Patrolling a Pipeline

Steve Alpern[1], Thomas Lidbetter[2], Alec Morton[3], and Katerina Papadaki[2]

[1]Warwick Business School, University of Warwick,
Coventry, CV4 7AL, United Kingdom
[2]Deparment of Mathematics, London School of Economics,
Houghton Street, London, WC2A 2AE, United Kingdom
[3]Department of Management Science, University of Strathclyde, 199 Cathedral
Street, Glasgow, G4 0QU, United Kingdom
steve.alpern@wbs.ac.uk
{t.r.lidbetter,k.p.papadaki}@lse.ac.uk
alec.morton@strath.ac.uk

Abstract. A pipeline network can potentially be attacked at any point and at any time, but such an attack takes a known length of time. To counter this, a Patroller moves around the network at unit speed, hoping to intercept the attack while it is being carried out. This is a zero sum game between the mobile Patroller and the Attacker, which we analyze and solve in certain cases.

Keywords: patrolling, zero-sum game, networks

1 Introduction

A game theoretic model of patrolling a graph was recently introduced in [1], in which an Attacker chooses a node of a graph to attack at a particular time and a Patroller chooses a walk on the nodes of the graph. The game takes place in discrete time and the attack lasts a fixed number of time units. For given mixed strategies of the players, the payoff of the game is the probability that the attack is intercepted by the Patroller: that is, the probability that the Patroller visits the node the Attacker has chosen during the period in which the attack takes place. The Patroller seeks to maximize the payoff and the Attacker to minimize it, so the game is zero-sum.

In [1], several general results of the game are presented along with solutions of the game for some particular graphs. This work is extended in [5], in which line graphs are considered. The game is surprisingly difficult to solve on the line graph, and the optimal policy for the Patroller is not always, as one might expect, the strategy that oscillates to and fro between the terminal nodes. Rather, depending on the length of time required for the attack to take place, it may be optimal for the Patroller to stay around the two ends of the line with some positive probability.

In this paper we present a new continuous game theoretic model of patrolling, in a similar spirit to [1], but on a continuous network, so that the attack may take place at any point of the network (not just at nodes). We also model time as being continuous, rather than discrete. This is a better model for a situation in which a pipeline may be disrupted at any point.

At first glance, this might appear to be a more complicated game to analyze. However, it turns out that continuity simplifies matters, and we are able to solve the game for Eulerian networks (Section 3) and for line networks (Section 4). The solution of the game on the line network is considerably easier to derive than for the discrete analogue, and we also show that the value oOf the latter game converges to that of the former as the number of nodes of the graph approaches infinity.

A game theoretical approach to patrolling problems has been successful in real life settings, for example in [6] and [7]. Other work on game theoretic models of patrolling a network include [2] and [4].

2 Definition of the game

We start by defining a continuous time patrolling game, where the Patroller moves at unit speed along a network Q with given arc lengths, and the Attacker can attack at any point of the network (not just at nodes). In this section we define the game formally and describe each of the players' strategy spaces.

The network Q can be viewed as a metric space, with $d(x,y)$ denoting the arc length distance, so we can talk about 'the midpoint of an arc' and other metric notions. We assume that the game has an infinite time horizon and that a Patroller pure strategy is a unit speed (Lipshitz continuous) path $w : [0,\infty) \to Q$, in particular, one satisfying

$$d\left(w\left(t\right), w\left(t'\right)\right) \leq |t - t'|, \text{ for all } t, t' \geq 0.$$

For the Attacker, a pure strategy is a pair $[x, I]$, where $x \in Q$ and $I \subset [0,\infty)$ is an interval of length r. It is sometimes useful to identify I with its midpoint y, where $I = I_y = [y - r/2, y + r/2]$. Thus $y \in [r/2, \infty)$.

The payoff function, taking the Patroller as the maximizer, is given by

$$P\left(w, \{x, y\}\right) = \begin{cases} 1 \text{ if } w\left(t\right) = x \text{ for some } t \in I_y, \\ 0 \qquad\qquad \text{otherwise.} \end{cases} \tag{1}$$

Hence the value, if it exists, is the probability that the attack is intercepted. Note that in this scenario the pure strategies available to both players are uncountably infinite, so the von Neuman minimax theorem no longer applies. Furthermore, the payoff function is not continuous (in either variable), so minimax theorems using that property also don't apply. For example, if w is the constant function x, then $P\left(w, [x, I]\right) = 1$, however an arbitrarily small perturbation of w or x can have $P\left(w', [x', I]\right) = 0$. However, in the examples we study in this paper we show that the value exists by explicitly giving optimal strategies for the players.

3 General results

We start by giving upper and lower bounds for the value of the game for general networks. First, we define the *uniform* attack strategy.

Definition 1. *The* ***uniform attack strategy*** *chooses to attack in the time interval* $[0, r]$ *at a uniformly random point on* Q*. More precisely, the probability the attack takes place in a region* A *of the network is proportional to the total length of* A*.*

We use the uniform attack strategy to deduce a simple lower bound on the value of the game. We denote the total length of Q by μ.

Lemma 1. *The uniform attack strategy guarantees that the probability* P *of interception is no more than* r/μ*.*

We also define a natural strategy for the Patroller. Recall that a *Chinese Postman Tour* (*CPT*) of the network Q is a minimum length tour that contains every point of Q. We denote the length of a CPT by $\bar{\mu}$. It is well known [3] that there are polynomial time algorithms (polynomial in the number of nodes of the network) that calculate $\bar{\mu}$. It is easy to see that $\bar{\mu} \leq 2\mu$, since doubling each arc of the network results in a new network whose nodes all have even degree and therefore contains an Eulerian tour.

Definition 2. *Fix a CPT,* $w : [0, \infty) \to Q$ *that repeats with period* $\bar{\mu}$*. The* ***uniform CPT strategy*** $\bar{w} : [0, \infty) \to Q$ *for the Patroller is defined by*

$$\bar{w}(t) = w(t + T),$$

where T *is chosen uniformly at random from the interval* $[0, \bar{\mu}]$*. In other words, the Patroller chooses to start the CPT at a random point along it.*

This strategy gives an upper bound on the value of the game.

Lemma 2. *The uniform CPT strategy guarantees that the probability* P *of interception is at least* $r/\bar{\mu}$*.*

Lemmas 1 and 2 give upper and lower bounds on the value of the game. If the network is Eulerian (that is, the network contains a tour that does not repeat any arcs) then $\mu = \bar{\mu}$ and Lemmas 1 and 2 imply that the value of the game is $r/\mu = r/\bar{\mu}$. We sum this up in the theorem below.

Theorem 1. *The value* V *of the game satisfies*

$$\frac{r}{\bar{\mu}} \leq V \leq \frac{r}{\mu}.$$

If the network is Eulerian then both bounds are tight, $V = r/\mu = r/\bar{\mu}$*, the uniform attack strategy is optimal for the Attacker and the uniform CPT strategy is optimal for the Patroller.*

Writing P^* for the probability the uniform CPT strategy intercepts the attack, we note that since it is true for any network that $\bar{\mu} \leq 2\mu$, we have

$$V \leq \frac{r}{\mu} \leq 2\left(\frac{r}{\bar{\mu}}\right) = 2P^*.$$

This shows that the value of the game is no more than twice the interception probability guaranteed by the uniform CPT strategy.

4 Solution on the line network

We now give a complete solution to the game on a line of unit length, that is the closed unit interval $[0,1]$. The Attacker picks a point $x \in [0,1]$ and an interval $I \subset [0,\infty)$ of length r. The Patroller picks a unit speed walk w on the unit interval, $w : \mathbb{R}^+ \to [0,1]$. The attack is intercepted if $w(t) = x$, for some $t \in I$. We assume $0 \leq r \leq 2$, otherwise the Patroller can always intercept the attacks by oscillating between the endpoints of the unit interval.

4.1 The Case $r > 1$

We begin by assuming the attack interval r is relatively large compared to the size of the line, in particular when $r > 1$. We shall see that the following strategies are optimal.

Definition 3. *Let the* ***diametrical*** *Attacker strategy be defined as follows: choose y uniformly in $[0,1]$ and attack equiprobably at one of the endpoints $x = 0$ or 1 during the time interval $I = [y, y + r]$.*

For the Patroller, the ***oscillation strategy*** *is defined as the strategy where the Patroller randomly picks a point x on the unit interval and a random direction and oscillates from one endpoint to the other.*

We note that the oscillation strategy is simply the uniform CPT strategy as defined in Definition 2, and thus ensures a probability $P \geq r/\bar{\mu} = r/2$ of interception, by Lemma 2.

We can show that the diametrical strategy ensures the attack will not be intercepted with probability any greater than $r/2$.

Lemma 3. *If $r \geq 1$ and the Attacker adopts the diametrical strategy then for any path w the attack is intercepted with probability $P \leq r/2$.*

We have the following corollary:

Theorem 2. *The diametric Attacker strategy and the oscillation strategy are optimal strategies and give value $V = r/2$.*

Proof. This follows directly from Lemma 2 and Lemma 3.

4.2 The Case $r \leq 1$

Now we consider the case of $r \leq 1$. In this case r is small compared to 1 (the size of the unit interval), thus the Patroller stays at the end with some probability and oscillates between the endpoints of the unit interval with the remaining probability.

Let q be the quotient and ρ the remainder when r divides 1. Thus $1 = rq + \rho$, where q is an integer and $0 \leq \rho < r$. Let $k = r + \rho$. We first define the Attacker strategies.

Definition 4. *Consider the following Attacker strategy, which we call r-**attack strategy**, that is performed at a random point in time, here we start it at time 0:*

1. *Attack at points $\mathcal{E} = \{0, r, 2r, \ldots, (q-1)r, 1\}$, starting attacks equiprobably between times $[0, r]$, each with total probability $\frac{r}{1+r}$. We call these the **external attacks**.*
2. *Attack at the midpoint of $(q-1)r$ and 1, which is the point $1 - \frac{r+\rho}{2} = 1 - \frac{r}{2}$, starting the attack equiprobably between times $\left[\frac{r-\rho}{2}, \frac{r+\rho}{2}\right]$ with total probability $\frac{\rho}{1+r}$. We call this the **internal attack**.*

The attacks are shown in Figure 1. The horizontal axis is time and the vertical axis is the unit interval.

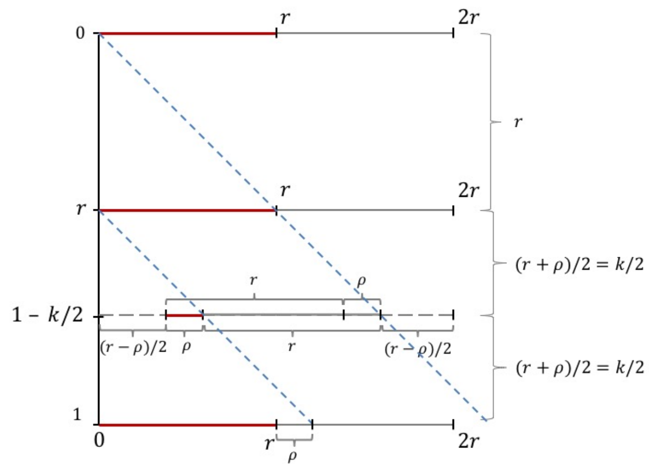

Fig. 1. The r-attack strategy is shown. The starting points of the attacks are shown in red.

Let $f(t)$ be the probability of interception at an external attack point if the Patroller is present there at time t. Let $g(t)$ be this probability for the internal

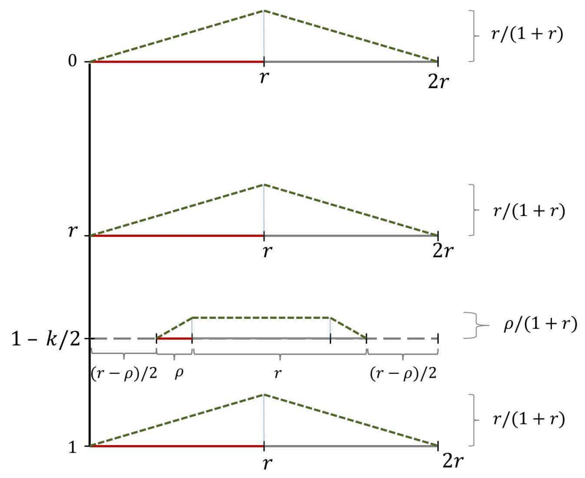

Fig. 2. The probability of interception at each point in time t is shown both for external attacks, $f(t)$, and for internal attacks, $g(t)$, for the r-attack strategy.

attack point. These probability functions for the r-attack strategy are shown in Figure 2.

The functions f and g are as follows:

$$f(t) = \begin{cases} \frac{t}{1+r}, & t \in [0, r] \\ \frac{2r-t}{1+r}, & t \in [r, 2r] \\ 0, & t \in [2r, \infty) \end{cases} \tag{2}$$

$$g(t) = \begin{cases} 0, & t \in \left[0, \frac{r-\rho}{2}\right] \\ \frac{t - \frac{r-\rho}{2}}{1+r}, & t \in \left[\frac{r-\rho}{2}, \frac{r+\rho}{2}\right] \\ \frac{\rho}{1+r}, & t \in \left[\frac{r+\rho}{2}, 2r - \frac{r+\rho}{2}\right] \\ \frac{2r - \frac{r-\rho}{2} - t}{1+r}, & t \in \left[2r - \frac{r+\rho}{2}, 2r - \frac{r-\rho}{2}\right] \\ 0, & t \in \left[2r - \frac{r-\rho}{2}, \infty\right) \end{cases} \tag{3}$$

We now define some Patroller strategies.

Definition 5. *Consider the Patroller strategies where the Patroller plays a mixture of oscillations of the interval* $[0, 1]$ *(the* ***big oscillations****) with probability* $\frac{1}{1+r}$*, and oscillations of the intervals* $\left[0, \frac{r}{2}\right]$ *and* $\left[1 - \frac{r}{2}, 1\right]$ *(the* ***small oscillations****) with probability of* $\frac{r}{2(1+r)}$ *on each. We call this* ***mixed-oscillation strategy****.*

The mixed oscillation strategy is shown in Figure 3. Note that the small oscillations have period r and thus intercept all attacks in the respective intervals. By attacking at 0 or 1 the Attacker secures $\frac{r}{2(1+r)} + \frac{r}{2} \times \frac{1}{1+r} = \frac{r}{1+r}$, since the big oscillation intercepts attacks at the endpoints with probability $\frac{r}{2}$. Any attacks in the open intervals $\left(0, \frac{r}{2}\right)$ and $\left(1 - \frac{r}{2}, 1\right)$, are dominated by attacks at endpoints.

Attacking in $\left[\frac{r}{2}, 1-\frac{r}{2}\right]$ secures an interception probability of $\frac{2r}{2} \times \frac{1}{1+r} = \frac{r}{1+r}$, since at points in $\left[\frac{r}{2}, 1-\frac{r}{2}\right]$, the big oscillation in each of its period time intervals of length 2, it intercepts attacks that start at two time intervals each of length r. Hence, $V \geq \frac{r}{1+r}$.

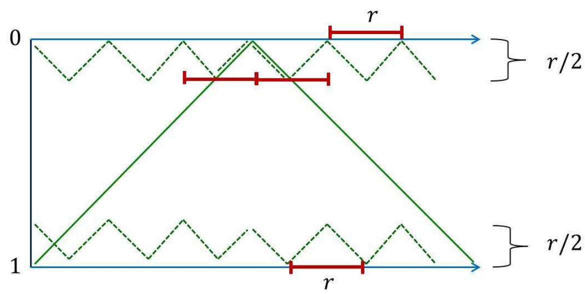

Fig. 3. The mixed oscillation strategy, where the horizontal axis is time and the vertical axis is the unit interval.

Theorem 3. *If $r \leq 1$, then the r-attack strategy and the mixed-oscillation strategy are optimal and the value of the game is $V = \frac{r}{1+r}$.*

4.3 Relation to Discrete Patrolling Game

The discrete analogue of our game, introduced in [1] was solved for line graphs in [5]. It is interesting (and reassuring) to find that the value of the discrete game converges to the value of the continuous game as the number of nodes tends to infinity.

We briefly describe the set-up of the discrete game. The game is played on a line graph with n nodes in a discrete time horizon $\mathcal{T} = \{1, 2, \ldots, T\}$. The Attacker chooses an *attack node* at which to attack and a set of m successive time periods in $\mathcal{T}$, which is when the attack takes place. The Patroller chooses a walk on the graph. As in the continuous case, the payoff of the game, which the Attacker seeks to minimize and the Patroller to maximize, is the probability that the Patroller visits the attack node while the attack is taking place.

The value of the game depends on the relationship between n and m, and the solution divides into 5 cases (see Theorem 6 of [5]). We are interested in fixing the ratio $r = m/n$ and letting n tend to infinity, therefore the solution of two of the cases of the game from [1] are irrelevant: in particular the case when $m = 2$, and the case when $n = m + 1$ or $n = m + 2$. The case $n < (m + 2)/2$ (corresponding to the case $r \geq 2$ in the continous case) is also uninteresting, since then the value is 1. Therefore we are left with two cases, whose solutions we summarize below.

Theorem 4 (From Theorem 6 of [5]). *The value V of the discrete patrolling game on the line is*

1. $V = m/(2n-2)$ *if* $(m+1)/2 \le n \le m+1$, *and*
2. $V = m/(n+m-1)$ *if* $n \ge m+3$, *or* $n = m+2$ *and* $m \ge 3$ *is odd.*

We now consider the behaviour of the value of the discrete game as $n \to \infty$, assuming that the ratio $r = m/n$ is fixed. In the first case of Theorem 4, as $n \to \infty$, the condition $(m+1)/2 \le n \le m+2$ becomes $1 \le r \le 2$ and we have

$$V = \frac{m}{2n-2} = \frac{r}{2-2/n} \to \frac{r}{2},$$

as $n \to \infty$. This corresponds to the solution of the continuous game as given in Theorem 2.

In the second case of Theorem 4, as $n \to \infty$, the condition on m becomes $r \le 1$ and we have

$$V = \frac{m}{n+m-1} = \frac{r}{1+r-1/n} \to \frac{r}{1+r},$$

as $n \to \infty$. Again, this corresponds to the solution of the continuous game as given in Theorem 3.

5 Conclusion

We have introduced a new game theoretic model of patrolling a continuous network in continuous time, analagous to the discrete patrolling game introduced in [1]. We have given general bounds on the value of the game and solved it in the case that the network is Eulerian or if it is a line.

We are optimistic that our results on the line network can be extended to a larger class of networks, such as stars or trees, and we conjecture that the value of the game is $r/\bar{\mu}$ for any tree network with diameter D such that $D \le r \le \bar{\mu}$, where $\bar{\mu}$ is the length of a CPT of the network.

References

1. Alpern, S., Morton, A., Papadaki, K.: Patrolling Games, Oper. Res. 59(5), 1246–1257 (2011)
2. Basilico, N., Gatti, N., Amigoni, F.: Patrolling security games: Definition and algorithms for solving large instances with single patroller and single intruder, Artif. Intell., 184:78–123 (2012)
3. Edmonds J., Johnson E.L.: Matching, Euler tours and the Chinese postman. Math. Program. 5(1), 88–124 (1973)
4. Lin, K.Y., Atkinson, M.P., Chung, T.H., Glazebrook, K.D.: A graph patrol problem with random attack times, Oper. Res. 61(3):694–710 (2013)
5. Papadaki, K., Alpern, S., Lidbetter, T., Morton, A., Patrolling a Border: Oper. Res. (in press) (2016)

6. Pita, J., Jain, M., Marecki, J., Ordóñez, F., Portway, C., Tambe, M., Western, C., Paruchuri, P., Kraus, S.: Deployed ARMOR protection: the application of a game theoretic model for security at the Los Angeles International Airport. In: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: industrial track, pp.125–132. International Foundation for Autonomous Agents and Multiagent Systems (2008)
7. Yang, R., Ford, B., Tambe, M., Lemieux, A.: Adaptive resource allocation for wildlife protection against illegal poachers. In: Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems, pp.453-460. International Foundation for Autonomous Agents and Multiagent Systems (2014)

Appendix: Omitted proofs

Proof of Lemma 1 The attack must be taking place during the time interval $[0, r]$. Let A be the set of points that the Patroller intercepts in this time interval. Then clearly A must have length no greater than r and so the probability the attack takes place at a point in A is r/μ. It follows that $P \leq r/\mu$. □

Proof of Lemma 2 Suppose the attack starts at time t_0 at some point $x \in Q$. Then the attack is certainly intercepted if $\bar{w}$ is at x at time t_0. Let $t_x \in [0, \bar{\mu}]$ be such that $w(t_0 + t_x) = x$, so that the attack is intercepted by $\bar{w}$ if $T = t_x$. Let A be the set of times $t \in [0, \bar{\mu}]$ such that $t_x - r \leq t \leq t_x$ or $t \geq t_x + \bar{\mu} - r$, so if $T \in A$, then the attack is intercepted by $\bar{w}$. But the measure of A is r, so the probability that T is in A is $r/\bar{\mu}$ and hence $P \geq r/\bar{\mu}$. □

Proof of Lemma 3 Take a Patroller path w. We can assume that w starts at an endpoint, otherwise it is weakly dominated by a strategy that does. To see this, suppose the Patroller starts at an interior point before traveling directly to an endpoint, arriving there at time $t < 1$. Now consider the Patroller strategy that is the same but in the time interval $[0, t]$ the Patroller remains at the endpoint. Then clearly the second strategy intercepts the same set of attacks as the first one. Without loss of generalization we assume w starts at $x = 0$.

We only need to consider the path in the time interval $[0, 1 + r]$, after which time the attack has been completed with probability 1. Since $r < 2$ the walk cannot go between the two ends more than twice, so there are three possibilities.

The first is that w stays at $x = 0$ for the whole time, in which case the probability the attack is intercepted is $P = 1/2 \leq r/2$.

The second possibility is that w stays at $x = 0$ for time t_1, then goes to $x = 1$ and stays there for time t_2. We can assume it takes the Patroller time 1 to go between the endpoints since any path taking longer than that would be dominated, so $t_1 + t_2 = r$. The attack is intercepted at $x = 0$ if it starts sometime during $[0, t_1]$, which has probability $(1/2)t_1$. It is intercepted at $x = 1$ if it ends sometimes during $[1 + r - t_2, 1 - r]$, which has probability $(1/2)t_2$. Hence $P = (1/2)(t_1 + t_2) = r/2$.

The final possibility is that w stays at $x = 0$ for time t_1, then goes directly to $x = 1$ for time t_2, then goes directly back to $x = 0$ for time t_3, in which case we must have $t_1 + t_2 + t_3 = r - 1$. This time the attack is intercepted at $x = 0$ in the case of either of the two mutually exclusive events that it starts in $[0, t_1]$ or ends in $[1 + r - t_3, 1 - r]$, which have total probability $(1/2)(t_1 + t_3)$. If the attack takes place at $x = 1$, it must be taking place during the whole of the time interval $[1, r]$. But w must reach $x = 1$ sometime during this time interval, since it must have time to travel from $x = 0$ to $x = 1$ and back again, and hence intercepts the attack with probability 1. So the overall probability the attack is intercepted is $(1/2)(t_1 + t_3) + 1/2 \leq (1/2)(t_1 + t_2 + t_3) + 1/2 = r/2$. □

Proof of Theorem 3 We already showed that $r/(1 + r)$ is a lower bound for the value and now we show that it is also an upper bound. Now, suppose that the Attacker plays the r-attack strategy. The Patroller could:

1. Stay at any attack point but will not win with probability greater than $\frac{r}{1+r}$.
2. Travel between consecutive external attacks and if possible try to reach the internal attack: Suppose the Patroller is at point 0 up to time t: If $t \in [0, r]$ and then leaves for point r, she will reach point r at times in the range $[r, 2r]$. This gives total interception probability $f(t) + f(t + r) = \frac{t}{1+r} + \frac{2r-(t+r)}{1+r} = \frac{r}{1+r}$. Note that if the Patroller continues to the next attack along the unit interval, if it is the internal attack she will reach it at times greater than $r + \frac{r+\rho}{2} = 2r - \frac{r-\rho}{2}$, when the internal attack has been completed, and if it is an external attack she will reach it at time greater than $2r$, where all external attacks have been completed. If $t \in [r, 2r]$ then all attacks at point 0 have been intercepted but the Patroller arrives at point r after all attacks have been completed, which gives interception probability of $\frac{r}{1+r}$.
3. Travel between last two external attacks, crossing internal attack in the middle (this is the same as doing a roundtrip from one of the last external attacks to the internal attack and back): Suppose the Patroller leaves point $(q - 1)r$ at time t, toward the internal attack point and the last external attack point 1: If $t \in [0, r - \rho]$, she will reach the internal attack point at times $\left[\frac{r+\rho}{2}, r - \rho + \frac{r+\rho}{2}\right] = \left[\frac{r+\rho}{2}, 2r - \frac{r+\rho}{s}\right]$, and she will reach the external attack at point 1 at times $[r + \rho, 2r]$. This sums to a probability of $f(t) + g\left(t + \frac{r+\rho}{2}\right) + f(t+r+\rho) = \frac{t}{1+r} + \frac{\rho}{1+r} + \frac{2r-(t+r+\rho)}{1+r} = \frac{r}{1+r}$. If $t \in [r - \rho, r]$, she will reach the internal attack point at times $\left[2r - \frac{r+\rho}{2}, 2r - \frac{r-\rho}{2}\right]$, and the external attack point 1 at times greater than $2r$. This sums to a probability of $f(t) + g\left(t + \frac{r+\rho}{2}\right) = \frac{t}{1+r} + \frac{2r-\frac{r-\rho}{2}-\left(t+\frac{r+\rho}{2}\right)}{1+r} = \frac{r}{1+r}$. Finally, if $t \in [r, 2r]$, the Patroller will intercept all attacks at point $(q - 1)r$ and will not make it in time for the internal attack nor the attack at point 1, this gives the desired probability.

□