Run or Hide? Fixed and Moving Targets

Fabio Drucker *

March 7, 2013

Abstract

We examine the time it takes for a random walk to hit an adversarial moving target. This problem was introduced by Aldous and Fill as the "cat and mouse" game. They conjectured that for every graph hitting an optimally moving target takes as long in expectation as hitting an optimally placed stationary target. We show this conjecture is false by presenting an instance of a graph where it takes more time for a random walk to hit a moving target than a stationary one. This solves the open question of Aldous and Fill on a "cat and mouse" game.

1 Introduction

We want to consider the following game on an arbitrary graph G. A mouse picks a vertex in the graph in which to place a cat, and a vertex in the graph in which to place itself. The cat begins a lazy random walk on the graph until it catches the mouse.

If once the mouse picks its spot it has to stay there forever, the best the mouse can do is picking the pair of vertices with maximum *hitting time*. On the other hand, the mouse could choose to take its own (non-random) walk on G, trying to evade the cat for longer than it is possible through hiding.

If the mouse can see the current position of the cat it is clear that in many graphs the mouse can avoid the cat forever. For example, in a cycle the mouse could always move in the same direction the cat just moved, thus always maintaining the distance between itself and the cat constant. The case in which the mouse cannot see the cat, however, is more complicated. Aldous and Fill [2] introduce this problem and conjecture than a mouse indeed cannot do better by moving around, and should instead pick the optimal hiding spot and stay there.

*Dartmouth College, Hanover NH.

Note that is important that the walk performed by the cat is lazy. If the cat takes a non-lazy walk, the mouse can avoid the cat forever in bipartite graphs, and it can easily be shown for several other non-bipartite graphs that the mouse can benefit by moving.

In this paper we present a graph that constitutes the first known example of a case where running away extends the expected time before the cat catches the mouse. We first derive the best stationary strategy for the mouse, and then present a simple moving strategy and show it can do better.

2 Known results

In this paper we show that in a particular graph the mouse can do better by moving than by staying stationary. A natural question to ask is for which graphs moving does not help this mouse, and how much better can moving be. Very little is known regarding this question.

Aldous and Fill [2] show that on an $n-$vertex graph the mouse cannot do better than $(en\tau_1)/(e-1)$, where τ_1 is the expected time until, for any vertices i, j, k, the probability that a walk that started at i is at k and the probability that a walk that started at j is at k differs at most by e^{-1}.

Sousi and Winkler [1] present the graph we present in this paper, and they additionally show that in toroidal grids the mouse cannot do better by moving than by remaining stationary. They also show that beyond its intrinsic interest as a game, the cat and mouse game is also relevant for estimating the mixing time of a finite state Markov chain. They prove that the mixing time of a chain can differ by at most a constant factor from the hitting time of a moving big set. A big set is a set of vertices such in the stationary distribution the combined probability of elements of the set is at least 1/4.

3 Better to run than to hide

We now present a graph $G = (V, E)$ where it takes longer to hit a moving than a static target. Our goal is to find a graph where a *distant vertex* (a vertex that maximizes expected hitting time) which can also be reached through a shortcut: a very short path that is followed with low probability. Then, the best stationary strategy will be to stay at the distant vertex. But we will be able to improve on this strategy by away from the shortcut path for a few steps, and only then moving to the distant vertex. Thus, if the walk uses the shortcut, the distant vertex is reached while the mouse is not there, so we improve on the stationary strategy. Otherwise, we move to the distant vertex and do as well as the stationary strategy. Thus, we do better than the stationary strategy with low probability and equally well otherwise.

We now present an example of one such graph $G = (V, E)$. The set of vertices V contains 12 sets of 4 vertices each, each set labeled 0 through 11. We label each vertex as $i(a, b)$, were i is the number assigned to its set, and $a, b \in \{0, 1\}$ (with no vertices sharing the same label). We think of each of these sets as a (super)vertex in a cycle, so we draw an edge between every pair of vertices in adjacent sets. We think of the vertices in a given set as arranged in a unit square, with (a, b) indicating the coordinates of that vertex in its square (a indicating the $x-$coordinate, b the $y-$coordinate). Formally, for every i, j such that $|i - j| = 1$, and every $a, b, c, d \in \{0, 1\}$, we draw an edge between $i(a, b)$ and $j(c, d)$. We call this edges *short*.

Additionally, we add some *long* edges that connect vertices not in adjacent sets: For every i even, we connect each vertex in i with the vertices in $i + 3 \mod 12$ that has the same x-coordinate and with the vertices in $i - 3 \mod 12$ that have the same y-coordinate.

To prove that in this graph it is better to run than to hide, we first show that the maximum hitting time is achieved between $0(0, 0)$ and $6(1, 1)$. Then we note that a random walk starting at $0(0, 0)$ has positive chance of reaching set 6 in 2 steps (using 2 long edges) but 0 chance of reaching cluster 5 in at most 2 steps. Thus, a strategy that stays in $5(1, 1)$ for the first 2 steps and then moves to $6(1, 1)$ has higher expected hitting time than a strategy that stays at $6(1, 1)$.

Lemma 3.1. *Let X be a lazy random walk on G. Then,*

$$\max_{x,y} \mathbf{E}_x[\tau_y] \;=\; \mathbf{E}_{0(0,0)[\tau_{6(1,1)}]} \;.$$

Proof. First we observe that a random walk X from $0(0, 0)$ to $i(a, b)$, with $i \neq 0$ can be divided into 2 stages: (1) a walk from $0(0, 0)$ to cluster i, and (2) a walk from the first vertex in i hit to $i(a, b)$. If $i \notin \{0, 3, 6, 9\}$, when X hits set i it has used at least one short edge (since using only long edges a walk starting at 0 can only reach sets $0, 3, 6$ and 9). Thus, for any such i the probability that the first vertex in i hit is $i(a, b) = 1/4$ for any vertex in $a, b \in \{0, 1\}$. If

$i = 6$, the same statement holds: If before reaching set 6 a short edge was used, then by the previously used argument both coordinates were randomized. Otherwise, the walk to set 6 consisted of two consecutive long edges, once from $0(0,0)$ to either set 3 or 9, and the other from that set to set 6. In this case, the first such edge randomizes one of the coordinates and the second randomizes the other coordinate, so every vertex in set 6 is equally likely to be hit first. Thus, by symmetry, stage 2 of this walk takes the same time in expectation for any values of a, b and for any $i \notin \{0,3,9\}$. Let z be the expected time for a walk starting at a uniformly random vertex in set i takes to hit $i(a,b)$.

Thus, to determine $\max_{i \notin \{0,3,9\}, a,b \in \mathbb{Z}_2} \mathbf{E}_{0(0,0)} \tau_{i(a,b)}$ we only need to determine which set takes longer for a random walk starting at $0(0,0)$ to hit. Let T_i be the first time in which X hits vertex i, and let $h(i) = \mathbf{E}[T_i]$. Then, $\mathbf{E}_{0(0,0)} \tau_{i(a,b)} = \mathbf{E}(T_i) + z$, for $i \notin \{0,3,9\}$.

Let $h(i) = \mathbf{E}[T_i]$. By taking advantage of symmetries in the graph we can calculate all values of $h(i)$ with a system of 6 linear equations, obtaining that

$$h(6) = 16, \ h(5) = h(7) = 16, \ h(4) = h(8) = 15,$$
$$h(3) = h(9) = 13, \ h(2) = h(10) = 13, \ h(1) = h(11) = 10.$$

Thus, $\arg\max_{i(a,b),\ i \notin \{0,3,9\}} \mathbf{E}_{0(0,0)} \tau_{i(a,b)} = 6(1,1)$, since no set takes longer to hit than set 6.

It remains to analyze the expected time for a walk starting at $0(0,0)$ to hit vertices in sets 3, 9 or 0 (other than $0(0,0)$). Since by symmetry it takes as long to hit a vertex in 3 as an equivalent vertex in 9, we analyze the time taken to hit vertices in 3 but not 9. In the first step the walk may use a long edge from 0 to 3, in which case the walk is at vertices $3(0,0)$ or $3(1,0)$ with probability 1/2 each, but cannot be at $3(0,1)$ or $3(1,1)$. Thus, when the walk first get to set 3 the distribution among vertices of 3 is not uniform, so we cannot use our previous argument. However, as soon as X uses an edge other than a long edge from 0 to 3, our previous argument does apply, and the next time the walk reaches set 3 it will be equally likely to hit any of its vertices. Let $E_{0,3}$ be the set of long edges between 0 and 3. Then, we can upper bound the expected hitting time to vertices in 3 by the time it takes for the walk to hit set 3 after using an edge not in $E_{0,3}$, plus z after that happens.

Let T'_3 be the first time X hits set 3 after using an edge not in $E_{0,3}$. The first step of the walk, the probability that an edge in $E_{0,3}$ is used is 1/6, since from any vertex in set 0 there are 2 edges to 0, 2 to 9, and 4 each to 11 and 1. In general, at each time step until we use an edge not in $E_{0,3}$ we have a 5/6 chance of using an edge not in $E_{0,3}$. Let p_i be the probability that the walk uses an edge not in $E_{0,3}$ for the first time at time i. Thus, $p_i = (1/6)^{i-1}(5/6)$. Then,

$$\begin{aligned}
\mathbf{E}[T_3'] &= \sum_{i=1}^{\infty} p_i i + \sum_{i=1,3,\ldots} p_i A_1 + \sum_{i=2,4,\ldots} p_i A_2] \\
&= \frac{6}{5} + A_1 \sum_{i=1,3,\ldots} p_i + A_2 \sum_{i=2,4,\ldots} p_i \\
&= \frac{6}{5} + \frac{6A_1}{7} + \frac{A_2}{7}, \qquad (1)
\end{aligned}$$

where A_1 is the expected time to hit set 3 if we first use an edge not in $E_{0,3}$ at an odd time step (when we are at vertex 0) and A_2 the expected time to hit set 3 if we do that at an even time step. Given this,

$$\begin{aligned}
A_1 &= \frac{2}{5}\mathbf{E}[1 \to 3] + \frac{2}{5}\mathbf{E}[11 \to 3] + \frac{1}{5}\mathbf{E}[9 \to 3]\ , \\
&= \frac{2}{5}h(2) + \frac{2}{5}h(4) + \frac{1}{5}h(6) = \frac{72}{5}\ . \\
A_2 &= \frac{2}{5}\mathbf{E}[4 \to 3] + \frac{2}{5}\mathbf{E}[2 \to 3] + \frac{1}{5}\mathbf{E}[6 \to 3]\ , \\
&= \frac{2}{5}h(1) + \frac{2}{5}h(1) + \frac{1}{5}h(3) = \frac{53}{5}\ .
\end{aligned}$$

Plugging this into (1) yields that $\mathbf{E}[T_3'] = 527/35 < 16$, completing the proof. We can perform a similar calculation for the expected time to get to set 0 after both coordinates have been randomized, and get that this expected time is $452/35 < 16$.

□

Theorem 3.2. *There exists a sequence f of vertices in G such that a random walk takes longer in expectation to hit $f(t)$ at time t than it does to hit any static vertex.*

Proof. Let $f(1) = f(2) = 5(1,1)$, and $f(t) = 6(1,1)$ for $t \geq 3$. We will show that it takes longer for a lazy random walk X from $0(0,0)$ to hit f than to hit $6(1,1)$.

First, we note that if X does not hit f at times 1 or 2, since for $t \geq 3$ $f(t) = 6(1,1)$, in expectation it takes the same amount of time to hit f as it does to hit $6(1,1)$.

We then only need to focus on the event of the walk hitting either f or $6(1,1)$ at times 1 or 2. During those times f is in set 5. There is no path of length at most 2 from set 0 to set

5, so the probability of X hitting f at time 1 or 2 is 0. On the other hand, X could reach $6(1,1)$ in two steps by using 2 long edges, first from set 0 to sets 3 or 9 and then from there to set 6. Recall that if X hits set 6 in two steps, it will hit $6(1,1)$ at step 2 with probability 1/4. Thus, there is a nonzero probability of X hitting $6(1,1)$ at time 2. □

References

[1] P. Sousi and P. Winkler, Mixing Time and Moving Targets, *http://arxiv.org/pdf/1210.5236*.

[2] D. Aldous and J. Fill, Reversible Markov Chains and Random Walks on Graphs, *http://www.stat.berkeley.edu/ aldous/RWG/book.html*.