

Continuous Patrolling Games

Steve Alpern* Thomas Lidbetter† Katerina Papadak‡

August 13, 2020

Abstract

The continuous patrolling game studied here was first proposed in Alpern *et al.* (2011), which studied a discrete time game where facilities to be protected were modeled as the nodes of a graph. Here we consider protecting roads or pipelines, modeled as the arcs of a continuous network Q . The Attacker chooses a point of Q to attack during a chosen time interval of fixed duration (the *attack time*, α). The Patroller chooses a unit speed path on Q and intercepts the attack (and wins) if she visits the attacked point during the attack time interval. Solutions to the game have previously been given in certain special cases. Here, we analyze the game on arbitrary networks. Our results include the following: (i) a solution to the game for any network Q , as long as α is sufficiently short, generalizing the known solutions for circle or Eulerian networks and the network with two nodes joined by three arcs; (ii) a solution to the game for all tree networks that satisfy a condition on their extremities. We present a conjecture on the solution of the game for arbitrary trees and establish it in certain cases.

* Warwick Business School, University of Warwick, Coventry CV4 7AL, UK, steve.alpern@wbs.ac.uk

†Rutgers Business School, Newark, NJ 07102, USA, tlidbetter@business.rutgers.edu

‡Department of Mathematics, London School of Economics, London WC2A 2AE, UK,
k.p.papadaki@lse.ac.uk

1 Introduction

Patrolling games were introduced at the end of Alpern *et al.* (2011) to model the operational problem of how to optimally schedule patrols to intercept a terrorist attack, theft or infiltration. That paper, contrasting with earlier adversarial patrolling (Stackelberg) versions, modeled the problem as a zero-sum game between an Attacker and a Patroller, who wish to respectively maximize and minimize the probability of a successful attack. The domain on which the game was played out was taken to be a graph, with attacks restricted to the nodes and taking a given integer number of periods. A patrol is a walk on the graph, and intercepts the attack if it visits the attacked node during the attack period. This could model a guard in an art museum who enters a room while a thief is in the midst of removing a valuable painting from the wall. That paper was able to make some key observations about their game, giving bounds on the value, but was unable to find the value precisely or give optimal strategies except in some very limited cases. Papadaki *et al.* (2016) solved the game for line graphs, but the solution was very complicated even for this apparently simple graph. In the Conclusion section of the original paper Alpern *et al.* (2011), an extension of the problem to continuous space and time was suggested:

“It may be natural to consider a continuous-time formulation of this problem. An attack takes place at any point of the network (not necessarily a node) on a continuous time interval of fixed length. The Patroller uses a unit speed path and wins if she is at the attacked point at some time during the attack interval. This would model, for example, the defense of a pipeline system, and would resemble to a greater extent the classical search game problem.” [p1256]

The purpose of this paper is to carry out the suggestion in the quote above for an arbitrary network. We allow attacks that have a prescribed duration α to occur at any point of a continuous network Q . A unit speed patrol on Q is said to intercept the attack (and win for the Patroller) if it arrives at the attacked point at some time during the attack. The value of the game is the probability of interception, with best play on both sides. We find that optimal play for the Attacker typically involves mixing pure attacks that take place at different times.

After this type of continuous game was first proposed in 2011, it has been solved for some special networks. The circle network (or any Eulerian network) is easy to solve: a periodic traversal of the Eulerian tour, starting at a random point, is optimal for the Patroller;

attacking starting at a fixed time at a uniformly random location is optimal for the Attacker (see Alpern *et al.* (2016) and Garrec (2019)). The line segment network was solved in Alpern *et al.* (2016). In Garrec (2019) a solution for some values of α is given for the network with two nodes connected by three unit length arcs, and a complete formulation of the general game is given, including a proof of the existence of the value. The present paper extends to some extent all three of these prior results to general classes of networks: Eulerian networks to networks without leaf arcs; the line segment network to trees; the three-arc network to networks with large girth - for small attack times.

Our main results and chapter organization are as follows. Section 3 presents several (mixed) strategies for the players that can be used or adapted to obtain solutions of the game for various classes of networks in later sections. We note that Eulerian networks have no leaves, and Section 4 generalizes the solution of the former to networks without leaves. In particular, as long as the attack time is sufficiently short, we show that the attack strategy that chooses a point uniformly at random is still optimal; an optimal strategy for the Patroller is to follow a double cover tour of the network which never traverses an arc consecutively in opposite directions (as described in Theorem 4). We also give a new algorithm for constructing such a tour in Theorem 3. In Section 5 we allow the network to have leaves, and modify the optimal strategies of the previous section to generate optimal strategies for arbitrary networks, as long as the attack time is sufficiently short (see Theorem 6).

Section 6 considers trees and in particular those that satisfy a condition we call the Leaf Condition. We give a precise definition of the condition, which requires some delicacy (Definition 8). In fact, any tree satisfies the Leaf Condition as long as the attack time is sufficiently short. Star networks (trees with only leaf arcs) also satisfy the Leaf Condition for sufficiently large attack times, and the only stars that do not satisfy the Leaf Condition are those that have an arc that is longer than half the total length of the network. In Theorem 8 we solve the game for all trees in the case that the Leaf Condition holds, giving a simple expression for the value of the game in terms of the length of the network, the attack time and another parameter. In Section 7 Conjecture 1 states that this expression is always equal to the value of the game on trees. We establish the conjecture for some stars that do not satisfy the Leaf Condition.

2 Literature Review

In addition to the papers discussed in the Introduction, which were the most relevant to continuous patrolling, there is a more extensive literature on adversarial patrolling. The problem of patrolling a perimeter has been analyzed by Zoroa *et al.* (2012) (where the attack location can move to adjacent locations) and Lin (2019), the latter in a continuous time context. Extensions of Alpern *et al.* (2016) where the costs of successful attacks are time and node dependent have been studied by Lin *et al.* (2013) (for random attack times), Lin *et al.* (2014) (with imperfect detection) and Yolmeh and Baykal-Gürsoy (2019) (which includes an application to an urban rail network).

Stackelberg approaches with the Patroller as first mover, have been pioneered in an artificial intelligence context by Basilico *et al.* (2012) (which includes an algorithm for large cases) and Basilico *et al.* (2017) (where the optimal strategy in certain cases is for the Patroller to stay in place until the sensor reveals an attack an unknown location).

More applied approaches to patrolling are of practical importance. Applications to scheduling randomized security checks and canine patrols at Los Angeles Airport have been developed and deployed in Pita *et al.* (2008). The United States Coast Guard also uses a game-theoretic system to schedule patrols in the Port of Boston (An *et al.* 2013) Recently, a game theoretic approach to schedule patrols to guard against poachers has been explored in Fang *et al.* (2016) (where the novel algorithm PAWS was introduced) and Xu *et al.* (2019) (where the success of deploying PAWS in the field is described) Patrolling to detect radiation and consequently nuclear threats was modeled in the novel paper of Hochbaum *et al.* (2014).

The possibility that the Attacker could know when the Patroller is nearby (perhaps at the same node), raised in Alpern *et al.* (2011), has recently been studied in Alpern and Katsikas (2019) and Lin (2019) in different contexts. In the former this knowledge helped the Attacker, in the latter, it did not. Multiple patrollers have been considered in the robotics and computer science literatures, where an important paper with a similar network structure to ours is Czyzowicz *et al.* (2017). A connection between patrols and inspection games is made in Baston and Bostock (1991) and between patrols and hide-seek games in Garrec (2019). Restricting the Patroller to periodic paths creates difficulties analyzed in Alpern *et al.* (2018).

3 Formal Definitions for Network and Game

In this section we define the continuous patrolling game and present definitions related to the connected network Q on which it is played. For Q , standard graph theoretic definitions must be modified for a network which is considered as a metric space and a measure space, not simply a combinatorial object.

To define Q , we begin with the usual combinatorial structure of arcs and nodes, with the addition of a *length* $\lambda(a)$ assigned to each arc. We can then identify an arc a with an open interval of length $\lambda(a)$, endowed with Lebesgue measure and Euclidean distance d , and consider λ as a measure on Q , called *length*. The total length of Q is denoted by $\mu = \lambda(Q)$. The topology on the interval arcs gives a topology on their union Q . A *path* in Q is a continuous function from a closed interval to Q . We take the metric $d(x, y)$ on Q as the minimum length of a path between x and y . A point x of Q is called a *regular* point if it has a neighborhood homeomorphic to an open interval. Regular points are the interior points of arcs. The remaining points of Q are the nodes, which are the boundary points of the arcs. The *degree* of a node y can be defined either combinatorially or topologically. Combinatorially, it is the number of adjacent nodes; topologically, it is the number of components of a neighborhood of y from which y has been deleted. Note that this means that strictly speaking, nodes of degree two are not permitted. A node of degree 1 is called a *leaf node*, and its adjacent arc is called a *leaf arc*. To ensure that every leaf arc has a single leaf node in its closure, we exclude the line segment network from consideration. In any case the continuous patrolling game has been solved for the line segment in Alpern *et al.* (2016).

A *circuit* in Q is a closed path (that is, with the same startpoint and endpoint) consisting of distinct adjacent arcs. A *tour* of Q is a closed path visiting all points of Q , and a tour of minimum length is called a *Chinese Postman Tour (CPT)*. The length of this path is denoted $\bar{\mu}$. It was shown by Edmonds and Johnson (1973) that a CPT can be found in polynomial time, with respect to the number of nodes. A closed path which is a circuit and a tour is called an Eulerian tour. As is well known, a connected network has an Eulerian tour if and only if it is Eulerian, defined as having nodes all of even degree. If we double every arc of a network Q , the resulting network is Eulerian with length 2μ , so Q has a tour of length 2μ and hence $\bar{\mu} \leq 2\mu$.

The continuous patrolling game is played on Q as follows. The Attacker chooses a point x in Q to attack, and a closed time interval J of given length α during which to attack it.

Since α is fixed, the *attack interval* $J = [\tau, \tau + \alpha]$ is determined by its starting time τ . The game and its value are determined by the pair (Q, α) . The Patroller chooses a path $S(t)$, where $t \geq 0$, which we call a *patrol*, satisfying

$$d(S(t), S(t')) \leq t|t'|, \text{ for all } t, t' \geq 0. \quad (1)$$

For simplicity, we shall call a path satisfying the 1-Lipshitz condition (1) a *unit speed path*. We don't specify an upper bound on the starting time of the attack, but in every case we have studied there is an optimal mixed attack strategy in which all its (pure strategy) attacks are over by time 4μ . A patrol is said to *intercept* an attack if it visits the attacked point while it is being attacked. The game is very simply defined: the maximizing Patroller wins (payoff $P = 1$) if her patrol intercepts the attack. Otherwise, the Attacker wins (payoff $P = 0$ to the Patroller). The payoffs to the Attacker are reversed, so the game has constant sum 1. To put this more concisely, if the patrol is S and the attack is at point x during the interval $J = [\tau, \tau + \alpha]$, then the payoff P to the maximizing Patroller is given by

$$P(S, (x, J)) = \begin{cases} 1 & \text{if } x \in S(J), \\ 0 & \text{otherwise.} \end{cases}$$

For mixed strategies, the expected payoff can be interpreted as the probability that the attack is intercepted. The value of the game, denoted V , is the interception probability, with best play on both sides.

Garrec (2019) used the fact that P is lower semicontinuous to establish the existence of a value V for this infinite game. We note that if $\alpha = 0$ then the Attacker can win almost surely by attacking uniformly on Q (according to λ) at a fixed time; if $\alpha \geq \bar{\mu}$, the Patroller can ensure a win by adopting a Chinese Postman Tour, starting anywhere at time 0 and repeating the tour with period $\bar{\mu}$. So to avoid the trivial cases where one of the player can always win, we assume $0 < \alpha < \bar{\mu}$.

Throughout the paper the complement $Q - Y$ of a set Y is denoted by Y^c .

3.1 The Uniform and the Independent Attack Strategies

Some networks, as we shall see in later sections, require Attacker strategies specifically suited to their structure, such as attacks on leaf nodes when the network is a tree. But there are also some general strategies that are available on any network. Here we define two of these and present the general bounds on the value that they give.

Definition 1 (Uniform attack strategy) A *uniform attack strategy* is a mixture of pure attacks that have a common attack time interval $J = [M, M + \alpha]$. The attacked point is chosen uniformly at random. That is, the probability that the attacked point lies in a set Y is given by $\lambda(Y) / \mu$.

We restate a lemma from Alpern et al. (2016) for completeness (the proof is in the Online Appendix).

Lemma 1 Against any patrol S , a uniform attack strategy is intercepted with probability not more than α/μ . Consequently $V \leq \alpha/\mu$ for any network.

We now define independence for sets and strategies.

Definition 2 (Independent set) A subset I of Q is called *independent* if the distance between any two of its points is at least α . For any subset Y of Q , the set $W \equiv W(Y)$ is the subset of Q consisting of all points at distance at most $\alpha/2$ from Y .

Definition 3 (Independent attack strategy) Given an independent set I of cardinality $|I|$ and the set $W \equiv W(I)$, the *independent attack strategy* is as follows for $p = \frac{|I|\alpha}{\lambda(W^c) + |I|\alpha}$.

1. With probability p attack at an element of I chosen equiprobably at a start time chosen uniformly at random in $J = [\emptyset, \alpha]$.
2. With probability $1 - p$ attack uniformly on W^c at start time $\alpha/2$.

The independent attack strategy randomizes over both time and space, unlike the strategy of the same name defined in Alpern et al. (2011) for the discrete patrolling game, which randomizes only over space. The following result gives an upper bound on the strategy's interception probability.

Theorem 1 Let I be an independent subset of Q of cardinality $|I|$. Then

$$V \leq \frac{\alpha}{\lambda(W^c) + |I|\alpha},$$

which the Attacker can ensure by adopting the independent attack strategy $\lambda(W^c) = 0$ we have $V \leq 1/|I|$. Furthermore, if I is the set of leaf nodes, and leaf arcs have lengths exceeding $\alpha/2$, then

$$V \leq \frac{\alpha}{\mu + |I|\alpha/2}.$$

Proof: Let S denote any patrol and suppose the independent attack strategy is adopted. If S remains in W during J , it intercepts the attack with probability at most p/l , where l is the cardinality of J . Similarly, since S has unit speed, if it remains in W^c during time J , it intercepts an attack with probability at most $(1-p)(\alpha/\lambda(W^c))$. The chosen value of p is the one that makes these probabilities both equal to $\alpha/\lambda(W^c) + l\alpha$.

Finally, suppose the patrol S starts in W^c at time 0, reaches a point $x \in l$ at some time t , $\alpha \leq t \leq 2\alpha$, early enough to intercept some attacks on l and late enough to intercept some attacks on W^c . Since the latest such a patrol can leave W^c at time $t - \alpha/2$, it can cover a set of length at most $(t - \alpha/2) - \alpha/2 = t - \alpha$ in W^c after the attacks at time $\alpha/2$, intercepting a fraction $(t - \alpha)/\lambda(W^c)$ of the attacks there. In addition, the patrol can intercept the attacks at x starting between $t - \alpha$ and α , so a fraction $(2\alpha - t)/l\alpha$ of the attacks at x , or $(2\alpha - t)/l\alpha$ of the attacks on l . Thus the maximum probability that a patrol arriving at l at time t can intercept an attack is given by

$$(1-p)\frac{t-\alpha}{\lambda(W^c)} + p\frac{2\alpha-t}{l\alpha} = \frac{\alpha}{\lambda(W^c) + l\alpha}.$$

By time symmetry, the same bound holds if the patrol starts at a point of l and ends up in W^c .

If $\lambda(W^c) = 0$ we have $V \leq 1/l$ trivially.

To prove the last assertion note that if l is the set of leaf nodes, and leaf arcs have lengths exceeding $\alpha/2$, then leaf nodes form an independent set l and $\lambda(l) = \alpha/2$. \square

3.2 A General Strategy Available to the Patroller

Some patrol strategies come from finding closed paths on the network with specific properties, and then have the Patroller go around them periodically starting at a random point. Normally the closed path will be a tour, but we give a more general definition in case it is not.

Definition 4 (Randomized periodic extension) If $S : [0, L] \rightarrow Q$ is a closed unit speed path, we can extend it to various patrols $S_\Delta : [0, \infty) \rightarrow Q$ of period L by the definition

$$S_\Delta(t) = S((t + \Delta) \bmod L), \text{ for all } t \geq 0.$$

Thus S_Δ is a periodic patrol that starts at the point $S(\Delta)$ at time 0. The **randomized periodic extension** \tilde{S} of S is defined as the random mixture of the pure patrols S_Δ , with Δ chosen uniformly in the interval (or circle) $[0, L]$. In the special case that S is a Chinese Postman Tour, with $L = \bar{\mu}$, we call \tilde{S} a Chinese Postman Tour strategy.

3.3 A Theorem on k -covering Tours

If a network Q has an Eulerian tour, its randomized periodic extension makes an effective patrolling strategy, because it visits all regular points equally often (once), so the Attacker is indifferent as to where to attack. If there is no Eulerian tour (the general case), we can still use this idea, if there is a tour which visits all regular points equally often. In Theorem 3 and Lemma 6, we will show that there is indeed such a tour which visits all regular points twice (a 2-cover), with some additional properties. This idea is formalized in the following.

Theorem 2 *Let $S: [0, L] \rightarrow Q$ be a closed unit speed tour that visits every point of Q at k times which are separated by at least $\alpha \pmod L$. Let \tilde{S} be the randomized periodic extension of S (from Definition 4). Then we have*

- (i) *\tilde{S} intercepts any attack with probability at least $k\alpha/L$.*
- (ii) *If $L = k\mu$, then the randomized periodic extension \tilde{S} (for the Patroller) and a uniform attack strategy (for the Attacker) are optimal and the value of the game is given by α/μ .*

Proof: For part (i), suppose the attack takes place at a point x in Q starting at some time τ . Let $t_i, i = 1, \dots, k$ be times separated by at least α , such that $S(t_i) = x$. The attack will be intercepted by S_Δ if Δ is in the set $Y = \bigcup [t_i - \tau - \alpha, t_i - \tau]$ (modulo L), since in this case the Patroller will visit $x = S(t_i)$ at some time in $[\tau, \tau + \alpha]$. The separation assumption ensures that these intervals are disjoint, and since they all have length α , the length (Lebesgue measure) of Y is given by $|Y| = k\alpha$. By the definition of \tilde{S} , the probability that $\Delta \in Y$ is equal to $|Y|/L = k\alpha/L$, as claimed in (i), so we have $V \geq k\alpha/L = k\alpha/k\mu = \alpha/\mu$ under the assumption of part (ii). By Lemma 1, we also have that $V \leq \alpha/\mu$, so the two inequalities give $V = \alpha/\mu$, with \tilde{S} and the uniform attack strategy optimal. \square

As suggested above in the introductory remarks of this subsection, taking $k=1$ in Theorem 2 gives another proof of the following elementary result of Alpern *et al.* (2016) and Garrec (2019).

Corollary 1 *If Q is Eulerian, with Eulerian tour S , then for $\alpha \leq \mu$ we have $V = \alpha/\mu$. ($V = 1$ if $\alpha \geq \mu$.) In this case the randomized periodic extension \tilde{S} and the uniform attack strategy are optimal for the Patroller and Attacker, respectively. Furthermore, for a Chinese Postman Tour S of any network Q , taking $k=1$ and $L = \bar{\mu}$ gives $V \geq \alpha/\bar{\mu}$.*

It is useful to note for applications to patrolling by m robots, that if in Theorem 2 we require that S visits every point at k times separated by time intervals $m\alpha$, then m Patrollers can intercept any attack with probability at least mka/L (or 1, if this is larger). To see this, pick Δ as above and let the path of the i 'th Patroller (robot) be defined by $S_i(t) = S(\Delta + i(L/m) + t)$. The arrival times at any point of Q are then separated by at least α . This shows that in our later lower bounds for V , these can be multiplied by the number of Patrollers, with an upper bound of 1.

3.4 Identifying Points of Q Helps the Patroller

We conclude this section with an observation on the effect of identifying points of Q on the value. Alpern *et al.* (2011) considered the effect of identifying two nodes of a graph. Here, we identify two *points* of the network Q , using the well known quotient topology. In Figure 1 we identify the arc midpoints C and D of the network Q to produce a new network Q' .

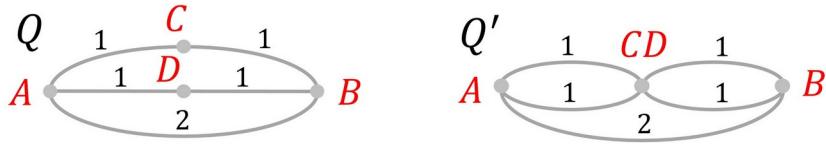


Figure 1: Identifying points C, D of Q to obtain Q' .

We may first look at two cases which have already been solved, the line segment $Q = [0, 1]$ and the circle $Q_{circle} = [0, 1] \bmod 1$ (which is obtained from the line segment by identifying the endpoints), with say $\alpha = 1/2$. From Alpern *et al.* (2016), we have $V(Q_{line}) = \alpha/(\mu + \alpha) = 1/3$. However as the circle is Eulerian, we have $V(Q_{circle}) = \alpha/\mu = 1/2$, which is larger. It is easy to show that identifying points cannot decrease the value. Of course if we further identify points on the circle, we get new points of degree 4, so the resulting Eulerian network retains the value of 12.

Lemma 2 Suppose Qd is the metric space obtained from Q , d by replacing the metric d with a smaller metric d' , that is, with $0 \leq d'(x, y) \leq d(x, y)$ for all $x, y \in Q = Q'$. Then $V(Q', d') \geq V(Q, d)$. Furthermore, if Q' is obtained from Q by decreasing the length of an arc or simply identifying two points x and y , the same result holds.

The proof of Lemma 2 is given in the Online Appendix. An application of it is given at the end of Section 4.

4 Networks Without Leaves

To extend Corollary 1 to general networks, we first note that Eulerian networks have no leaf arcs, so we attempt to find such a tour S satisfying the hypothesis of Theorem 2 for networks without leaf arcs. It turns out that taking $k = 2$ in Theorem 2 is high enough. We can find such a tour (see Theorem 4) if α is sufficiently small with respect to the *girth* g of Q , defined for networks as the minimum length of a circuit in Q , and if Q has no circuits then $g = \infty$ (For networks with unit length arcs, this coincides with the usual integer definition of the girth of a graph.) Our first result is the following.

Theorem 3 *For any network Q there is a tour S which covers every arc twice and for which no arc is traversed consecutively in opposite directions, except for leaf arcs.*

The way we will prove Theorem 3 is to double every arc of Q to create a network \hat{Q} . Then \hat{Q} is Eulerian and has an Eulerian tour. We note that in Euler's Theorem (finding an Eulerian tour in graphs of even degree), we can control to some extent the construction of the tour. The following refinement of Euler's Theorem (Lemma 3) is based on some simple modifications of the traditional proof and shows that we can control the pairing of entered and exited passages of the tour at every node. Formally, a passage at a node x is a pair (x, a) , where a is an arc incident to x . So a node of degree d has d passages and every arc is part of two passages.

Lemma 3 *Let Q be a connected Eulerian network such that at every node the passages are identified in pairs (they are “paired”). Then there is an Eulerian tour S of Q satisfying*

S never enters and leaves a node via paired passages. (*)

Proof: This proof mimics the usual proof of Euler's Theorem. We first construct a circuit C satisfying condition (*), which we call a $*$ -circuit, using the following rules:

1. Start at any node x and leave by any passage P (we let P' be the paired passage of P).
2. Always leave a node by an untraversed passage not paired with the arriving passage.
3. If, after arriving at a node, there are three untraversed passages with exactly two of them paired, leave by one of this pair.

4. If, after arriving at node x , there are two untraversed passages, leave by passage if P is untraversed.
5. If there are no remaining untraversed passages after arriving at a node, stop.

To simply obtain a circuit (not necessarily satisfying (*)) starting and ending at x , we would follow the usual method of simply leaving a node by any *untraversed passage*, a simpler form of Rule 2. The existence of an untraversed passage (at any node other than the starting node x) follows from the fact that after arriving at a node an odd number of passages will have been traversed, so an odd number (hence not 0) are untraversed. We show that the full form of Rule 2 along with the other rules ensure that we can always leave a node in a way that satisfies (*) whether the node is the initial node x or another node y .

We first check that after arriving at a node y other than the starting node x , there cannot be only one remaining untraversed passage which is paired with the arriving passage. Since every node has even degree and degree two nodes are not permitted, the node y must have been previously arrived at. After this previous arrival at y , there must have been three untraversed passages with exactly two of them paired. But Rule 3 ensures the circuit left by one of those two passages so after arriving by the other one on the final visit, the last untraversed passage must have a different label.

To check that the final arriving passage at the initial node x is not P (note that if P had not been traversed before the penultimate visit to x , Rule 4 ensures that it will be traversed on that visit, and it will not be the final arriving passage).

If C is a tour (contains all the arcs), we are done. Otherwise, since Q is connected, there is a node z with some passages in C and some not in C (see Figure 6). Suppose that C leaves z beginning via passage a and ends at z via passage b . We create a new $*$ -circuit starting at z , called C' , using the same rules and using only passages not in C . Suppose C' begins with a passage called d (which we can choose) and ends with a passage called e (which we cannot control). The combined circuit CC' which starts at z and traverses C and then C' will satisfy (*) except possibly for the transitions b, d and e, a between the two circuits, so we need $d \neq b$ and $e \neq a$ (this means d is not paired with b and e is not paired with a). The arc d is chosen as follows.

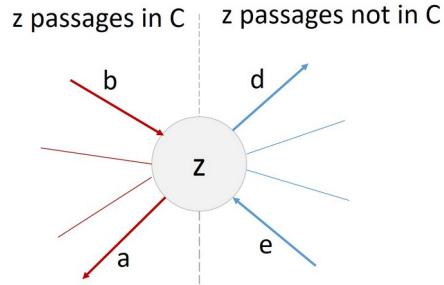


Figure 2: How to join two $*$ -circuits at node z .

1. If a' is not in C , take $d = a'$. This ensures that $d = a' \neq b$ since $a \neq b$. Also $e \neq d = a'$.
2. If a' is in C , take $d \neq b$. We know that also $e \neq a'$ because a is in C .

If the circuit CC' is not a tour, we iteratively continue to add new circuits until we end up with a tour, noting that the process is guaranteed to end since every new circuit contains at least one new arc and there are a finite number of arcs. \square

Now we are ready to prove Theorem 3.

Proof of Theorem 3. Let \hat{Q} be the Eulerian network obtained from Q by doubling every arc. (This has the effect of replacing leaf arcs with loops of double the length, since degree two nodes are not permitted.) At every node of \hat{Q} we pair passages that correspond to the same passage of Q . Now apply Lemma 3 to \hat{Q} to obtain an Eulerian circuit \hat{S} of \hat{Q} satisfying condition (*). This corresponds to S_2 , a *double cover* of Q (a tour of Q where every arc is traversed twice), in which consecutive arcs are distinct, except for leaf arcs. For loops, an arc may be repeated consecutively, but always in the same direction both times. \square

Theorem 3 is not new; it was proved by Sabidussi (1977). See also Klavzar and Rus (2013) and Eggleton and Skilton (1984). Our proof based on the new result, Lemma 3, is elementary.

The proof of Lemma 3 gives rise to an algorithm for constructing an Eulerian tour of \hat{Q} satisfying condition (*), and hence a tour of Q of the form described in the statement of Theorem 3 (named S_2). Indeed, by following the rules listed in the proof of Lemma 3, we obtain a circuit C in \hat{Q} satisfying (*); by recursively applying the rules to the connected components of $\hat{Q} - C$ and appending these circuits to C at appropriate points, we can obtain an Eulerian tour of \hat{Q} satisfying (*).

We illustrate the creation of the $*$ -circuit described above for the network K_4 depicted in Figure 3. Doubling each arc, we give the extra arc the same label as the original arc

but with a prime. Applying the rules of the proof of Lemma 3, starting at the bottom left node, we obtain a circuit: $a, b, c, d, e, \alpha, f, d'$. Removing this circuit leaves the network consisting of arcs b, e and f' , which is already a circuit. Adding this circuit at the first possible opportunity, we obtain the Eulerian tour $a, b, e, f', b, c, d, e, \alpha, f, d'$.

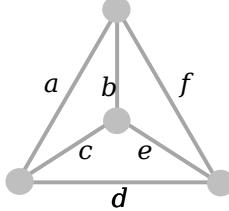


Figure 3: The network K_4 .

Theorem 4 Let Q be a network without leaf arcs. Then for $\alpha \leq g$, where g is the girth, we have the following:

1. The value of the game is $V = \alpha/\mu$.
2. For the Attacker, any uniform attack strategy is optimal.
3. For the Patroller, the randomized periodic extension \tilde{S}_2 is optimal, for any tour S_2 given by Theorem 3.

Proof: Let S_2 be a tour of Q given by Theorem 3. Note that it has length $L = 2\mu$. Since there are no leaf arcs, any two consecutive arcs of S_2 are distinct. Suppose some point x of Q is reached by S_2 at consecutive times t and s with $t < s$. Let Z denote the restriction of S_2 to the interval $[t, s]$. Then Z is a circuit of length $s - t$ and hence $s - t \geq g$, by the definition of girth. Hence $V = \alpha/\mu$, by Theorem 2(ii) with $k = 2$ and since $\alpha \leq g$. \square

For the network K_4 depicted in Figure 3, assuming all arcs have length 1, the girth g is 3. So for $\alpha \leq 3$, the uniform attack strategy is optimal and the Patroller strategy S_2 is optimal, where S_2 is the tour $a, b, e, f', b, c, d, e, \alpha, f, d'$.

As a further example, consider Q to be a network with two nodes A and B connected by three arcs of lengths $a \leq b \leq c$. Then $g = a + b$ and $\mu = a + b + c$, so we have by Theorem 4 that the value is $V(\alpha) = \alpha / (a + b + c)$ for $\alpha \leq a + b$. This network, with $a = b = c = 1$ (and hence $g = 2$), was studied by Garrec (2019), who found (among other results) that $V(\alpha) = \alpha/3$ for $\alpha \leq 2$ and $V(\alpha) \leq f(\alpha) \equiv 1 - (1/3)(2 - \alpha/2)^2$ for $\alpha \in [2, 103]$. Since $f(\alpha) < \alpha/3$ for $\alpha \in [2, 103]$ ($f(\alpha) = \alpha/3$ for $\alpha = 2$ and $f'(\alpha) = 4 - \alpha)/6 < 1/3$ for $\alpha > 2$), the Patroller cannot

obtain an interception probability of $\alpha/3$ for α in this interval, so the bound $\alpha \leq g = 2$ in Theorem 4 is tight.

However, it is useful to observe that we can sometimes improve on the upper bound. Suppose we have a network Q with two nodes connected by five arcs labeled as 1, 2, 3, 4, 5, with arc i having length i . The girth is given by $g = g(Q_5) = 1 + 2 = 3$. However, suppose we obtain a double cover (with $k=2$) S of Q described by the sequence $[2, 3, 4, 5, 1, 2, 3, 4, 5]$, where unprimed arcs go from, say, node A to node B and primed arcs go from node B to node A . The shortest return time to a regular point is for a point x near node B on the arc of length 5. After leaving x , going to nearby B , the patrol traverses arcs of lengths $1 + 2 + 3 + 4 = 10$ before going back to x from B . Note that S returns to A after gaps of 3, 7, 6, 5 and 9 so at two time points separated by 14 (at the start and after the gap of 6). Also B is visited twice separated by a gap of 14. So for the network Q we have $V = \alpha/\mu$ for $\alpha \leq 10$ rather than just for $\alpha \leq 3$. This observation leads to combinatorial questions about the maximum shortest circuit in a k -cover of a network Q . As noted above based on Garrec's analysis of the three arc network, in certain cases $V = \alpha/\mu$ fails for all $\alpha > g$. We note that our example Q generalizes easily to the following.

Theorem 5 *Let Q be a network with two nodes connected by an odd number of arcs. Then $V = \alpha/\mu$ for $\alpha \leq \mu - L$, where L is the length of the longest arc.*

If Q is a network with two nodes connected by an even number of arcs, then Q is Eulerian and thus $V = \alpha/\mu$ for all α .

We conclude this section with an application of our earlier result on identifying points. Consider the two networks Q and Q' drawn in Figure 1, with $\alpha = 3$. We would like to show that $V(Q') = \alpha/\mu = 3/6 = 1/2$. We know from Lemma 2 that $V(Q') \leq \alpha/\mu = 1/2$. So we only need $1/2$ as a lower bound on $V(Q')$. However we cannot apply Theorem 4 because it is not true that α is less than or equal to the girth of Q' , which is 2. However we know either from Garrec (2019) or from Theorem 4 (which applies because $3 = \alpha < g = 4$) that $V(Q) = \alpha/\mu = 1/2$. So by viewing Q' as coming from Q by identifying points C and D , Lemma 2 gives $V(Q') \geq V(Q) = 1/2$.

5 Brief Attacks on Arbitrary Networks

We now extend Theorem 4 to networks with leaves. We begin with a modified Patroller strategy based on the tour \mathcal{S} of Theorem 3.

Definition 5 Let S be a tour given by Theorem 3. We denote by S^α the tour that follows the same trajectory as S but stops for time α whenever it reaches a leaf node.

Lemma 4 Let Q be a network with l leaf nodes and girth g . Then

$$V \geq \frac{\alpha}{\mu + l\alpha/2}, \text{ for } \alpha \leq g.$$

Proof: Tour S_2^α takes total time $2\mu + l\alpha$. Note that every point of Q is visited by S_2^α at two times differing by at least α . So by Theorem 2 part (i) with $k = 2$, $L = 2\mu + l\alpha$, we have $V \geq 2\alpha / (2\mu + l\alpha)$. (We observe that instead of stopping for time α , the tour S_2^α could do anything in this time interval, such as going away from the node a distance $\alpha/2$ and returning.) \square

Definition 6 (Generalized girth) We define the **generalized girth** g^* of a network Q by considering a leaf arc of length L to be a circuit of length $2L$. So g^* is the minimum of circuit lengths of Q and twice the length of any leaf arc.

In particular $g^* \leq g$, with equality if there are no leaf arcs or if all leaf arcs have length greater than $g/2$. Note that if $\alpha \leq g^*$ we know in particular that all leaf arcs have length at least $\alpha/2$ and hence Theorem 1 applies. Thus we have the following Attacker estimate.

Lemma 5 Let Q be a network with l leaf nodes and generalized girth g^* . Then by adopting the independent attack strategy on the set I of leaf nodes, the Attacker can ensure that the interception probability is less than $\frac{\alpha}{\mu + l\alpha/2}$ for $\alpha \leq g^*$. Hence,

$$V \leq \frac{\alpha}{\mu + l\alpha/2}, \text{ for } \alpha \leq g^*.$$

Proof: As noted above, the assumption on α ensures that all leaf arcs have length at least $\alpha/2$, so the result follows from Theorem 1. \square

Since $g^* \leq g$, Lemmas 4 and 5 apply when $\alpha \leq g^*$ and hence we have the following extension of Theorem 4 to networks with leaf arcs.

Theorem 6 Let Q be a network with l leaf nodes and generalized girth g^* . Then

$$V = \frac{\alpha}{\mu + l\alpha/2}, \text{ for } \alpha \leq g^*.$$

For the Patroller, an optimal strategy is S_2^α as defined above. For the Attacker, an optimal strategy is the independent attack strategy, taking I to be the independent set of leaf nodes.

It is useful for later comparisons to specialize this result to trees.

Corollary 2 *If Q is a tree with l leaf arcs, then*

$$(i) V \geq \frac{\alpha}{\mu + l\alpha/2},$$

(ii) *with equality if all leaf arcs have length at least $p\alpha$.*

Proof: To establish (ii), note that trees have no circuits, so the generalized girth g^* is twice the length of its smallest leaf arc, so by assumption, $\alpha \leq g^*$. The result now follows from Theorem 6. For (i), consider the patrol S_2^α . Note that between any two visits by S_2^α to a point of Q , a leaf node is visited. Hence the return times exceed the time α that S_2^α stops at that node, and the result follows from Theorem 2(i) with $k=2$ and $L=2\mu + l\alpha$. \square

For example, consider the tree Q depicted in Figure 4. The number of leaf arcs is $l=5$, the generalized girth is $g^*=2$ and total length is $\mu=9$, so by Theorem 6, the value of the game is $d(9 + 5\alpha/2)$ for $\alpha \leq 2$.

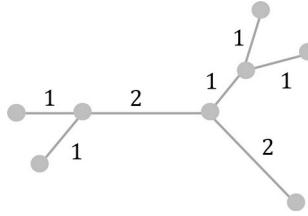


Figure 4: The tree Q .

6 Solving the Game for Trees

In Corollary 2 we gave some preliminary results for trees. Lemma 4 gave a lower bound on the value of the game based on the Patroller strategy S_2^α . Furthermore, for $\alpha \leq g^*$, where g^* is the generalized girth, we showed in Theorem 6 that the independent attack strategy ensures that this lower bound is tight. Note that for a tree, g^* is twice the length of the shortest leaf arc. In this section, we extend these results and give optimal Patroller and Attacker strategies for some values of α which are greater than g^* . We start by defining the *extremity set* E , a subset of Q that is essential in describing optimal Patroller and Attacker strategies.

6.1 The Extremity Set E

The relationship between the network Q and the duration α of the attack interval determines the type of optimal player strategies. In this section we define the extremity set E that helps us explore this relationship for trees.

If B is a set of points then we denote by \bar{B} the topological closure of B . If Q is a tree network, then its minimum tour time is 2μ , as every arc must be traversed twice. If x is a regular point of tree network Q , then $Q - \{x\}$ has two connected components $Q = Q_1(x)$ and $Q_2 = Q_2(x)$, whose lengths satisfy $\lambda(Q_1) + \lambda(Q_2) = \lambda(Q) = \mu$. We introduce a subset $E(Q)$ of Q .

Definition 7 (The extremity set E) Let Q be a tree. The extremity set $E \equiv E(Q, \alpha)$ is defined as the set of all regular points $x \in Q$ such that

$$\min_{i=1,2} \lambda(Q_i(x)) < \alpha/2. \quad (2)$$

Note that $\min_{i=1,2} \lambda(Q_i) \leq \mu/2$ and if additionally $\mu < \alpha$ then (2) holds for all regular points, which implies that $\bar{E} = Q$. The extremity set E consists of regular points whose minimum return time during a CPT is less than the attack duration α . It can be partitioned into maximal connected sets that we call *components* of E and we denote by E

Example 1 We illustrate the extremity set E on the tree network of Figure 4 that has $\mu = 9$. Figure 5 shows how E changes for increasing values of α on this network. As α increases the components grow starting from points near the five leaf nodes of the tree. Initially there are five components (cases $\alpha = 1, 2, 3, 4$), but eventually points near non-leaf nodes become members of E and the number of components increase to seven (cases $\alpha = 5, 6, 7, 8$). Note that in case $\alpha = 8$ the closure of E is equal to the whole network. The results from the previous sections (Theorem 6, Corollary 2) solve the game for cases $\alpha \leq g^* = 2$, but in this section we extend the results to cover all cases of $\alpha \leq 4$.

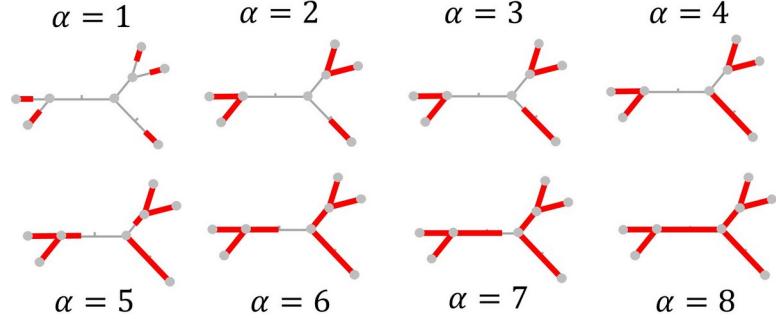


Figure 5: The extremity set $E(Q, \alpha)$, shown in thick red lines, for the tree Q of Figure 4 and $\alpha = 1, \dots, 8$.

Example 2 Figure 6 depicts a star network. We assume that the extremity set E is as it appears on the figure in red thick lines; we make no assumptions on the value of α or the length of the arcs; from the shape of E we draw some conclusions. Here, E decomposes into four components: (A, B) , (A, C) , (D, F) , (E, G) . We claim that $\lambda(DF) = \lambda(EG) = \alpha/2$; this is because on leaf arc AD (similarly for AE) if $\lambda(DF) < \alpha/2$ there would be a point X on the right of F whose distance from D would be $< \alpha/2$, implying $\lambda(DX) < \alpha/2$ and thus contradicting $X \notin E$. Similarly, if $\lambda(DF) > \alpha/2$ there would be a point X on the left of F where $\lambda(DX) > \alpha/2$ contradicting $X \in E$. Thus, components E_j that are strict subsets of a leaf arc and whose closure contains the leaf node will have length $\leq \alpha/2$. However, components E_j whose closure is the entire leaf arc (like AB and AC) must have length $\leq \alpha/2$; if they had length $> \alpha/2$ then there would be point X on the component AB near node A where $\lambda(BX) > \alpha/2$ contradicting $X \in E$.



Figure 6: A tree, with its extremity set E in thick red.

6.2 The E -patrolling Strategy \mathcal{S} for Trees

We will see that for some trees, the uniform CPT strategy is still optimal for the Patroller, but its optimality depends on the size of the attack duration, α . As mentioned earlier, for a tree a CPT is simply any depth-first search which returns to its start point after completing

its search, so that $\bar{\mu} = 2\mu$; every point of the tree except the leaf nodes is visited at least twice by a CPT. This means the leaf nodes and regular points near them are left “less protected” by a uniform CPT than the other points, and for sufficiently small values of α , there will be points in the tree whose two closest visit times (modulo $\bar{\mu}$) are at least time α apart, meaning that they are, in a sense “twice as protected” as the leaf nodes (in all that follows, arithmetic on time will be performed modulo the length of the tour in question).

This motivates the introduction of a new Patroller strategy \mathcal{S} for trees that we call the E -patrolling strategy. To describe it, we use the set extremity set $\bar{E} \equiv E(Q, \alpha)$ that we defined earlier; in particular, we use the closure of \bar{E} of E and its components $\bar{E}^1, \dots, \bar{E}^k$, each of which is a subtree of Q . We have $\lambda(\bar{E}) = \lambda(E)$ but by using the components of \bar{E} rather than the components of E , we include the nodes and thereby unite adjacent components of E into a single component of \bar{E} . For example, in Figure 6 there are four components of E but only three components of \bar{E} , since in \bar{E} the lines AB and AC join to form a single component BAC.

Let Q be a tree with $\bar{E} \neq Q$. We first construct a CPT S with the additional property that every component \bar{E}^j is searched in a single CPT of \bar{E}^j , which we call C_j ; note that some CPTs of Q might search different subsets of \bar{E}^j during non-consecutive time intervals - we exclude this possibility by construction.

To obtain a CPT of Q with this property, we begin at any regular point not in \bar{E} and go in either direction. When arriving at any node, we leave by a passage not already traversed, if there is such a passage (This is the usual depth-first construction and ensures we obtain a CPT.) Furthermore, if the node belongs to some component \bar{E}^j and there are untraversed passages staying in that component, we take one of these. For example, in Figure 6 if we start on FA going right, and tour the leaf arc to B from A, we must then take the passage to C (staying in component BAC) rather than the other untraversed passage out of A going to E. This ensures that the CPT say ABAEACADA (in which the component BAC of \bar{E} is not traversed in a single CPT of BAC) will not be constructed, but rather one like [ABACA]EADA, where the bracketed expression is a CPT of the component BAC.

Then we make two types of additions at every component. If $\lambda(\bar{E}^j) \geq \alpha/2$, we follow the CPT C_j of \bar{E}^j in S by another identical one, before continuing with S . Note that this local CPT takes time $\geq \alpha$, so the time between the first and second CPT of \bar{E}^j reaching any (regular) point is at least α .

If $\lambda(\bar{E}^j) < \alpha/2$ then we cannot simply tour \bar{E}^j twice in succession because points of \bar{E}^j will not be visited at least two times that are at least time α apart. Instead we wait

until S returns to \bar{E}^j after the first occurrence of G in S , and then insert another C_j . We must check that these two times that S visits \bar{E}^j are separated by time at least α . Let $[t_1, t_2]$ be the time interval during which S tours \bar{E}^j so that $S(t_1) = S(t_2) = x$, say, and $t_2 - t_1 = 2\lambda(\bar{E}^j)$. We claim that $Q - \bar{E}^j$ has at least two components. If not, then x , which is on the boundary of \bar{E}^j , must be in the interior of an arc. Let $x' = S(t_1 - \varepsilon) = S(t_2 + \varepsilon) \notin E$, where $d(x, x') < \alpha/2 - \lambda(\bar{E}^j)$. Then $S((t_1 - \varepsilon, t_2 + \varepsilon))$ is a component of $Q - x'$ with length $\lambda(\bar{E}^j) + \varepsilon < \alpha/2$, so $x' \in E$, a contradiction.

Thus, $Q - \bar{E}^j$ has at least two components and they must all have length greater than $\alpha/2$, since any component with length at most $\alpha/2$ would be a subset of \bar{E} , and could not be disjoint from \bar{E}^j . So the next time after t_2 that S arrives at x is $t_3 \geq t_2 + \alpha$, and the next time after t_3 that S arrives at x is at least $t_3 + \alpha$. Then S is updated by adding another tour of C_j at time t_3 .

Observe that each additional local CPT of \bar{E}^j takes time $2\lambda(\bar{E}^j)$, so the total length of the resulting tour S^E is $2\mu + 2(\sum_j \lambda(\bar{E}^j)) = 2(\mu + \lambda(E))$ and by construction it reaches every point of Q at two times separated by at least α (modulo the length of the tour). Note that if $\bar{E} = Q$, we simply take $\bar{S} = S$. The optimal periodic strategy is thus \bar{S} . For the network of Figure 6, taking S as *ABACADA* we could have $\bar{S} = \text{ABACAFD } [FDF][ABACA]GE[GEG]A$, where the brackets indicate the three inserted local CPT's of the components \bar{S} . Note that two of these are inserted right after their first occurrence, but the third one [ABACA] is inserted nonconsecutively. Our construction would not work directly on the CPT *ABAEACADA*.

Thus we have established the following result by explicit construction.

Lemma 6 *Let Q be a tree. Then there is a tour S^E , an E -patrolling strategy, of length $2(\mu + \lambda(E))$ such that every point x of Q is visited at least twice at times that differ by at least α .*

We conjecture that E -patrolling strategies are always optimal for trees, and we later confirm the conjecture in some special cases. For now we give a general bound on the value of the game obtained by using an E -patrolling strategy. Let $v^* \equiv \alpha / (\mu + \lambda(E))$.

Theorem 7 *Let Q be a tree. Any E -patrolling strategy intercepts any attack with probability at least v^* .*

Proof: Follows from Lemma 6 and Theorem 2 part (i) with $k = 2$, $S = S^E$, and $L = 2(\mu + \lambda(E))$. \square

Note that when $\alpha \leq g^*$ we have $\lambda(E) = \alpha/2$, and the result of Theorem 7 becomes the same as the result of Corollary 2. In that case, the patrolling strategy S_2^α gives the same lower bound as an E -patrolling strategy.

6.3 The E -attack Strategy

In the previous section we showed that on a tree, any E -patrolling strategy intercepts any attack with probability at least v^* . Here, we define the E -attack strategy, whose attacks are intercepted with probability at most v^* on some trees. The condition that allows this strategy to be defined and to be optimal is given in Definition 8. It is useful to note that while for patrolling strategies we looked at the components of the closure \bar{E} of E , for the attack strategy given here we look at the components of E itself.

Definition 8 (Leaf Condition) Let Q be a tree. We say that (Q, α) satisfies the **Leaf Condition** if E consists of all points on every leaf arc within distance $\alpha/2$ of its leaf node.

For example, in Figure 5 the cases that satisfy the Leaf Condition are the first four ($\alpha = 1, 2, 3, 4$), where E consist of five components. All of these five components are subsets of leaf arcs and they are within $\alpha/2$ from the leaf node. Note that the Leaf Condition implies that every component E_j of E corresponds to a leaf node; this is easy to check in Figure 5. Cases $\alpha=5, 6, 7, 8$ have seven components; five of these components are subsets of leaf arcs but two of them are subsets of non-leaf arcs and thus (Q, α) does not satisfy the Leaf Condition.

Definition 9 (E -attack strategy) Suppose (Q, α) satisfies the Leaf Condition, where Q is a tree. Let x_j denote the leaf node contained in the closure of the component E_j of E , and let $e_j = \lambda(E_j)$ and let $M = \max_j \lambda(E_j)$ be the maximum length of a component of E . We define the E -attack strategy as follows:

1. With probability $\lambda(E^c)/(\mu + \lambda(E))$, attack a uniformly random point of E^c at time M .
2. With probability $2\alpha/(\mu + \lambda(E))$, attack at leaf node x at a start time chosen uniformly in the interval $[M - e_j, M + e_j]$.

Note that the Leaf Condition implies that $\sum_j e_j = \lambda(E)$, therefore the sum of the probabilities from 1. and 2. above sum to 1. Also, unlike the uniform attack strategy, the E -attack strategy is not synchronous. That is, the attack does not start at a fixed, deterministic time.

Example 3 We revisit Figure 6, where the leaf arcs have lengths 2, 1, 6, 6 and $\alpha=6$. We illustrate the E -attack strategy on this star network in Figure 7. Here $\mu=15$; the extremity set E is shown in thick red lines. E consists of four components that are subsets of leaf arcs and whose points are within ρ_2 from the leaf node, thus the Leaf Condition is satisfied. Also, note that $\lambda(E)=9$ and $\mu+\lambda(E)=24$. The E -attack strategy then attacks as follows: with equal probabilities $6/24$ it attacks at nodes D and E with a starting time chosen uniformly on $[0, 6]$; with probabilities $4/24$, $2/24$ it attacks leaf nodes B , C with a starting time chosen uniformly on $[1, 5]$, $[2, 4]$ respectively; with probability $1/24$ it attacks uniformly on set E^c at time $M=3$.

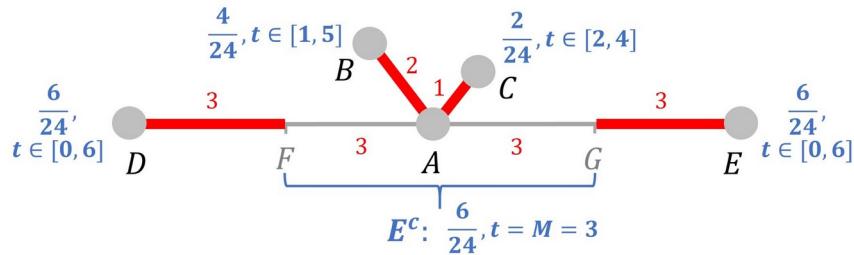


Figure 7: The E -attack strategy on an asymmetric star with arcs lengths 2,1,6,6 with $\alpha=6$. The set E is shown in thick red lines.

We next prove that for trees Q , the E -attack strategy is optimal if (Q, α) satisfies the Leaf Condition.

Lemma 7 Let Q be a tree and suppose (Q, α) satisfies the Leaf Condition. Then the E -attack strategy is intercepting by any patrol with probability at most $v^* = \alpha / (\mu + \lambda(E))$.

The proof of Lemma 7 is in the Online Appendix. If we combine the results of Theorem 7 and Lemma 7 on patrolling and attack strategies for trees, we obtain the following exact result for the value of the game.

Theorem 8 Let Q be a tree and suppose (Q, α) satisfies the Leaf Condition. Then any E -patrolling strategy is optimal, the E -attack strategy is optimal, and the value of the game is $V = v^*$.

Example 4 We revisit the network Q from Figure 7 with $\alpha=6$ and $\mu=15$. We first consider patrolling strategies. The S_2^α patrolling strategy is ADDABBACCAEEA, where repeating a node means it stays there for duration α ; this tour has length $2(6) = 54$. From Corollary 2

we have $V \geq \frac{\alpha}{\mu + \lambda(E)/2} = 6/27$. An E -patrolling strategy is $ADF DABACABACAEGEA$ with length $2\mu + 2\lambda(E) = 48$; from Theorem 7 we have $V \geq v^* = \frac{\alpha}{\mu + \lambda(E)} = 6/24$. As we can see, an E -patrolling strategy, which is defined only for trees offers an improvement over the \S patrolling strategy, which is a more general strategy.

Now, we consider attacker strategies. Let I be the set of leaf nodes. The sets E and $W \equiv W(I)$ are shown in Figure 8 with solid thick red and dashed thick green lines respectively. Note that (Q, α) satisfies the Leaf Condition. The E -attack strategy is demonstrated in Figure 7 and it gives a lower bound, $v^* = \frac{\alpha}{\mu + \lambda(E)} = 6/24$, from Theorem 8, which is optimal. The bound given by Theorem 1 $\frac{\alpha}{\lambda(W^c) + \lambda(E)} = \frac{\alpha}{\mu + \lambda(E)/2} = 6/27$ does not hold in this case because I is not an independent set or, equivalently, leaf arcs do not have lengths exceeding α .

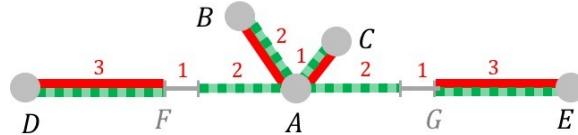


Figure 8: Star with arc lengths 6,6,2,1 and $\alpha=6$. The solid thick red line is the set E and the thick dashed green line is the set $W \equiv W(I)$, where I is the set of leaf nodes; note that here I is not an independent set.

A star is a network consisting entirely of leaf arcs. We call a star *balanced* if no arc comprises more than half of its total length; otherwise we say that it is *skewed*. It is easy to check that balanced stars satisfy the Leaf Condition. All symmetric stars (whose arcs are all the same length) are balanced. An example of a skewed star is a star with arc lengths 1, 1, 6 as shown in Figure 9: $\mu = 8$ and one of the arcs has length 6, which is more than half of μ .

It is also easy to see that if Q is a star (which may be balanced or skewed) whose longest arc has length at most $\alpha/2$, then $\bar{E} = Q$ and hence Q satisfies the Leaf Condition. So Theorem 8 gives the following.

Corollary 3 Let Q be a star. Then the E -attack strategy and any E -patrolling strategy are optimal and the value of the game is $V = v^* = \alpha / (\mu + \lambda(E))$ if either

- (i) Q is balanced or
- (ii) α is at least twice the length of the longest arc of Q .

We study the skewed star of Figure 9 in Section 7.

Note that if Q is the line segment network, then by adding an artificial node in the center, we can apply Corollary 3, part (i), recovering the result for the value of this game, given previously in Alpern *et al.* (2016) (though the optimal strategies given here are different).

7 Trees Not Satisfying the Leaf Condition

In the last section we considered patrolling games on trees. We showed (Theorem 7) that the E -patrolling strategy intercepts any attack with probability at least $v^* = \alpha / (\mu + \lambda(E))$ and that (Lemma 7) for trees satisfying the Leaf Condition, the E -attack strategy avoids interception with probability at least v^* . Thus for trees we have $V \geq v^*$ if the Leaf Condition is satisfied, but what happens when it is not satisfied? In this section we present some trees Q and attack durations α for which the Leaf Condition fails but nevertheless $V \geq v^*$. We do this by specifying particular attack strategies which are optimal on these trees.

We conjecture the following on trees:

Conjecture 1 *Let Q be a tree network, then for any α the E -patrolling strategy is optimal and the value of the game is $V = v^* = \alpha / (\mu + \lambda(E))$.*

In Section 7.1 we find such strategies for a range of values of α on a skewed star with lengths 6, 1, 1 and $\neq 8$ as shown in Figure 9. It is easy to see that this star satisfies the Leaf Condition only for $\alpha \leq 4$ and $\alpha \geq 12$. In what follows we introduce attack strategies for this star that guarantee v^* for the attacker for $4 \leq \alpha \leq 8$, and thus verify the conjecture. We refer to this skewed star as the 6–1–1 star. Finally, in Section 7.2 we give an attack strategy on a non-star tree that also guarantees the value v^* for the attacker and thus verifies the conjecture.

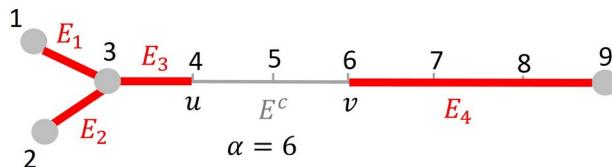


Figure 9: The 6–1–1 star. The extremity set E is shown in thick red lines and E^c in grey.

7.1 Optimal Attack Strategies on the 6-1 Star, $4 \leq \alpha \leq 8$

We define an attack strategy that we will show is optimal for the 6-1-1 star for $4 \leq \alpha \leq 8$. To aid notation we let $\theta = 2(\mu + \lambda(E))$; this gives $v^* = 2\alpha/\theta$. Further, we note that for the 6-1-1 star with $4 \leq \alpha \leq 8$ it is easy to check that $\lambda(E) = \alpha$. We denote the left and right boundary points of E^c with E by u and v respectively; since $4 \leq \alpha \leq 8$, both of these points are on the long arc or on its boundary.

We note that the Leaf Condition for this star holds for $\alpha = 4$ but not for $4 < \alpha \leq 8$, thus the E -attack strategy is not defined for the latter set of values. Thus, we define a new attack strategy. For $\alpha = 4$ both strategies can be used.

Definition 10 (6-1-1-attack) The 6-1-1-attack strategy is defined as follows:

Left attacks: With probability $2\alpha/\theta$, attack equiprobably at nodes 1 and 2 starting uniformly at times in $[1, 1+\alpha]$.

Middle attacks: With probability $2\lambda(E^c)/\theta$, attack at a uniformly random point of E^c (set of points between u and v), starting equiprobably at times 0 or $\alpha/2 + 2$.

Right attacks: With probability $2\alpha/\theta$, attack at node 9 at starting times in $[0, \alpha+2]$ described as follows: conditional on the attack taking place here, the starting time is given by the following cumulative probability distribution function,

$$f(y) = \begin{cases} \frac{y}{2\alpha} & \text{if } 0 \leq y \leq 2, \\ \frac{1}{\alpha} + \frac{y-2}{\alpha} & \text{if } 2 \leq y \leq \alpha, \\ \frac{\alpha-1}{\alpha} + \frac{y-\alpha}{2\alpha} & \text{if } \alpha \leq y \leq \alpha+2. \end{cases}$$

(This is uniform on $[2, \alpha]$ with conditional probability $(\alpha-2)/\alpha$ and uniform on $[0, 2] \cup [\alpha, \alpha+2]$, with conditional probability 2α).

Note that the total probability of attack is

$$\frac{2\alpha}{\theta} + \frac{2\lambda(E^c)}{\theta} + \frac{2\alpha}{\theta} = \frac{4\lambda(E) + 2(\mu - \lambda(E))}{\theta} = \frac{2(\mu + \lambda(E))}{\theta} = 1$$

Proposition 1 For the 6-1-1 star, shown in Figure 9, with $4 \leq \alpha \leq 8$, we have $V = V^* = 2\alpha/\theta$. The 6-1-1-attack strategy avoids interception with probability $v = \alpha/(\mu + \lambda(E))$.

The proof of Proposition 1 can be found in the Online Appendix. Proposition 1 provides a counterexample to a conjecture in Alpern *et al.* (2016). The conjecture was that for trees, if α is at least the diameter of the network, the value of the game is $\bar{b} = \alpha/(2\mu)$. For the 6-1-1 star, the diameter is 7, and by Proposition 1, for $\alpha < 8$, the value is $\alpha(\mu + \lambda(E))$. This is not equal to $\alpha/(2\mu)$, since $\lambda(E) < \mu$ in that range of α , disproving the conjecture.

The 6-1-1 strategy does not guarantee the value of v^* for the attacker for $\alpha > 8$. Note that for $\alpha > 8$ there are no middle attacks. Suppose that the Patroller stays at node 2 until time $\alpha - 1$, then arrives at node 1 at time $\alpha + 1$ and goes back to node 2 at time $\alpha + 3$ and then heads to node 9 and arrives there at time $\alpha + 10$. This patrol will intercept the left attacks with conditional probability of 1. The attacks at node 9 end at time $2\alpha + 2$ which is greater than $\alpha + 10$ if and only if $\alpha \geq 8$. Thus there will be positive probability p that the left attacks are intercepted. Thus, the total interception probability will be $\frac{2\alpha}{\theta} 1 + \frac{2\alpha}{\theta} p > \frac{2\alpha}{\theta} = v^*$.

7.2 A Non-star Tree with $\bar{E} = Q$ satisfying Conjecture 1.

We now consider the tree depicted in Figure 10 with unit length arcs and $\alpha = 6$. This gives $\bar{E} = Q$ and thus $\lambda(E) = \mu$. Here $\mu = 6$ and thus $v^* = \alpha/2\mu = 1/2$.

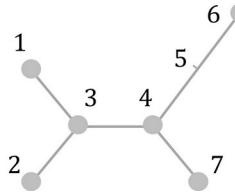


Figure 10: A tree with $\mu = 6$.

We propose the following Attacker strategy for this specific tree with $\alpha = 6$.

- At each leaf node 1 and 2 attack with probability $6/24$ at a start time chosen uniformly in the interval $[0, 6]$ (total attack probability $12/24$).
- At leaf node 6 attack with attack start time uniformly: in the interval $[0, 2]$ with probability $2/24$, in the interval $[2, 4]$ with probability $4/24$, in the interval $[4, 6]$ with probability $2/24$ (total attack probability $8/24$).
- At leaf node 7 attack takes place with probability $4/24$ at a start time chosen uniformly in the interval $[1, 5]$ (total attack probability $4/24$).

It is easy to verify that the probability of interception guaranteed by this strategy is $v^* = 1/2$, thus verifying the conjecture; the proof is along the same lines as that of Section 7.1.

8 Conclusions

This paper is the first to analyze continuous patrolling games on arbitrary networks. These games model the problem of defending transportation networks, pipelines or other networks which can be attacked anywhere along their length. In Section 3 we developed a number of very general techniques which can be used by the players to estimate the efficacy of various types of strategies on arbitrary networks. In Section 4 we solved the patrolling game for all networks without leaves, using a periodic patrolling path which covers every arc exactly twice (a double-cover), for attack lengths α not exceeding the girth g of the network. That result (Theorem 4) was then extended in Section 5 to general networks (allowing leaves) by considering a notion of generalized girth g^* (Theorem 6). For trees, g^* is simply twice the length of the shortest leaf arc, so Theorem 6 may not be very useful for trees with a short leaf arc. To remedy this, Section 6 developed a concept of the extremity set E of a tree and strategies for both players which are defined in terms of E . We then defined the Leaf Condition for a tree, which required (among other things) that the extremity set E is a subset of the leaf arcs. Conjecture 1 says that the value of the patrolling game on a tree is given by $v^* = \alpha / (\mu + \lambda(E))$ and that the strategies mentioned above based on E are optimal. Our main result for trees is that Conjecture 1 holds for all trees satisfying the Leaf Condition (Theorem 8). We then showed that the Leaf Condition (and hence Conjecture 1) holds for all stars where the length of any leaf arc is not more than half the total length of the star. Some stars without this property, as well as a certain non-star tree, are shown to satisfy Conjecture 1 (but not the Leaf Condition) by explicit construction of strategies in Section 7.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. CMMI-1935826.

References

Alpern S, Morton A, Papadaki, K (2011) Patrolling games. *Oper. Res.* 59(5):1246–1257.

Alpern S, Lidbetter T, Morton A, Papadaki K (2016) Patrolling a Pipeline. In *International Conference on Decision and Game Theory for Security 2016*, 129–138, Springer International Publishing.

Alpern S, Lidbetter T, Papadaki K (2018) Optimizing Periodic Patrols against Short Attacks on the Line and Other Networks. *Eur. J. Oper. Res.* 273(3):1065–1073.

Alpern S, Katsikas S (2019) The Uniformed Patroller Game. arXiv preprint arXiv:1908.01859.

An B, Ordóñez F, Tambe M, Shieh E, Yang R, Baldwin C, DiRenzo J III, Moretti K, Maule B, Meyer G (2013) A deployed quantal response-based patrol planning system for the U.S. Coast Guard. *Interfaces* 43(5):400–420.

Basilico N, Gatti N, Amigoni F (2012) Patrolling security games: Definition and algorithms for solving large instances with single patroller and single intruder. *Artif. Intell.* 184:78–123.

Basilico N, De Nittis G, Gatti N (2017) Adversarial patrolling with spatially uncertain alarm signals. *Artif. Intell.* 246:220–257.

Baston VJ, Bostock FA (1991) A generalized inspection game. *Nav. Res. Log.* 38:171–182.

Czyzowicz J, Gasieniec L, Kosowski A, Kranakis E, Krizanc D, Taleb N (2017) When Patrolmen Become Corrupted: Monitoring a Graph Using Faulty Mobile Robots. *Algorithmica* 79:925–940. <https://doi.org/10.1007/s00453-016-0233-9>

Edmonds J, Johnson EL (1973) Matching, Euler tours and the Chinese postman. *Math. Program.* 5(1):88–124.

Eggleton RB, Skilton DK (1984) Double tracings of graphs. *Ars Combin.* A 17:307–323.

Fang F, Nguyen TH, Pickles R, Lam WY, Clements GR, An B, Singh A, Tambe M, Lemieux A (2016) Deploying PAWS: Field optimization of the protection assistant for wildlife security. In *Twenty-Eighth IAAI Conference*.

Garrec T (2019) Continuous patrolling and hiding games. *Eur. J. Oper. Res.* 277(1):42–51.

Glicksberg IL (1952) A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *P. Am. Math. Soc.* 3(1):170–174.

Hochbaum DS, Lyu C, Ordóñez F (2014) Security routing games with multivehicle Chinese postman problem. *Networks* 64(3):181–191.

Klavzar S, Rus J (2013) Stable traces as a model for self-assembly of polypeptide nanoscale polyhedrons. *MATCH Commun. Math. Comput. Chem.* 70:317–330.

Lin KY, Atkinson MP, Chung TH, Glazebrook KD (2013) A graph patrol problem with random attack times. *Oper. Res.* 61(3):94–710.

Lin KY, Atkinson MP, Glazebrook KD (2014) Optimal patrol to uncover threats in time when detection is imperfect. *Nav. Res. Log.* 61(8):557–576.

Lin KY (2019) Optimal patrol on a perimeter. arXiv preprint arXiv:1905.03600.

Papadaki K, Alpern S, Lidbetter T, Morton A (2016) Patrolling a border. *Oper. Res.* 64(6):1256–1269.

Pita J, Jain M, Marecki J, Ordóñez F, Portway C, Tambe M, Western C, Paruchuri P, Kraus S (2008) Deployed ARMOR protection: The application of a game theoretic model for security at the Los Angeles international airport. *Proc. 7th Internat. Joint Conf. on Autonomous agents multiagent systems* (International Foundation for Autonomous Agents and Multiagent Systems, Southland, SC), 125–132.

Sabidussi G (1977) Tracing graphs without backtracking. Université Catholique de Louvain. Centre for Operations Research and Econometrics [CORE].

Xu L, Gholami S, McCarthy S, Dilkina B, Plumptre A, Tambe M, Singh R, Nsubuga M, Mabonga J, Driciru M, Wanyama F, Rwetsiba A, Okello T, Enyel E (2019) Stay ahead of poachersillegal wildlife poaching prediction and patrol planning under uncertainty with field test evaluations. arXiv preprint arXiv:1903.06669.

Yolmeh A, Baykal-Gürsoy M (2019) Patrolling Games on General Graphs with Time-Dependent Node Values. *Mil. Oper. Res.* 24(2):17–30.

Zoroa N, Fernández-Sáez M, Zoroa P (2012) Patrolling a perimeter. *Eur. J. Oper. Res.* 222(3):571–582.