voecking@ls2.cs.uni-dortmund.de)

Received 2 May 2002; revised 22 October 2002

We analyse a randomized pursuit-evasion game played by two players on a graph, 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

† Partially supported by the IST Programme of the EU under contract number IST-1999-14186 (ALCOM-FT), DFG grant Me 872/8-1 and DFG Sonderforschungsbereich 376 'Massive Parallelität: Algorithmen, Entwurfsmethoden, Anwendungen'.
‡ Partially supported by the IST Programme of the EU under contract number IST-1999-14186 (ALCOM-FT).
§ A preliminary version of this paper appeared in proceedings of the 29th International Colloquium on Automata, Languages, and Programming (ICALP), 2002.

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

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 regard to all (possibly randomized) rabbit strategies. Our main result is a hunter strategy for general graphs with an escape length of only $\mathcal{O}(n \log(\mathrm{diam}(G)))$ 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 they 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 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 also efficient against an evader who may only move along the edges of the graph.

One approach 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 not necessarily restricted to the graph is caught, is $\mathcal{O}(n \log n)$. 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

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

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 neither player can '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. Strategies are defined as follows.

Definition. 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 $\boldsymbol{\mathcal{S}}$ 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.

For two pure strategies $\mathcal{H}$ and $\mathcal{R}$ of hunter and rabbit, respectively, the *escape length* $\mathrm{el}(\mathcal{H}, \mathcal{R}) := \min\{t \in \mathbb{N}_0 \mid \mathcal{H}_t = \mathcal{R}_t\}$ is the number of rounds until the rabbit is caught. Similarly, $\mathrm{el}(\boldsymbol{\mathcal{H}}, \boldsymbol{\mathcal{R}})$ denotes the *expected escape length* for two mixed strategies $\boldsymbol{\mathcal{H}}$ and $\boldsymbol{\mathcal{R}}$.

We analyse for both players the best expected escape length the player can guarantee for himself, regardless of what the other player does. This means that we analyse hunter strategies that aim to minimize the escape length with respect to a worst case rabbit strategy, and that we analyse rabbit strategies that maximize the escape length with respect to a worst case hunter strategy.

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 every $t \in \mathbb{N}_0$. 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 and 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 affect the results of the paper.

1.2. Previous work

Search games have a long history in the field of games theory. In 1965 Isaacs introduced the so-called *Princess–Monster* game [10]. In this game a (highly intelligent) monster tries to capture a princess in a totally dark room $\mathcal{D}$ with arbitrary shape. Both the monster and the princess are aware of the boundary of the room and the monster catches the princess if their mutual distance becomes smaller than some threshold (which is small in comparison with the extension of $\mathcal{D}$). The monster moves at a known speed using simple

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

motion, that is, the monster moves along continuous trajectories inside $\mathscr{D}$. The princess moves along continuous trajectories but at arbitrary speed.

Since the general game seemed to be hard to analyse, Isaacs also introduced a simpler Princess–Monster game where both the princess and the monster are moving on a closed curve taken as a circle. This game was analysed several years later by Alpern [2] and Zelekin [19]. Finally, Gal presented an analysis of the Princess–Monster game in a convex multidimensional region [7].

The Hunter vs. Rabbit game is a *discrete* variant of the Princess–Monster game that is played in rounds. The most important difference between the two variants is that, in our case, the rabbit (the princess) can use short-cuts not known to the hunter (the monster), that is, the rabbit is allowed to 'jump' from a vertex to any other vertex of the graph. Further, the rabbit is only caught if, at the end of a round, it is on the same node as the hunter. During the motion the rabbit cannot be caught.

A first study of a variant of the Hunter vs. Rabbit game in which the graph is not known to the players 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 the area of mobile *ad hoc* networks, related models are used to design communication protocols (see, *e.g.*, [3, 4]). 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 delivery 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 [16, 17] the graph was considered to be a system of tunnels in which a 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 [12] showed that, if ℓ guards are sufficient to capture the fugitive, then this can be done without re-contamination, *i.e.*, if at any point in time the fugitive is known not to be on edge e, then there is no chance of him entering edge e without being caught in the remainder of the game. Megiddo, Hakimi, Garey, Johnson and Papadimitriou [14] 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 [11].

Pursuit-evasion problems in the plane were introduced by Suzuki and Yamashita [18]. They gave necessary and sufficient conditions for a simple polygon to be searchable by a single pursuer. Later Guibas, Latombe, LaValle, Lin and Motwani [8] 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, Lee and Chwa [15] 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.

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

Efrat, Guibas, Har-Peled, Lin, Mitchell and Murali [5] 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 $\mathrm{diam}(G)$. Observe that $\Omega(n)$ is a lower bound on the escape length against restricted as well as against unrestricted rabbit strategies on every graph with n vertices (the rabbit chooses its first vertex uniformly at random and does not move during the game). 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 anticlockwise, and he then 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 $\Omega(n \cdot \log(d))$. 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, 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.

1.4. Basic concepts

An alternative way to look at the Hunter vs. Rabbit game is to view it as a two-person matrix game in the game theory framework. Entries in the payoff matrix describe escape length; the rabbit is the min player and the hunter is the max player. A game has a unique *value* if both players achieve the same payoff when they use optimal strategies. We

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

note that in our game the payoff matrix is infinitely large, as the number of strategies for each player is infinite. Nevertheless, we show that the game has a value. For this purpose we present upper and lower bounds for the *security value* for each player, *i.e.*, the largest expected payoff a player can guarantee himself, regardless of what the other player does. If we now restrict the number of rounds in the game by an integer t (if the escape length is larger than t then the entry in the payoff matrix is t) then the number of strategies is finite, and hence this restricted game has a unique value. For $t \to \infty$ we obtain that this value is strictly increasing. By the upper bound on the security value it follows that the Hunter vs. Rabbit game has a unique value.

The strategies will be analysed 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 $\mathscr{H}$ and an m-round rabbit strategy $\mathscr{R}$ for a phase. We want to determine the probability that the rabbit is caught during the phase. Therefore we introduce the indicator random variables $\mathrm{hit}(t), 0 \leqslant t < m$ for the event $\mathscr{H}_t = \mathscr{R}_t$ that the pure hunter strategy $\mathscr{H}$ and the pure rabbit strategy $\mathscr{R}$ chosen according to $\mathscr{H}$ and $\mathscr{R}$, respectively, meet in round t of the phase. Furthermore, we define indicator random variables $\mathrm{fhit}(t), 0 \leqslant t < m$ describing first hits, *i.e.*, $\mathrm{fhit}(t) = 1$ if and only if $\mathrm{hit}(t) = 1$ and $\mathrm{hit}(t') = 0$ for every $t' \in \{0, \ldots, t-1\}$. Finally we define $\mathrm{hits} = \sum_{t=0}^{m-1} \mathrm{hit}(t)$.

The goal of our analysis is to derive upper and lower bounds for $\mathbf{Pr}[\mathrm{hits} \geqslant 1]$, the probability that the rabbit is caught in the phase. To analyse the quality of an m-round rabbit strategy we fix a pure hunter strategy $\mathscr{H}$ and derive a lower bound on the probability $\mathbf{Pr}[\mathrm{hits} \geqslant 1]$ using the following proposition, which follows trivially from the definitions.

Proposition 1.1. *Let $\mathscr{R}$ be an m-round rabbit strategy and let $\mathscr{H}$ be a pure m-round hunter strategy. Then*

$$\mathbf{Pr}[\mathrm{hits} \geqslant 1] = \frac{\mathbf{E}[\mathrm{hits}]}{\mathbf{E}[\mathrm{hits} \mid \mathrm{hits} \geqslant 1]}.$$

Similarly, to analyse the quality of an m-round hunter strategy we fix a pure rabbit strategy and apply the following proposition, which is known as the second moment method.

Proposition 1.2. *Let $\mathscr{H}$ be an m-round hunter strategy and let $\mathscr{R}$ be a pure m-round rabbit strategy. Then*

$$\mathbf{Pr}[\mathrm{hits} \geqslant 1] \geqslant \frac{\mathbf{E}[\mathrm{hits}]^2}{\mathbf{E}[\mathrm{hits}^2]}.$$

Proof. We consider the conditional expectations $\mathbf{E}[\mathrm{hits} \mid \mathrm{hits} \neq 0]$ and $\mathbf{E}[\mathrm{hits}^2 \mid \mathrm{hits} \neq 0]$. For these we have

$$\mathbf{E}[\mathrm{hits}^2 \mid \mathrm{hits} \neq 0] - \mathbf{E}[\mathrm{hits} \mid \mathrm{hits} \neq 0]^2 = \mathbf{Var}[\mathrm{hits} \mid \mathrm{hits} \neq 0] \geqslant 0.$$

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

By using $\mathbf{E}[\text{hits} \mid \text{hits} \neq 0] = \frac{\mathbf{E}[\text{hits}]}{\mathbf{Pr}[\text{hits} \neq 0]}$ and $\mathbf{E}[\text{hits}^2 \mid \text{hits} \neq 0] = \frac{\mathbf{E}[\text{hits}^2]}{\mathbf{Pr}[\text{hits} \neq 0]}$, we get

$$\frac{\mathbf{E}[\text{hits}^2]}{\mathbf{Pr}[\text{hits} \neq 0]} \geqslant \frac{\mathbf{E}[\text{hits}]^2}{\mathbf{Pr}[\text{hits} \neq 0]^2},$$

which yields the lemma since $\mathbf{Pr}[\text{hits} \geqslant 1] = \mathbf{Pr}[\text{hits} \neq 0]$. □

Note that in both cases a bound against *all* pure strategies of the other player implies the same bound against mixed strategies, as well.

2. Efficient hunter strategies

In this section we prove that, for a graph G with n nodes and diameter $\text{diam}(G)$, there exists a hunter strategy such that for every rabbit strategy the expected escape length is $\mathcal{O}(n \cdot \log(\text{diam}(G)))$. 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 nth harmonic number, which is defined as $\sum_{i=1}^{n} \frac{1}{i}$. Clearly, by repeating this strategy until the rabbit is caught we get a hunter strategy such that, for every rabbit strategy, the expected escape length is $\mathcal{O}(n \cdot \log(n))$. 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 2.1. *Every 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

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

whether the rabbit gets caught during the initialization phase. Therefore there is no need to specify the exact route taken by the hunter to get to the starting position.

After the first $\lceil n/2 \rceil$ rounds the *main phase* starts, which lasts for n rounds. 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. This strategy is called the RandomSpeed strategy. Clearly, it takes exactly $\lceil \frac{3}{2} n \rceil = O(n)$ rounds. The following analysis shows that it achieves the desired probability of catching the rabbit when simulated on the n-node cycle.

Theorem 2.2. *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 2.1.

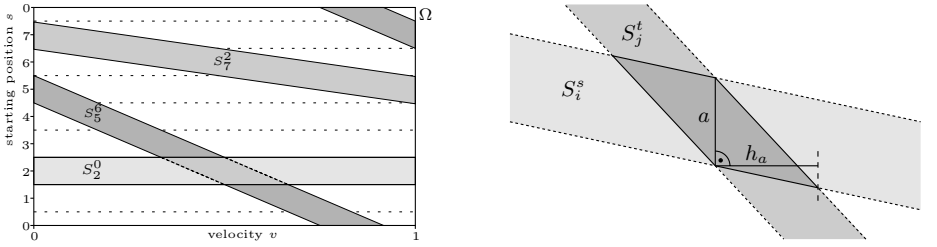
Since the rabbit strategy is oblivious in the sense that it does not know the random choices made by the hunter, we can assume that the rabbit strategy is fixed in the beginning before the hunter starts. Consider an arbitrary pure rabbit strategy $\mathscr{R} = \mathscr{R}_0, \mathscr{R}_1, \ldots, \mathscr{R}_{n-1}$, *i.e.*, $\mathscr{R}_t$ 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 that 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 $\mathscr{R}$, the probability $\mathbf{Pr}[\text{hits} \geqslant 1] = \mathbf{Pr}[\text{hunter catches rabbit}]$ is larger than $\frac{1}{2H_n + 1}$. For this purpose we estimate $\mathbf{E}[\text{hits}]$ and $\mathbf{E}[\text{hits}^2]$ and use Proposition 1.2 to derive a bound for $\mathbf{Pr}[\text{hits} \geqslant 1]$. 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 if and only if his random choice $\omega \in \Omega$ is in the set $S^t_{\mathscr{R}_t}$. By identifying $S^t_{\mathscr{R}_t}$ with its indicator function, we can write $\text{hits}(\omega) = \sum_{t=0}^{n-1} S^t_{\mathscr{R}_t}(\omega)$.

The following interpretation of the sets S_i^t will help in deriving bounds for $\mathbf{E}[\text{hits}]$ and $\mathbf{E}[\text{hits}^2]$. 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 if and only if 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:

$$\mathbf{E}[\text{hits}] = \mathbf{E}\left[\sum_{t=0}^{n-1} S^t_{\mathscr{R}_t}\right] = \sum_{t=0}^{n-1} \mathbf{E}\left[S^t_{\mathscr{R}_t}\right] = \sum_{t=0}^{n-1} \int_{\Omega} \frac{1}{n} S^t_{\mathscr{R}_t}(\omega)\, d\omega = 1 \tag{2.1}$$

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press



(a) (b)

Figure 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

Note that $\int_\Omega S^t_{\mathscr{R}_t}(\omega)\,d\omega$ 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 $\mathbf{E}[\text{hits}^2]$. By definition of hits, we have

$$\mathbf{E}[\text{hits}^2] = \mathbf{E}\left[\left(\sum_{t=0}^{n-1} S^t_{\mathscr{R}_t}\right)^2\right] = \mathbf{E}\left[\sum_{s=0}^{n-1}\sum_{t=0}^{n-1} S^s_{\mathscr{R}_s}\cdot S^t_{\mathscr{R}_t}\right]$$
$$= \sum_{s=0}^{n-1}\sum_{t=0}^{n-1}\int_\Omega \frac{1}{n} S^s_{\mathscr{R}_s}(\omega)\cdot S^t_{\mathscr{R}_t}(\omega)\,d\omega. \tag{2.2}$$

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

Lemma 2.3. *The area of the intersection between two stripes S_i^s and S_j^t with $s,t\in\{0,\dots,n-1\}$, is at most $\frac{1}{|t-s|}$.*

Proof. Without loss of generality 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_i^s and S_j^t, the length of the straight line 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

$$\sum_{t=0}^{n-1}\int_\Omega S^s_{\mathscr{R}_s}(\omega)\cdot S^t_{\mathscr{R}_t}(\omega)\,d\omega \leqslant \sum_{t=0}^{s-1}\frac{1}{|t-s|} + \int_\Omega S^s_{\mathscr{R}_s}(\omega)\cdot S^s_{\mathscr{R}_s}(\omega)\,d\omega + \sum_{t=s+1}^{n-1}\frac{1}{|t-s|}$$
$$= \sum_{t=1}^{s}\frac{1}{t} + \int_\Omega S^s_{\mathscr{R}_s}(\omega)\,d\omega + \sum_{t=1}^{n-s-1}\frac{1}{t} \leqslant 2H_n + 1.$$

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

Plugging the above inequality into equation (2.2) yields $\mathbf{E}[\text{hits}^2] \leqslant 2H_n + 1$. Combining this with Proposition 1.2 and equation (2.1), we get $\mathbf{Pr}[\text{hunter catches rabbit}] \geqslant \frac{1}{2H_n+1}$, 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 2.4. *Let $G = (V, E)$ denote a graph and let* $\text{diam}(G)$ *denote the diameter of this graph. Then there exists a hunter strategy on G that has expected escape length $\mathcal{O}(|V| \cdot \log(\text{diam}(G)))$.*

Proof. We *cover* the graph with $r = \Theta(n/d)$ cycles $C_1, \ldots, C_r$ of length d where $d = \Theta(\text{diam}(G))$, 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.

Now the hunter strategy 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. For the purpose of analysis, we assume that on every cycle the nodes are enumerated consecutively from 1 to d. Instead of directly calculating the probability that the hunter catches the rabbit, we first analyse the probability that at some point in time both of them are on a node with the same number. Let X denote the indicator random variable for this event. We observe that the probability for $X = 1$ is identical to the probability that the hunter catches the rabbit on a cycle of length d. Consequently,

$$\mathbf{Pr}[X = 1] = \Omega(1/H_d).$$

Now we use the fact that the hunter catches the rabbit if and only if they are on a node with the same number *and* they are on the same cycle. If, during a phase, hunter and rabbit are more than once on a node with the same number, we consider only the first time. At this time, the probability that hunter and rabbit are also on the same cycle is $\frac{1}{r}$. We obtain

$$\mathbf{Pr}[\text{hunter catches rabbit} \mid X = 1] \geqslant \frac{1}{r}.$$

We conclude

$$\mathbf{Pr}[\text{hunter catches rabbit}] \geqslant \mathbf{Pr}[\text{hunter catches rabbit} \mid X = 1] \cdot \mathbf{Pr}[X = 1] = \Omega(d/nH_d).$$

□

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

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 3.1. *For the cycle of length n, there is a mixed, unrestricted rabbit strategy such that, for every restricted hunter strategy, the escape length is* $\Omega(n\log(n))$.

The rabbit strategy is based on a nonstandard 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 an arbitrary 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 $X_0 = 0$. At a *step* $t \geqslant 1$, 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 $\mathscr{P}$:

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

for every $k \geqslant 1$ and $\mathbf{Pr}[x_t = 0] = \frac{1}{2}$. Observe that $\mathbf{Pr}[|x_t| \geqslant k] = (k+1)^{-1}$, for every $k \geqslant 0$. The following lemma gives a property of this random walk that will be crucial for the proof of our lower bound.

Lemma 3.2. *There is a constant* $c_0 > 0$, *such that, for every* $t \geqslant 1$ *and* $\ell \in \{-t,\ldots,t\}$, $\mathbf{Pr}[X_t = \ell] \geqslant c_0/t$.

Proof. We will prove the lemma using two claims. The first claim shows a simple monotonicity property of the random walk and the second claim shows that, with at least

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

constant probability, the particle does not move more than distance $\mathcal{O}(\tau)$ within τ steps. (Observe that this does not imply that the expected distance travelled in τ steps is $\mathcal{O}(\tau)$. In fact, it is well known that, under the heavy-tailed distribution $\mathscr{P}$, $\mathbf{E}[|x_t|] = \infty$, so that even the expected distance travelled in only one step is undefined.) □

Claim 3.3. (Monotonicity) *For every* $t \geqslant 0, \ell \geqslant 0, \mathbf{Pr}[X_t = \ell] \geqslant \mathbf{Pr}[X_t = \ell + 1]$.

Proof. We use induction on t. For $t = 0$ the claim is obviously true, as $X_0 = 0$. Assume by induction hypothesis that the claim holds for $t = r$. Define

$$P_i(j) = \mathbf{Pr}[X_r = i - j \wedge x_{r+1} = j] = \mathbf{Pr}[X_r = i - j]\,\mathbf{Pr}[x_{r+1} = j],$$

for $i \geqslant 0, j \in \mathbb{Z}$. Then

$$\mathbf{Pr}[X_{r+1} = \ell] = \sum_{j \in \mathbb{Z}} P_\ell(j) = \sum_{j \geqslant 0} [P_\ell(j) + P_\ell(-j-1)],$$

and

$$\mathbf{Pr}[X_{r+1} = \ell + 1] = \sum_{j \in \mathbb{Z}} P_{\ell+1}(j) = \sum_{j \geqslant 0} [P_{\ell+1}(-j) + P_{\ell+1}(j+1)].$$

As a consequence,

$$\mathbf{Pr}[X_{r+1} = \ell] - \mathbf{Pr}[X_{r+1} = \ell + 1] = \sum_{j \geqslant 0} Y(j),$$

where $Y(j) = P_\ell(j) + P_\ell(-j-1) - P_{\ell+1}(-j) - P_{\ell+1}(j+1)$. Since $\mathscr{P}$ is symmetric, $Y(j) = (\mathbf{Pr}[X_r = \ell - j] - \mathbf{Pr}[X_r = \ell + j + 1])(\mathbf{Pr}[x_{r+1} = j] - \mathbf{Pr}[x_{r+1} = j + 1])$. Observe that both factors are always positive by the induction hypothesis, symmetry, and some shifting. Hence, the claim is proved. □

Claim 3.4. *For every* $\tau \geqslant 0$ *and* $t \in \{0, \ldots, \tau\}$, $\mathbf{Pr}[|X_t| \leqslant \tau] \geqslant 1/2e^4$.

Proof. Observe that the variance of $\mathscr{P}$ is unbounded. Nevertheless, we can use the Chebyshev inequality for bounding the distance travelled by the particle, as follows. Now we truncate the random variables x_i. For this purpose let us fix τ. For simplicity of notation, assume that τ is a multiple of four. For each random variable x_i, we define an auxiliary random variable y_i taking integer values in the range $[-\frac{\tau}{4}, \frac{\tau}{4}]$ such that $\mathbf{Pr}[y_i = k] = \mathbf{Pr}[x_i = k \mid |x_i| \leqslant \frac{\tau}{4}]$. Observe that

$$\mathbf{Pr}[y_i = k] = \frac{(\tau/4 + 2)}{(\tau/4 + 1)} \mathbf{Pr}[x_i = k].$$

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

Since y_t is bounded, it has a finite variance, which can be estimated as follows:

$$\begin{aligned}\mathbf{Var}[y_t] &= \sum_{i=1}^{\tau/4} i^2 \,\mathbf{Pr}[|y_t| = i] \\ &= \frac{(\tau/4+2)}{(\tau/4+1)} \sum_{i=1}^{\tau/4} i^2 \frac{1}{(i+1)(i+2)} \leqslant \frac{\tau}{2}.\end{aligned}$$

Next define $Y_t = \sum_{i=1}^{t} y_i$. Since the random variables y_j are independent, $\mathbf{Var}[Y_t] = \sum_{j=1}^{t} \mathbf{Var}[y_t] \leqslant \frac{\tau^2}{2}$. Furthermore, observe that $\mathbf{E}[Y_t] = 0$, for all $t \geqslant 0$. Hence, applying the Chebyshev inequality gives

$$\mathbf{Pr}[|Y_t| \geqslant \tau] \leqslant \mathbf{Pr}[|Y_t| \geqslant \sqrt{2 \cdot \mathbf{Var}[Y_t]}] \leqslant \frac{1}{2}.$$

Finally, we apply this bound to the original random variables X_t and obtain

$$\begin{aligned}\mathbf{Pr}[|X_t| \leqslant \tau] &\geqslant \mathbf{Pr}\left[|X_t| \leqslant \tau \;\middle|\; \forall_{i=1}^{t} |x_i| \leqslant \frac{\tau}{4}\right] \mathbf{Pr}\left[\forall_{i=1}^{t} |x_i| \leqslant \frac{\tau}{4}\right] \\ &= \mathbf{Pr}[|Y_t| \leqslant \tau] \left(1 - \frac{1}{(\tau/4+2)}\right)^t \\ &\geqslant \frac{1}{2}\left(\frac{1}{e}\right)^4.\end{aligned}$$

Thus Claim 3.4 is proved. □

Fix $t \geqslant 1$. We will use Claim 3.4 in order to prove $\mathbf{Pr}[t \leqslant |X_t| < 4t] \geqslant c_0$, for a suitable constant c_0. Afterwards, we will apply Claim 3.3 to this bound and obtain the lemma.

The probability that there exists $k \in \{1, \ldots, t\}$ with $|x_k| > 2t$, and the first such k, say k^*, fulfils $|x_{k^*}| \leqslant 3t$, is at least

$$\left(1 - \left(1 - \frac{1}{2t+2}\right)^t\right)\left(1 - \frac{2t+2}{3t+2}\right) \geqslant \frac{1}{5}\left(1 - e^{-1/4}\right)$$

since, for every $m \in \mathbb{N}$, we have $\mathbf{Pr}[|x_k| \geqslant m] = (m+1)^{-1}$. Given this event, we observe that there exists $r \in \{1, \ldots, k^*\}$ with $|X_r| \in \{t, \ldots, 4t-1\}$. Let us denote by r^* the smallest such r. Applying Claim 3.4 gives $\mathbf{Pr}[|X_t - X_{r^*}| \leqslant t] \geqslant 1/2e^4$, since $X_t - X_{r^*}$ and X_{t-r^*} are identically distributed. It further follows by symmetry that $\frac{1}{2} \cdot \mathbf{Pr}[|X_t - X_{r^*}| \leqslant t] = \mathbf{Pr}[0 \leqslant X_t - X_{r^*} \leqslant t] = \mathbf{Pr}[-t \leqslant X_t - X_{r^*} \leqslant 0]$. Now observe that if $t \leqslant |X_{r^*}| < 4t$ then $t \leqslant |X_{r^*} - t| < 4t$ or $1 \leqslant |X_{r^*} + t| < 4t$. And so we obtain

$$\begin{aligned}&\mathbf{Pr}[t \leqslant |X_t| < 4t] \\ &\quad\geqslant \frac{1}{2} \cdot \mathbf{Pr}[(|X_t - X_{r^*}| \leqslant t) \mid \exists r \leqslant t : (t \leqslant |X_r| < 4t)] \cdot \mathbf{Pr}[\exists r \leqslant t : (t \leqslant |X_r| < 4t)] \\ &\quad\geqslant \frac{1}{4e^4} \cdot \mathbf{Pr}[\exists r \leqslant t : (t \leqslant |X_r| < 4t)] \\ &\quad\geqslant \frac{1}{20e^4}(1 - e^{-1/4}).\end{aligned}$$

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

If we now define

$$c_0 \stackrel{\text{def}}{=} \frac{1}{60e^4}\left(1 - e^{-1/4}\right),$$

then applying Claim 3.3 gives $\mathbf{Pr}[|X_t| = t] \geqslant c_0/t$. Finally, applying the same claim again, we obtain $\mathbf{Pr}[|X_t| = \ell] \geqslant c_0/t$, for $0 \leqslant \ell \leqslant t$. This completes the proof of Lemma 3.2. □

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 3.1.

Lemma 3.5. *The probability that the hunter catches the rabbit within n rounds is* $\mathcal{O}(1/H_n)$.

Proof. Fix any n-round hunter strategy $\mathscr{H} = \mathscr{H}_0, \mathscr{H}_1, \ldots, \mathscr{H}_{n-1}$. Because of Proposition 1.1, we only need to estimate $\mathbf{E}[\text{hits}]$ and $\mathbf{E}[\text{hits} \mid \text{hits} \geqslant 1]$. First, we observe that $\mathbf{E}[\text{hits}] = 1$. This is because the rabbit chooses its starting position uniformly at random so that $\mathbf{Pr}[\text{hit}(t) = 1] = 1/n$ for $0 \leqslant t < n$, and hence $\mathbf{E}[\text{hit}(t)] = \mathbf{Pr}[\text{hit}(t) = 1] = 1/n$. By linearity of expectation, we obtain $\mathbf{E}[\text{hits}] = \sum_{t=0}^{n-1} \mathbf{E}[\text{hit}(t)] = 1$. Thus, it remains only to show that $\mathbf{E}[\text{hits} \mid \text{hits} \geqslant 1] \geqslant c_1 H_n$ for some constant c_1. 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 3.6. *For every* $\tau \in \{0, \ldots, \frac{n}{2} - 1\}$, $\mathbf{E}[\text{hits} \mid \text{fhit}(\tau) = 1] \geqslant c_1 H_n$, *for a suitable constant* c_1.

Proof. Assume hunter and rabbit meet at time τ for the first time, *i.e.*, $\text{fhit}(\tau) = 1$. Observe that the hunter has to stay somewhere in interval $[\mathscr{H}_\tau - (t-\tau), \mathscr{H}_\tau + (t-\tau)]$ 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 3.2 implies, for every $t > \tau$, that $\mathbf{Pr}[\text{hit}(t)] \geqslant c_0/(t-\tau)$. Consequently, $\mathbf{E}[\text{hits} \mid \text{fhit}(\tau) \geqslant 1] \geqslant 1 + \sum_{t=\tau+1}^{n-1} c_0/(t-\tau)$, which is $\Omega(H_n)$ since $\tau < n/2$. □

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

$$\begin{aligned}
\mathbf{E}[\text{hits} \mid \text{hits} \geqslant 1] &= \sum_{\tau=0}^{n-1} \mathbf{E}[\text{hits} \mid \text{fhit}(\tau) = 1] \cdot \mathbf{Pr}[\text{fhit}(\tau) = 1 \mid \text{hits} \geqslant 1] \\
&\geqslant \sum_{\tau=0}^{n/2-1} \mathbf{E}[\text{hits} \mid \text{fhit}(\tau) = 1] \cdot \mathbf{Pr}[\text{fhit}(\tau) = 1 \mid \text{hits} \geqslant 1] \\
&\geqslant c_1' H_n \sum_{\tau=0}^{n/2-1} \mathbf{Pr}[\text{fhit}(\tau) = 1 \mid \text{hits} \geqslant 1]
\end{aligned}$$

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

for some suitable constant $c_1' = 2c_1$. Finally, observe that

$$\sum_{\tau=0}^{n/2-1} \mathbf{Pr}[\mathrm{fhit}(\tau) = 1 \mid \mathrm{hits} \geqslant 1] + \sum_{\tau=n/2}^{n-1} \mathbf{Pr}[\mathrm{fhit}(\tau) = 1 \mid \mathrm{hits} \geqslant 1] = 1.$$

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 $\mathbf{E}[\mathrm{hits} \mid \mathrm{hits} \geqslant 1] \geqslant c_1 H_n$. In the other case, we 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 3.5 is proved. □

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 3.7. *For every positive integer n, d with $d < n$, there exists a graph G with n nodes and diameter d and a rabbit strategy such that, for every hunter strategy on G, the escape length is $\Omega(n \cdot \log(d))$.*

Proof. For simplicity, we assume that n is odd, $d = 4d'$ and $N = (n-1)/2$ is a multiple of d'. The graph G consists of a *centre* $s \in V$ and N/d' subgraphs called loops. Each *loop* consists of a cycle of length $2d' + 1$ and a simple path of $d' + 1$ nodes such that the first node of the simple path is identified with one of the nodes on the cycle and the last node is identified with s. Thus, all loop subgraphs share the centre 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 3.5. 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 $\Omega(n \cdot \log(d))$. □

4. Trees and directed graphs

In the previous sections, we have seen a restricted hunter strategy such that, for every unrestricted rabbit strategy, the expected escape length is $\mathcal{O}(n \cdot \log(\mathrm{diam}(G))$. Furthermore, we have seen that this bound is optimal against unrestricted rabbits on cycles and several other networks of smaller diameter. This gives rise to the question whether, for every hunter strategy, there is a rabbit strategy such that the escape length is $\Omega(n \cdot \log(\mathrm{diam}(G)))$. We can answer this question negatively. In fact, in the following section we present a hunter strategy on a complete binary tree such that, for every unrestricted rabbit strategy, the expected escape length is $\mathcal{O}(n)$.

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

Subsequently, in Section 4.2 we investigate the Hunter vs. Rabbit game on strongly connected directed graphs. We show that there exists a directed graph and a rabbit strategy such that every restricted hunter needs $\Omega(n^2)$ rounds to catch a restricted rabbit. Furthermore, for every strongly connected directed graph, there is a hunter strategy such that, for every unrestricted rabbit strategy, the expected escape length is $\mathcal{O}(n^2)$.

4.1. A linear time algorithm for binary trees

In this section, we investigate whether there exist graphs for which there is a hunter strategy against unrestricted rabbits with escape length $o(n \cdot \log(\mathrm{diam}(G))$. The following theorem answers this question positively. It gives an example of an n-node network with diameter $\Theta(\log n)$ and escape length $\mathcal{O}(n)$.

Theorem 4.1. *For the complete binary tree T of height h and $m = 2^h$ leaf nodes, there is a hunter strategy such that, for every (unrestricted) rabbit strategy, the expected escape length is $\mathcal{O}(m)$.*

Proof. For simplicity, we assume that h is a power of 2. Furthermore, we initially assume that the rabbit visits only leaf nodes. (Finally, we will remove this assumption.)

We define the level of a node v of T as the height of the subtree T_v rooted at v. The hunter strategy is called *sparse random DFS*, and is defined as follows. The hunter repeats the following four times (starting at the root of T): he chooses a node with height $h/2$ at random, visits it, and applies the same strategy recursively to the subtree T_v. The recursion stops at subtrees of height 2, *i.e.*, subtrees with 4 leaf nodes. Here, for four times, the hunter chooses a leaf node uniformly at random and checks whether the rabbit hides on this leaf node.

The corresponding 4-ary recursion tree is called the *search tree* T_S. Let h_S denote the height of T_S and let L denote the number of leaf nodes of T_S. It is straightforward to see that $h_S = \log h = \log\log m$ and $L = 4^{h_S} = \log^2 m$. Observe that each leaf of T_S corresponds to a visited leaf of T. Furthermore, each edge of T_S corresponds to a path in T that the hunter has to follow in order to reach the root of the selected subtree on the next recursion level. Figure 2 shows a picture of the embedding of the recursion tree T_S into the tree T. Of course, the hunter needs some number of rounds in order to follow the paths that simulate the edges of T_S. Observe that the length of these paths decreases by a factor of two with every level of recursion. However, the number of edges in T_S per recursion level increases by a factor of four with each level. Hence, the leaf level of T_S dominates the execution time, which leads to the following observation. □

Observation 4.2. *The hunter can perform the sparse random DFS in $\mathcal{O}(L)$ rounds, where $L = \log^2 m$ is the number of visited leaf nodes of T.*

Next we investigate the probability that the hunter catches the rabbit within one pass of the described search algorithm.

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

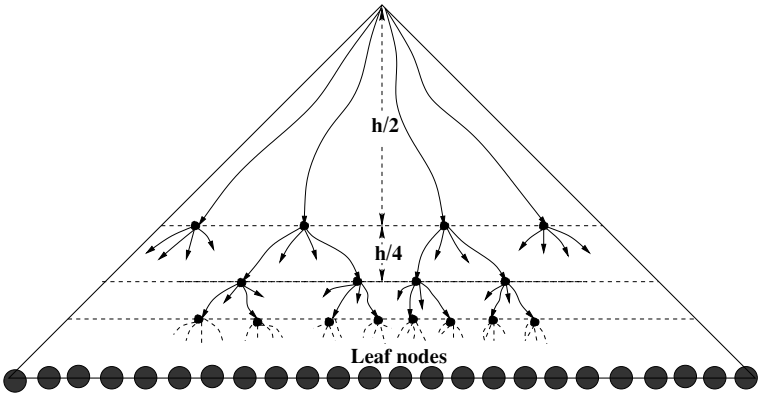


Figure 2. Embedding of the recursion tree T_S into the tree T

Lemma 4.3. *The probability that sparse random DFS finds the rabbit is* $\Omega(L/m)$.

Proof. For $1 \leqslant i \leqslant L$, let r_i denote the round in which the ith leaf node is visited. Let $\mathrm{hit}(i)$ denote a 0/1 random variable which is one if and only if the hunter hits the rabbit in round r_i, and $\mathrm{fhit}(i) = 1$ if this is the first hit. Clearly, for every i, $\mathbf{E}[\mathrm{hit}(i)] = \frac{1}{m}$. Using linearity of expectation, we obtain $\mathbf{E}[\mathrm{hits}] = L/m$. Now, applying Proposition 1.1 yields that the lemma can be shown by proving $\mathbf{E}[\mathrm{hits} \mid \mathrm{hits} \geqslant 1] = \mathcal{O}(1)$. As $\mathbf{E}[\mathrm{hits} \mid \mathrm{hits} \geqslant 1] \leqslant \max_{1 \leqslant i \leqslant L} \{\mathbf{E}[\mathrm{hits} \mid \mathrm{fhit}(i) = 1]\}$, we only need to show $\mathbf{E}[\mathrm{hits} \mid \mathrm{fhit}(i) = 1] = \mathcal{O}(1)$, for $1 \leqslant i \leqslant L$.

Fix an arbitrary $i \in \{1, \ldots, L\}$. We assume that $\mathrm{fhit}(i) = 1$, that is, the hunter meets the rabbit at leaf i of the search tree T_S, and this is the first hit. For a level $\ell \in \{1, \ldots, h_S\}$ of the search tree, let $T_S(\ell)$ denote the complete 4-ary subtree of height ℓ that contains i. If the mapping of i to a leaf of T is fixed, then so is the mapping of the root nodes of the subtrees $T_S(\ell)$, $\ell \in \{1, \ldots, h_S\}$. This partially determines the search tree T_S and hence the leaf nodes visited in addition to i later in the search. We show that the search tree still contains 'enough' randomness such that $\mathbf{E}[\mathrm{hits} \mid \mathrm{fhit}(i) = 1]$ is not too large.

Consider a fixed subtree $T_S(\ell)$ for some value $\ell \in \{1, \ldots, h_S\}$. Let $v(\ell)$ denote the root of $T_S(\ell)$ and let $w(\ell)$ denote the corresponding node in T according to the embedding of T_S in T. We first bound the expected number of hits made by the hunter during the search on the subtree $T_S(\ell)$ *not* including the hits made in $T_S(\ell - 1)$. During this part of the search $3 \cdot 4^{\ell-1}$ leaf nodes of T are visited. These leaf nodes are all contained in the subtree of T rooted at $w(\ell)$. Altogether, this subtree contains 2^{2^ℓ} leaf nodes, since $w(\ell)$ is on level 2^ℓ in T. As each of these nodes is visited with equal probability, the expected number of hits is at most $\frac{3 \cdot 4^{\ell-1}}{2^{2^\ell}}$.

We get an upper bound on $\mathbf{E}[\mathrm{hits} \mid \mathrm{fhit}(i) = 1]$ by summing this value for all subtrees $T_S(1), \ldots, T_S(h_S)$. Hence,

$$\mathbf{E}[\mathrm{hits} \mid \mathrm{fhit}(i) = 1] \leqslant 1 + \sum_{\ell=1}^{h_S} \frac{3 \cdot 4^{\ell-1}}{2^{2^\ell}} \leqslant 3.$$

Hence, Lemma 4.3 is proved. □

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

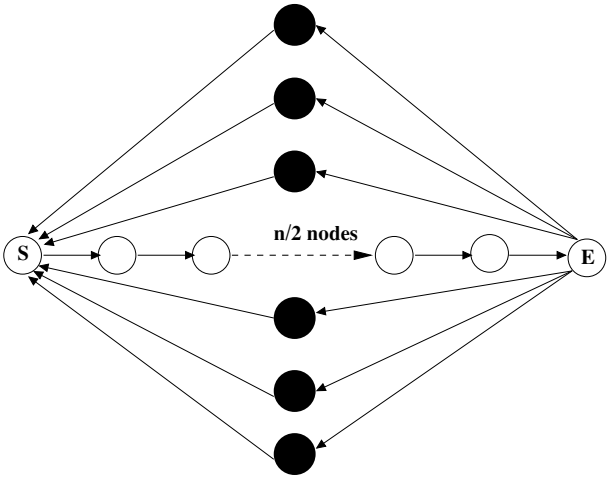


Figure 3. A good graph for the rabbit

Combining Observation 4.2 and Lemma 4.3, we conclude that the escape length is $\mathcal{O}(n)$. Finally, it remains to show how to deal with rabbit strategies that hide on internal nodes of T. To solve this problem we define a *virtual tree* T' which is a complete binary tree of height $h+1$. We embed T' into T such that every node in T hosts at least one leaf of T' and adjacent nodes in T' are hosted by adjacent nodes in T. (The latter requirement means that the *dilation* of the embedding is one.) Then the hunter simulates the random DFS for T' on T. In this way the rabbit cannot avoid the leaves of T', and Theorem 4.1 follows.

It remains to describe the embedding of T' into T. Let T'_1 and T'_2 denote the two disjoint subtrees of height h of T'. We map every node of T'_1 to its counterpart in the isomorphic tree T. Additionally, we map the root of T' to the root of T. If T does not consist of a single node, we apply the same rule recursively with trees T'_2 and T^*, where T^* denotes the subtree of T induced by its internal nodes. In this way, every node of T receives at least one leaf node of T' (and possibly several other internal nodes). □

4.2. Directed graphs

Now we want to consider the Hunter vs. Rabbit game on directed graphs. We slightly alter the definition of restricted strategy for this purpose. In a directed graph $G=(V,E)$ we call a pure strategy $\mathscr{S}$ *restricted*, if either $[\mathscr{S}_t, \mathscr{S}_{t+1}\rangle \in E$ or $\mathscr{S}_t = \mathscr{S}_{t+1}$ holds for every $t \in \mathbb{N}_0$.

Theorem 4.4. *Let G denote an arbitrary directed strongly connected graph with n nodes. Then there is a restricted hunter strategy on G such that, for every unrestricted rabbit strategy, the expected escape length is $\mathcal{O}(n^2)$. Furthermore, there is a directed graph with n nodes, where there exists a restricted rabbit strategy such that, for every restricted hunter strategy, the expected escape length is $\Omega(n^2)$.*

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

Proof. The hunter strategy is defined as follows. In every n rounds, the hunter goes to a node in the graph chosen uniformly at random (this is possible in n steps because the graph is strongly connected) and the hunter meets the rabbit with probability $\Omega(1/n)$. This proves the claim.

We now want to construct a graph and a rabbit strategy such that, for every restricted hunter strategy, the expected escape length is $\Omega(n^2)$. The graph has a directed path of $n/2$ nodes, starting with node S and ending with node E. For each of the remaining nodes (let us call them black nodes), there is an arc from E and an arc to S. Our construction is illustrated in Figure 3 The rabbit initially chooses one of the black nodes at random and stays there for ever. Now, it is easy to see that, if the hunter fails to find the rabbit in a black node, he has to spend $n/2$ rounds to check another black node. This shows a lower bound of $\Omega(n^2)$ even against a stationary rabbit. Hence the theorem is proved. □

References

[1] Aleliunas, R., Karp, R. M., Lipton, R. J., Lovász, L. and Rackoff, C. (1979) Random walks, universal traversal sequences, and the complexity of maze problems. In *Proc. 20th IEEE Symposium on Foundations of Computer Science* (FOCS), pp. 218–223.

[2] Alpern, S. (1974) The search game with mobile hider on the circle. *Differential Games and Control Theory*, pp. 181–200.

[3] Chatzigiannakis, I., Nikoletseas, S., Paspallis, N., Spirakis, P. and Zaroliagis, C. (2001) An experimental study of basic communication protocols in ad-hoc mobile networks. In *Proc. 5th Workshop on Algorithmic Engineering*, pp. 159–171.

[4] Chatzigiannakis, I., Nikoletseas, S. and Spirakis, P. (2001) An efficient communication strategy for ad-hoc mobile networks. In *Proc. 20th ACM Symposium on Principles of Distributed Computing* (PODC), pp. 320–322.

[5] Efrat, A., Guibas, L. J., Har-Peled, S., Lin, D. C., Mitchell, J. S. B. and Murali, T. M. (2000) Sweeping simple polygons with a chain of guards. In *Proc. 11th ACM–SIAM Symposium on Discrete Algorithms* (SODA), pp. 927–936.

[6] Franklin, M. K., Galil, Z. and Yung, M. (2000) Eavesdropping games: A graph-theoretic approach to privacy in distributed systems. *J. Assoc. Comput. Mach.* **47** 225–243.

[7] Gal, S. (1979) Search games with mobile and immobile hider. *SIAM J. Control Optim.* **17** 99–122.

[8] Guibas, L. J., Latombe, J.-C., LaValle, S. M., Lin, D. and Motwani, R. (1999) A visibility-based pursuit-evasion problem. *Internat. J. Comput. Geom. Appl.* **9** 471–493.

[9] Hatzis, K., Pentaris, G., Spirakis, P. and Tampakas, V. (1998) Implementation and testing eavesdropper protocols using the DSP tool. In *Proc. 2nd Workshop on Algorithm Engineering* (WAE), pp. 74–85.

[10] Isaacs, R. (1965) *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*, Wiley, New York/London/Sydney.

[11] Kirousis, L. M. and Papadimitriou, C. H. (1986) Searching and pebbling. *Theoret. Comput. Sci.* **47** 205–218.

[12] LaPaugh, A. S. (1993) Recontamination does not help to search a graph. *J. Assoc. Comput. Mach.* **40** 224–245.

[13] LaValle, S. M. and Hinrichsen, J. (1999) Visibility-based pursuit-evasion: The case of curved environments. In *Proc. IEEE International Conference on Robotics and Automation* (ICRA), pp. 1677–1682.

[14] Megiddo, N., Hakimi, S. L., Garey, M. R., Johnson, D. S. and Papadimitriou, C. H. (1988) The complexity of searching a graph. *J. Assoc. Comput. Mach.* **35** 18–44.

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press

[15] Park, S.-M., Lee, J.-H. and Chwa, K.-Y. (2001) Visibility-based pursuit-evasion in a polygonal region by a searcher. In *Proc. 28th International Colloquium on Automata, Languages and Programming* (ICALP), pp. 456–468.
[16] Parsons, T. D. (1976) Pursuit-evasion in a graph. In *Theory and Applications of Graphs* (Y. Alavi and D. Lick, eds), *Lecture Notes in Mathematics*, Springer, pp. 426–441.
[17] Parsons, T. D. (1978) The search number of a connected graph. In *Proc. 9th Southeastern Conference on Combinatorics, Graph Theory and Computing*, pp. 549–554.
[18] Suzuki, I. and Yamashita, M. (1992) Searching for a mobile intruder in a polygonal region. *SIAM J. Comput*. **21** 863–888.
[19] Zelikin, M. I. (1972) A certain differential game with incomplete information (Russian). *Dokl. Akad. Nauk* **202** 998–1000.

https://doi.org/10.1017/S0963548303005625 Published online by Cambridge University Press