

Adversarial Fence Patrolling: Non-Uniform Policies for Asymmetric Environments*

Yaniv Oshrat, Noa Agmon, Sarit Kraus

Department of Computer Science, Bar-Ilan University
oshblo@zahav.net.il, agmon@cs.biu.ac.il, sarit@cs.biu.ac.il

Abstract

Robot teams are very useful in patrol tasks, where the robots are required to repeatedly visit a target area in order to detect an adversary. In this work we examine the *Fence Patrol* problem, in which the robots must travel back and forth along an open polyline and the adversary is aware of the robots' patrol strategy. Previous work has suggested non-deterministic patrol schemes, characterized by a uniform policy along the entire area, guaranteeing that the minimal probability of penetration detection throughout the area is maximized. We present a patrol strategy with a non-uniform policy along different points of the fence, based on the location and other properties of the point. We explore this strategy in different kinds of tracks and show that the minimal probability of penetration detection achieved by this non-uniform (variant) policy is higher than former policies. We further consider applying this model in multi-robot scenarios, exploiting robot cooperation to enhance patrol efficiency. We propose novel methods for calculating the variant values, and demonstrate their performance empirically.

1 INTRODUCTION

The problem of multi-robot patrol has been thoroughly investigated during the past few years (Agmon et al. 2008; Hefferan, Cliff, and Fitch 2016; Talmor and Agmon 2017; Alam et al. 2017; Ivanová and Surynek 2017; Alam et al. 2019). In this problem, a team of robots is required to continuously travel in an area, in our case a line (fence), while monitoring it in order to detect changes in state. Many of the studies in this area have concentrated on assuring optimization of frequency criteria (Elmaliach, Agmon, and Kaminka 2009; Collins et al. 2013). However, we follow several studies that consider the problem of **multi-robot patrol in adversarial environments**: There is an adversary that wishes to penetrate this line without being detected. We assume that the adversary has full knowledge about the patrolling robots

(their capabilities, their locations and patrolling behavior), and would take the optimal steps (in his view) in order to penetrate successfully. Our goal is to set the patrol in a way that would minimize the adversary's probability of success. Naturally, this problem is applicable in many security domains.

In this paper we focus on the problem of patrolling along an open polyline track. Former studies presented policies that assigned a uniform policy (probability value) to all of the segments in the track. We suggest assigning each segment with a distinct probability value, according to the segment's position in the track. We show analytically and empirically that the results obtained by this approach are significantly better than those of the former approach.

Moreover, previous studies (e.g. Agmon, Kraus, and Kaminka 2008) concentrated on symmetric tracks, i.e. all of the parts of the track are similar in their parameters. This similarity holds for both the physical properties of the line and the algorithmic attributes of the robots' motion. In this work we challenge this symmetry, from practical considerations (reality often presents asymmetric structures that should be dealt with) and for optimality calculations. We show that, in many cases, non-uniform policies yield better results than uniform ones.

2 RELATED WORK

Systems of multiple robots cooperating in an arena to perform a joint mission were studied in recent years in various contexts. The arena might be an area that should be watched over, cleaned or monitored (Chevalyere 2004; Alam et al. 2017). Alternatively the arena is a line that should be monitored in order to detect intrusions or to prevent them (Lin 2019). The offered solutions differ in the ways the quality of patrol is measured: Many of the works focused on frequency-based patrolling – optimizing the frequency of robot visits in the various locations of the arena (Almeida et al. 2004; Elmaliach, Agmon, and Kaminka 2009). Sometimes there are vital points that should be monitored in a more frequent manner than other points (Collins et al. 2013). Other works concentrate on adversarial patrol schemes (Basilico et al. 2009; Sless, Agmon, and Kraus 2014; Alam et al. 2019). This approach (also addressed in

*This research has been supported in part by the Ministry of Science and Technology, Israel and the Japan Science and Technology Agency (JST), Japan.
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

this paper) assumes an adversary that wishes to penetrate a line - either an open polyline or a closed perimeter. The adversary has the time and knowledge to study and understand the properties of the patrol, and to aim its penetration trials at its weakest points. In this context, an efficient patrol would be the one that makes it harder for the adversary to locate or exploit its vulnerabilities for penetration. Non-deterministic approaches may prevent the adversary from assuring the success of the penetration by choosing a smart penetration point. Performance evaluation of models and algorithms is done according to these different definitions of quality (Portugal and Rocha 2013). In this work we evaluate the performance by the concept of Probability of Penetration Detection (ppd) which we wish to maximize (Agmon, Kaminka, and Kraus 2011).

Another approach to adversarial patrolling considers the game theoretic approach based on stochastic Stackelberg game models. The backbone of this method is defender-attacker Stackelberg games in which the defender (leader) first commits to a randomized security policy, and the attacker (follower) uses surveillance to learn about the policy before attacking. The application might be either attack-timing indifferent (Basilico, Gatti, and Amigoni 2012), or with temporal preferences incorporated using exponential discounting (Vorobeychik, An, and Tambe 2012; Vorobeychik et al. 2014). A recent comprehensive survey of many patrol models is brought in (Huang et al. 2019).

In previous works dealing with optimizing ppd, the problem of patrolling along a fence (an open polyline) has been studied under the assumption that the robots use one probability variable p in all of their decisions (Agmon, Kaminka, and Kraus 2008). Although more complex possibilities were mentioned, they were neither investigated nor analyzed. In this paper we revisit the problem and consider more complex strategies of multiple probability variables.

3 MODELS AND NOTATIONS

We are given a team of k identical robots that are required to patrol along an open polyline. The line (track) is divided into N segments, $S = \{s_1, \dots, s_N\}$, and each robot travels through one segment in one time-cycle. Note that the segments are not required to be of the same length or the same orientation: Some of them may be longer than others, contain turns, obstacles or other special characteristics, according to the conditions of the target area. However, all segments are set with proper length and robot velocity such that each robot traverses each segment in one time-cycle. We presume that robots travel on a pre-designated track and their movement is not interrupted by passersby.

In the *disjoint sectors policy* (Agmon, Kaminka, and Kraus 2008; 2011), the N segments are divided into k disjoint sectors, $C = \{c_1, \dots, c_k\}$, where each sector $c_i \in C$ contains $d = N/k$ segments (d is assumed to be an integer). Each sector is patrolled by exactly one robot, and the robots do not enter each others' sectors. In this policy we focus our analysis of the system on the analysis of a single sector, since the behavior is equivalent in all sectors. In section 5 of this paper we also discuss the *non-disjoint policy*, in which sectors have one or more overlapping segments which are

monitored by more than one robot. In this policy there may be several ways to divide a given track into sectors.

The robots have directionality associated with their movement. In each time-cycle, a robot has to decide where it should go - to either continue in its current direction or to turn around and go in the opposite direction. Turning around might be costly, i.e. if the robot decides to change its direction, it takes τ time-units to perform the turning and start the patrol in the opposite direction. In this paper, for simplicity, we demonstrate the movement with $\tau = 0$ (a realistic scenario for rail-mounted robots, for example), but the model can be applied for any $\tau \in \mathbb{N}$: (Agmon, Kaminka, and Kraus 2011) present several movement models, including ones with $\tau > 0$, and show how to convert calculations from one model to the other.

A *robot configuration* ($Rconfig_j$) is defined as the pair $\langle s_i, dir \rangle$, where s_i indicates the location of robot j (the segment in which it resides), and $dir \in \{right, left\}$ is the robot's direction. For example, $Rconfig_7 = \langle s_2, right \rangle$ means that robot number 7 is located in the second segment in the sector, and is moving to the right.

The patrol algorithm of the robots is characterized by a probability, i.e., in each step a robot continues in its current direction with some probability p , and turns around with probability $1 - p$. There are various models for assigning the probability to a specific robot configuration and they are elaborated on below. **Finding the optimal value of p for each robot configuration is the essence of the adversarial patrol problem.**

Penetration configuration is defined as the tuple $\langle s_{adversary}, Rconfig_1, \dots, Rconfig_j \rangle$, where $j \geq 1$, $s_{adversary}$ indicates the segment through which the adversary is going to penetrate, and $Rconfig_i$ indicates the configuration of robot i just before the penetration begins. For example, the penetration configuration $\langle s_1, \langle s_3, right \rangle, \langle s_7, left \rangle \rangle$ describes the situation in which there are two robots in the track, the adversary tries to penetrate segment 1, starting as robot No. 1 is in segment 3 moving to the right and robot No. 2 is in segment 7 moving to the left.

The adversary has to decide on its preferred penetration configuration before time-cycle 0. It may take several time-cycles for this penetration configuration to form, and the adversary will wait for it to arrive. At time-cycle 0 the adversary will be ready, in place, and in time cycle 1 the penetration will begin. We assume that the time it takes the adversary to penetrate, *penetration time*, is not instantaneous, and lasts $t > 0$ time-units, from time-cycle 1 to time-cycle t . Note that in time-cycle 0 the adversary has not yet penetrated, and the robots are still in the starting segments that were chosen by the adversary. At time-cycle 1 the penetration commences, and the robots are no longer at their starting positions - all of them have already moved, either to the left or to the right, according to the stochastic decision each of them made independently.

The robots succeed in their detection mission if (at least) one of them traverses a segment while the adversary is in this segment. If t time-units have passed since the adversary commenced its penetration and none of the robots have traversed this segment, then the penetration is classified as

successful and the robots fail. The adversary does not have another chance, nor is the adversary allowed to stop the penetration attempt after it has commenced and retry later.

The concept of Probability of Penetration Detection (ppd) is essential to our analysis: In general, ppd is the probability of detecting a penetration. When applied to a specific penetration configuration, ppd_{conf} is the probability that an adversary, commencing penetration in configuration $conf$, is detected by some robot during its penetration attempt. When referring to a discreet segment s_i , ppd_{s_i} is the minimal ppd of all the penetration configurations with this segment as the penetrated segment. In other words, ppd_{s_i} is the minimal probability that a patrol path of any robot will pass through segment s_i during the time that a penetrator is passing through that segment. Note that ppd_{conf} and ppd_{s_i} are functions of the probability values of the track.

In this work we concentrate on assigning optimal probability values to track segments in order to maximize the minimal ppd_{s_i} , $\forall s_i \in S$, namely MaxMinppd . With these values, the adversary, although knowing all of the patrol parameters and choosing the weakest point, i.e., the segment s_j with the lowest ppd_j value, as its penetration attempt arena, would have the minimal probability of succeeding in his attempt to penetrate the line.

4 UNIFORM VS. NON-UNIFORM P POLICY IN A DISJOINT MODEL

In an open polyline track there are two extreme segments (at both ends of the line) in which the robot movement is constrained, as it cannot continue in its current direction. In these segments $p = 0$ and the robot turns around with the probability $1 - p = 1$, no matter what the probability values in the non-extreme segments are. All of the non-extreme segments, however, may be assigned a p value of our choice in order to maximize Minppd in the track. In earlier works (Agmon, Kaminka, and Kraus 2008; 2011) the case of *uniform p policy* was analyzed, i.e., the probability value p in all of the non-extreme segments is equal (Fig. 1). The robot would “flip a coin” with probability p to decide whether to continue in its current direction or to turn around, regardless of the current segment and direction. This policy is practiced in the case of a closed perimeter track and is very reasonable there, since in a closed perimeter there is full symmetry between all of the segments and there is no way to distinguish between them (Agmon, Kraus, and Kaminka 2008). Nevertheless, this is not the case in the open polyline track, whereas there are extreme segments which force the robot to change its direction, hence the distance of a specific segment from the edge influences its ppd.




Figure 1: The uniform p policy in an open polyline.

For this reason it seems that a *non-uniform P policy*, in which $P = \{p_1, \dots, p_d\}$ and each segment s_i is assigned the probability value p_i , would be more appropriate and

might yield better results (i.e., higher MaxMinppd in the track). Although mentioned in the aforementioned works, no analysis of the non-uniform P policy in an open polyline was introduced. Specifically, neither a practical method to calculate the values of P nor a comparison between the results of both policies was presented.

In the non-uniform P policy there is a set $P = \{p_1, \dots, p_d\}$ of probabilities rather than a single p value in the uniform p policy. In our track notation, the leftmost segment is segment s_1 and the rightmost one is segment s_d . Both extreme values of P are preset: $p_1 = 1$, $p_d = 0$. The rest of the values p_2, \dots, p_{d-1} should be determined such that the MaxMinppd of the track is maximal. When its current direction is to the right, a robot that resides in segment s_i uses p_i as its coin value to decide whether to continue in its current direction or to turn around. When its current direction is to the left, a robot in segment s_i uses p_{d+1-i} (its mirror-image index) as its coin value (Fig. 2).




Figure 2: The non-uniform P policy in an open polyline. When the current direction is to the right the robot uses p_i as its coin value in segment s_i . When the direction is to the left the robot uses p_{d+1-i} in segment s_i .

4.1 Calculating the ppd

The ppds of the segments in the track are probability functions: $\forall i \in \{1, \dots, d\}$, $\text{ppd}_{s_i} : [0, 1]^d \rightarrow [0, 1]$. In order to compute the ppd functions, given a set P of track segment probabilities, we deploy the method described in section 4.2 in (Agmon, Kaminka, and Kraus 2011) as Algorithm 1. We build a Markov Chain G in which, for each segment s_i in the original track, two states of G are created: One for moving to the right and the other for moving to the left. If the robot reaches the segment through which the adversary tries to intrude within t time-units, then the adversary is caught.

From Markov chain G we derive the respective stochastic matrix M . The algorithm as described in (Agmon, Kaminka, and Kraus 2011) is set for the uniform p policy for the closed perimeter track, hence it results in a set of d single-variable ppd functions. In the variant P policy of an open sector with d segments, we have d^2 configurations: d possible segments for the adversary to penetrate, while the robot might reside in d possible segments at time-cycle 0. ppd values might be calculated regardless of robot direction: Due to symmetry considerations, calculations may be made assuming a specific direction (e.g. to the right) and these values hold also for the opposite direction. Therefore, we use the algorithm d times, for every segment through which the adversary may intrude, and combine the sets of functions we get as a result to form a $d \times d$ matrix of non-linear ppd functions with $d - 2$ variables p_2, p_3, \dots, p_{d-1} (recall: $p_1 = 1, p_d = 0$ are constants).

4.2 Calculating P for MaxMinppd in non-uniform P policy

The algorithms described in former works to calculate the value of p for MaxMinppd were designed for the uniform p policy (one-variable functions) and are not applicable here (d -variable functions). Numerical methods may be used to find an approximation to the MaxMinppd, e.g. the downhill simplex (Nelder-Mead) method (Nelder and Mead 1965) or the sequential quadratic programming (SQP) method (Bonnans et al. 2006). In our implementation we used the Nelder-Mead and the SLSQP methods of the function optimize.minimize() in the Python SciPy library for scientific computing (Gao and Han 2012; Kraft 1988). These numeric methods yielded good results in less than 24 hours in the lower values of track configurations (d and t).

For higher values ($t \geq 8$), using numeric methods for identifying (semi-)optimal P might take too long for practical use. For these cases we propose an Asymmetric Track Adversarial Patrol Search (ATAPS) heuristic that scans parts of the possible range of values in adjustable resolution and finds a policy within a given time frame. ATAPS receives as input a track configuration (d and t values) and a time-frame – the given time for calculation, at the end of which the result is required. According to the time-frame it decides which one of the three methods would be most appropriate for the case: (1) Numeric method as described above; (2) exhaustive search of the space using relevant resolution (adjust the range $r = [left, right] \subseteq [0, 1]$ and the resolution $\tau \in [0, 1]$, and evaluate $(r_{right} - r_{left})/\tau$ possible values for each p_i); or (3) random search using random restart hill climbing search. It also decides which verification method – the aforementioned matrix calculation or a fast simulation tool of patrols that we implemented – is faster for the specific task. These decisions are made considering the number of steps required for the calculation and the time needed for each step. For each possible set $P = \{p_1 \dots p_d\}$ ATAPS evaluates the ppds using the chosen tool. An outline of the ATAPS algorithm is provided in the supplemental material.

Fig. 3 demonstrates various cases and the consequences of applying different calculation methods to them. We measure the calculation time (in seconds) of each method, and the cases are shown by their *sector configuration* – the pair $\langle d, t \rangle$ (for simplicity denoted as d/t) where d is the number of segments in a sector and t is the penetration time. In track configuration 6/7 (Fig. 3a) the numeric

methods (Nelder-Mead, SLSQP) are the most efficient: They yield a high Minppd and do so in a short span of time, compared to an exhaustive search with matrices verification (XS+Matrices) and exhaustive search with simulation verification (XS+Simulation). In configuration 7/8 (Fig. 3b) we present results of eXhaustive Search with High Resolution (XSHR: $r = [0.5, 1], \tau = 0.03$) and eXhaustive Search with Low Resolution (XSLR: $r = [0.5, 1], \tau = 0.035$). The exhaustive search is faster than the numeric methods, but yields a lower Minppd. Random Search (RS) is even faster but yields even lower results. In configuration 8/10 (Fig. 3c) all search methods (exhaustive and random) yield results in practical time ($r = [0.5, 1], XSHR_\tau = 0.035, XSLR_\tau = 0.04$), and the random search does it faster and better than the exhaustive search, whereas the numeric methods fail to return a result at all even after 9 days of running.

4.3 Advantage of the Non-Uniform P Policy: Experimental Results

In open polyline tracks we are only interested in t values in the range $d - 1 \leq t \leq 2d - 3$. Lower t values will allow the adversary to find a segment to penetrate with a probability of 0 of being caught. Higher values of t mean that deterministic movement would catch the adversary with a probability of 1 and would be better than our stochastic model.

The results presented in the figures throughout the paper were obtained in computational experiments performed using the methods described in Section 4.2. Fig. 4 shows the results of the non-uniform p policy in comparison with the uniform p policy. It is clear that the Minppd of the non-uniform p policy is higher than that of the uniform p policy for all sizes of sectors d and all relevant t values. It can be seen that the advantage of the non-uniform p policy is greater in the “difficult” cases, when t is small, i.e. the adversary needs only a short period of time to penetrate. In reality, such cases occur when the fence is weak or low, or when the adversary has a knack for penetration.

It should be noted that as d increases, the MaxMinppd in the track decreases. For practical use, high values of d (i.e. fences with many segments equipped with few robots) would not be appropriate since the adversary would have a very low probability of getting caught. For example, facing a fence with 7 segments and a penetration time of 7 time-units, a smart adversary who cleverly chooses his penetration segment would only have a 25% probability of being

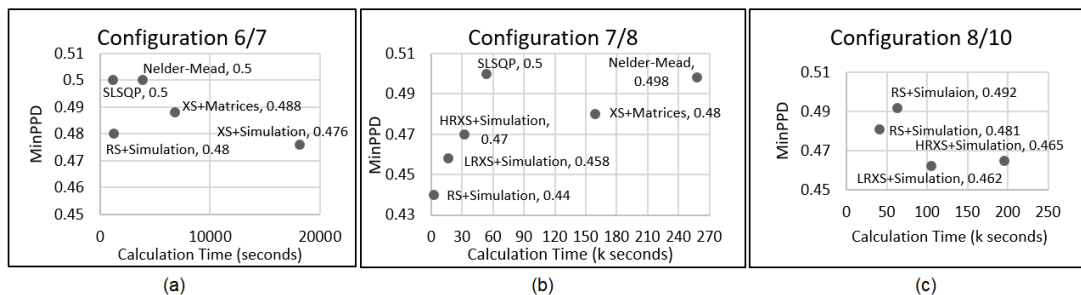


Figure 3: Minppd calculation results using different methods as a function of calculation time (in seconds).




Figure 4: Comparison between the uniform p model and the variant P model in different polyline track configurations. Each configuration is characterized by its d and t values (e.g., configuration 5/7 means an open polyline with 5 segments and a penetration time of 7 time-units).

caught even when the robots deploy the better non-uniform P policy. Nevertheless, a little improvement in fence properties that slightly increases the penetration time from 7 to 8 time-units would increase the probability of catching the adversary to 50% using a non-uniform P policy, much better than 30.3% with a uniform p policy.

4.4 Advantage of the Non-Uniform P Policy: Theoretical Study

As shown in Fig. 4, the non-uniform P policy yields better results (i.e. higher Minppd) than the uniform p policy for all track configurations that were checked. In the following theorem and supporting lemmas we show that, when considering the basic configuration (i.e. minimal relevant t), the non-uniform P policy will always overtake the uniform p policy. To prove this property, we first consider the best segment for the adversary to penetrate in such settings.

Lemma 1. *In a uniform p policy, when moving with an optimal p value that yields MaxMinppd , the Minppd will occur at the extreme segment/s, i.e. the best segment through which to intrude (from the adversary's point of view) is one (or both) of the extreme segments.*

Proof. Assume, towards contradiction, that the minimal probability of penetration detection is $\text{ppd}_{(j,i)}$, where (j,i) is the configuration in which the robot resides in segment s_j , the adversary intrudes in segment s_i , and s_i is not an extreme segment, i.e. $1 < i < d$. $\text{ppd}_{(j,i)}$ is the sum of all path-probabilities from s_j to s_i whose length $\leq t$. In this configuration, observe the extreme segment on the other side of s_i from s_j (if s_j is itself one of the extreme segments, then observe the other extreme segment), segment s_1 in Fig. 5. $\text{ppd}_{(j,1)}$ is the sum of all path-probabilities from s_j to s_1 whose length $\leq t$. Note that each path in $\text{ppd}_{(j,1)}$ is an extension of a path in $\text{ppd}_{(j,i)}$ and $\text{ppd}_{(j,1)}$ has a lower probability than its corresponding path beginning in $\text{ppd}_{(j,i)}$ since it was multiplied by $p < 1$ at least once. Therefore $\text{ppd}_{(j,1)} < \text{ppd}_{(j,i)}$, thus a contradiction. ■

In the following lemma we show that if the adversary chose one of the extreme segments as his penetration point,




Figure 5: Open polyline, s_1 and s_d are extreme segments.

then the best situation for him to commence the penetration is when the robot is in the farthest non-extreme segment and its direction is towards the far extreme segment.

Lemma 2. *In the uniform p policy, in a track of length d with the optimal p value for MaxMinppd and $t = d - 1$, MaxMinppd is achieved when the robot is in s_{d-1} and $p = (d - 3)/(d - 2)$.*

The outline of the proof shows that due to the low t value there is only one possible path from s_{d-1} to the penetrated segment (while other segments have more than one), and this path's ppd is lower than any other path's in the track. The full proof is presented in the supplemental material.

Note that for $d > 4$, the optimal p value in the uniform p policy will always be $p > 0.5$.

Theorem 1. *In an open polyline track with $t = d - 1$, given an optimal p value in the uniform p policy, decreasing p_{d-2} (while leaving all other p_i s as they are) will increase the Minppd .*

The intuition behind the proof is that ppd_{d-1} is the Minppd in the case of $t = d - 1$, and decreasing p_{d-1} increases ppd_{d-1} . We show that ppd_{d-1} is still the Minppd of the track. The outcome of the theorem is that a non-uniform P policy can always yield better results (i.e. higher MaxMinppd) than the best result of the uniform p policy in the case of $t = d - 1$. Again, the full proof is presented in the supplemental material due to lack of space.

4.5 Tracks with an irregular segment

In this section we discuss a case in which there is an irregular segment in the track – one specific segment is different from all the other segments in the sense that it has a shorter penetration time: Due to a lower fence, bad visibility by guards or other assisting conditions, the adversary needs less time to penetrate this segment. The penetration time of this segment is $t_{\text{irreg}} < t$. A track with at least one segment that has a different penetration time value than the other segments is denoted an *asymmetric track*. The specific location of the irregular segment is significant. An irregular segment located at the extremity of the track is much more problematic than one located in the middle of the track.

Lemma 3. *When moving with optimal P values that yield MaxMinppd , a Minppd will occur at the extreme segment/s, i.e. one of the best segments through which to intrude (from the adversary's point of view) is one (or both) of the extreme segments.*

Proof. The proof is similar to the proof of Lemma 1, which considered uniform p , with the exception that in the non-uniform model P values might be 1, and in this case the extreme segment indeed has a Minppd value, but other segments adjacent to the extreme segment might have the same value as well. ■

This lemma shows that, in an open polyline track, the most difficult areas to monitor are the extremities. Unlike a closed perimeter, where every segment might be visited by robots coming from both sides, an extreme segment of an open polyline track is guarded by just one robot. As shown in Fig. 6, the non-uniform P model has a significant advantage over the uniform p model in asymmetric tracks because assigning different P values can help balance this inequality between segments. Using segment-specific probability values might enhance the presence of the robot in the extremities at the expense of reducing the presence in the middle area, thus yielding a more balanced set of ppds with a higher Minppd. Fig. 6 demonstrates that the closer the irregular segment is to the track’s extremities, the more substantial the advantage of the non-uniform model.

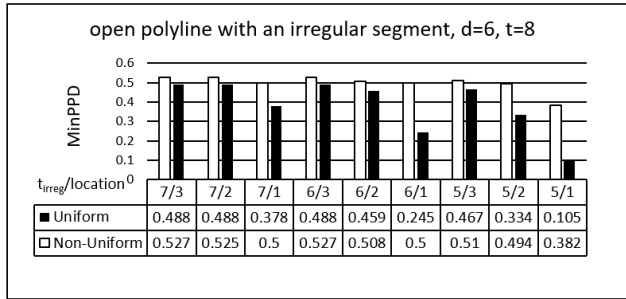


Figure 6: Comparison between the uniform p model and the non-uniform P model in an open polyline track, $d = 6$, $t = 8$, with a single irregular segment. The configuration notes the irregular segment penetration time value according to its location in the track (7/1 means a segment with $t = 7$ which is located in the first segment in the track).

5 NON-DISJOINT MODEL

5.1 Dual robot track

As mentioned above, former works dealing with optimizing Minppd in adversarial patrol (Agmon, Kaminka, and Kraus 2008; 2011) concentrated on a disjoint model, i.e. dividing the track into disjoint sectors with one robot patrolling along each sector. Other work, focusing on visit frequency as the patrol-quality measurement, suggested overlap of the assigned patrol areas of the robots (Elmaliach, Shiloni, and Kaminka 2008). We now introduce the non-disjoint division into the MaxMinppd arena, and show that this model may achieve higher Minppd values.

Consider a track with d segments. As long as $t \geq d - 1$, one robot can patrol the track and achieve positive Minppd. If t is smaller, Minppd = 0 and an additional robot must be added to the system. A similar situation happens when the Minppd is positive but is too low for practical use (for example, a single robot in a $d = 9$, $t = 8$ track would yield Minppd = 0.057, which is usually too low to be considered as proper monitoring). In the disjoint model, the track is divided into two disjoint sectors with one robot in each sector, and each sector is managed independently. The Minppd of the track is the minimal Minppd of the two sectors.

In the non-disjoint model we divide the track into two equal partly-overlapping sectors, making one segment (if d is odd) or two segments (if d is even) patrolled by both robots. The robots cooperate in monitoring the mutual segment(s), while each of the other segments is monitored by a single robot (see Fig. 7).




Figure 7: Dual robot patrol in an open polyline with $d = 9$: Top – non-disjoint scheme; bottom – disjoint scheme. The dark segment is monitored by two robots.

For example, consider a track with $d = 9$ and $t = 6$. This track cannot be properly monitored by one robot because the adversary can easily penetrate with probability 1 (e.g. commencing penetration of sector 9 when the robot is in sector 1). The disjoint model would add a robot and divide the track into two disjoint sectors, one of length 4 and the other of length 5. In the first sector the ppd is 1 in all of the segments, and in the second the Minppd is 0.48 (using uniform p policy) or 0.52 (using non-uniform p policy, as described in Fig. 4), so a knowledgeable adversary would penetrate the second sector in the optimal robot configuration. A better solution would be to form two sectors, each of them 5 segments long, and to include the middle segment in both sectors. Analysis shows that this configuration achieves Minppd = 0.62, much higher than in the former model. The reason for the significant improvement is that the weak points of an open polyline track are its extreme segments (as proven in Lemma 3). In the non-disjoint model the robots “help” each other in monitoring one of the extreme segments (the mutual middle one) which is patrolled by both of them, thus allowing an enhanced presence in the other extreme segments.

Calculation of the optimal P values for the non-disjoint model is done again using the ATAPS heuristic, but with essential adaptations: The two sectors are mirror-symmetric to each other, but the movement in a sector in one direction is **asymmetric** to the movement in the opposite direction, since the extreme segment in the inner part of the track is monitored by both robots and needs less presence by each of them than the other (non-middle) sector-ends. Therefore we look for $2(\lceil d/2 \rceil - 2) p$ values. In the aforementioned example ($d = 9$, $t = 5$, two sectors with 5 segments each) we need to calculate 6 p values, 3 for the movement from the edge towards the middle and another 3 for the movement from the middle towards the edge (again, the p values of the extreme segments of each sector are preset). Fig. 8 shows the results of the calculations by ATAPS.

Naturally, when using the disjoint schemes, the non-uniform model always yields a higher Minppd than the uniform model. The non-disjoint model produces even higher results in some of the cases. There are two cases, however, in which the non-disjoint model does not overpower the disjoint model: (a) In even d tracks when the t values are low:

Track configuration (d/t)	8/4	8/5	9/4	9/5	9/7	10/5	10/6	10/7	11/5	11/6	11/7	11/8
uniform disjoint policy	0.466	0.516	0.148	0.334	0.522	0.334	0.477	0.522	0.105	0.245	0.378	0.489
non-uniform disjoint policy	0.5	0.527	0.25	0.5	0.586	0.5	0.523	0.586	0.177	0.326	0.5	0.53
non-uniform non-disjoint policy	0.354	0.573	0.318	0.5	0.638	0.332	0.468	0.606	0.222	0.368	0.507	0.613

Figure 8: Comparison between different patrol schemes in an open polyline track with two robots. Each configuration is characterized by its d and t values.

Odd d tracks are divided into two uneven disjoint sectors (e.g., a 9-segment track is divided into a 4-segment sector and a 5-segment sector). The longer of these two newly formed sectors is the weak link, whereas it competes with a sector of the same length but with “help” from the other robot in the non-disjoint model (e.g. two 5-segment sectors with a mutual segment). Unlike this, even d tracks are divided into two equal-length sectors, both of them shorter than the sectors of the non-disjoint division, thus having an advantage. In low t values (e.g. 10/5) this advantage suffices to yield a better result, but in higher t values (e.g. 10/7) the non-disjoint model’s robot cooperation overpowers this advantage and produces a higher Minppd. (b) In special cases, when the single robot mode achieves its MaxMinppd in multiple points in the track: There are specific tracks where the MaxMinppd in the track is achieved in many places in the track, i.e. there are multiple penetration configurations that have the same MaxMinppd. In such cases, trying to increase the ppd in the non-mutual edge automatically decreases the ppd in other segments and yields a lower total Minppd in the track. An example of such a case is an open polyline with $d = 5$ and $t = 5$ (see Fig. 4). As a result, the dual robot mode with $d = 9$ and $t = 5$, which is divided in the non-disjoint division into two $d = 5, t = 5$ segments, cannot overpower the disjoint model and they both achieve the same MaxMinppd (see Fig. 8). A similar situation occurs in the $d = 6, t = 7$ case that affects the dual robot mode with $d = 11, t = 7$: Here the non-disjoint model manages to achieve an advantage over the disjoint model, but it is a minor advantage.

5.2 Numerous robots track

A similar concept can be applied to tracks with more than two robots. The non-disjoint model is even more efficient in these tracks, since all of the non-extreme sectors have their extreme segments monitored in cooperation with their neighbours. Only the extreme sectors of the track have real single-robot-monitored extreme segments, and we can make these sectors shorter in order to compensate for the lack of neighbours. For example, a track with $d = 11, t = 4$ cannot be properly monitored by less than 3 robots (due to the low t value). The disjoint model would divide the track into two 4-segment sectors and one 3-segment sector, and its weak point would have Minppd = 0.5 (see Fig. 4). Using the non-disjoint model, we divide the track into a middle sector of 5 segments and two extreme sectors of 4 segments each, whereas at every connection point there is one mutual segment (see Fig. 9). The middle sector is indeed longer but it gets “a little help from his friends” at both edges, thus achieving Minppd = 0.528. When choosing a patrol scheme




Figure 9: Triple robot patrol in an open polyline track with $d = 11$: Top – non-disjoint scheme; bottom – disjoint scheme. Dark segments are monitored by two robots each.

for a given track, there are several parameters that should be taken into consideration. Some of them, like the movement model, have no cost aspect, and the decision would be made according to the results only (i.e. the model that yields the highest Minppd is preferable). On the other hand, the number of robots to deploy in the track is an important factor that has costly consequences: Adding more robots would significantly improve the quality of the patrol, but is more expensive and complicated. Improving the quality of the patrol without adding robots to the track is an important practical feature.

To wrap up the concepts presented in this article, review Fig. 10. The table presents 6 different solutions to patrolling an open polyline track with 11 segments and a penetration time of 5. The solutions differ in the number of robots deployed, in the way the track is divided into sectors, in the schemes that are used to govern the movement of the robots and, of course, in the quality of the patrol – the Minppd. The contribution of the ideas from this work can be recognized in the table.

No.	Number of Robots	Scheme	Division	MinPPD	Total MinPPD
1	2	uniform, disjoint	5/5	0.334	0.105
			6/5	0.105	
2	2	non-uniform, disjoint	5/5	0.5	0.177
			6/5	0.177	
3	2	non-uniform, non-disjoint	6/5	0.223	0.223
			6/5		
4	3	uniform, disjoint	4/5	0.516	0.516
			4/5	0.516	
			3/5	1	
5	3	non-uniform, disjoint	4/5	0.527	0.527
			4/5	0.527	
			3/5	1	
6	3	non-uniform, non-disjoint	4/5	0.587	0.587
			5/5		
			4/5		

Figure 10: Optional solutions to patrol an open polyline track $d = 11, t = 5$. In light grey – dual robot configurations; in dark grey – triple robot configurations. Division indicates the way the track is divided into segments.

6 Conclusion

We presented a non-uniform P model for the movement of robots in adversarial patrolling of an open polyline. We showed that this model is significantly better than the former uniform model in regards to achieving maximal Minppd. We also suggested ways to calculate the optimal p values in order to apply the model to specific scenarios. In addition, multi-robot cooperation schemes were presented, showing that collaborative patrol between neighbouring robots can yield better results than deploying them disjointly.

References

- Agmon, N.; Sadov, V.; Kaminka, G. A.; and Kraus, S. 2008. The impact of adversarial knowledge on adversarial planning in perimeter patrol. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*, 55–62. International Foundation for Autonomous Agents and Multiagent Systems.
- Agmon, N.; Kaminka, G.; and Kraus, S. 2008. Multi-robot fence patrol in adversarial domains. In *Proceedings of the Tenth Conference on Intelligent Autonomous Systems (IAS-10)*, 193–201.
- Agmon, N.; Kaminka, G. A.; and Kraus, S. 2011. Multi-robot adversarial patrolling: Facing a full-knowledge opponent. *Journal of Artificial Intelligence Research* (42):887–916.
- Agmon, N.; Kraus, S.; and Kaminka, G. A. 2008. Multi-robot perimeter patrol in adversarial settings. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, 2339–2345. IEEE.
- Alam, T.; Rahman, M. M.; Bobadilla, L.; and Rapp, B. 2017. Multi-vehicle patrolling with limited visibility and communication constraints. In *Military Communications Conference (MILCOM), MILCOM 2017-2017 IEEE*, 465–470. IEEE.
- Alam, T.; Rahman, M. M.; Carrillo, P.; Bobadilla, L.; and Rapp, B. 2019. Stochastic multi-robot patrolling with limited visibility. *Journal of Intelligent & Robotic Systems* 1–19.
- Almeida, A.; Ramalho, G.; Santana, H.; Tedesco, P.; Menezes, T.; Corruble, V.; and Chevalleyre, Y. 2004. Recent advances on multi-agent patrolling. In *Brazilian Symposium on Artificial Intelligence*, 474–483. Springer.
- Basilico, N.; Gatti, N.; Rossi, T.; Ceppi, S.; and Amigoni, F. 2009. Extending algorithms for mobile robot patrolling in the presence of adversaries to more realistic settings. In *Web Intelligence and Intelligent Agent Technologies, 2009. WI-IAT'09. IEEE/WIC/ACM International Joint Conferences on*, volume 2, 557–564. IEEE.
- Basilico, N.; Gatti, N.; and Amigoni, F. 2012. Patrolling security games: Definition and algorithms for solving large instances with single patroller and single intruder. *Artificial Intelligence* 184:78–123.
- Bonnans, J.-F.; Gilbert, J. C.; Lemaréchal, C.; and Sagastizábal, C. A. 2006. *Numerical optimization: theoretical and practical aspects*. Springer Science & Business Media.
- Chevalleyre, Y. 2004. Theoretical analysis of the multi-agent patrolling problem. In *Intelligent Agent Technology, 2004. (IAT 2004). Proceedings. IEEE/WIC/ACM International Conference on*, 302–308. IEEE.
- Collins, A.; Czyzowicz, J.; Gasieniec, L.; Kosowski, A.; Kranakis, E.; Krizanc, D.; Martin, R.; and Morales Ponce, O. 2013. Optimal patrolling of fragmented boundaries. In *Proceedings of the twenty-fifth annual ACM symposium on Parallelism in algorithms and architectures*, 241–250. ACM.
- Elmaliach, Y.; Agmon, N.; and Kaminka, G. A. 2009. Multi-robot area patrol under frequency constraints. *Annals of Mathematics and Artificial Intelligence* 57(3-4):293–320.
- Elmaliach, Y.; Shiloni, A.; and Kaminka, G. A. 2008. A realistic model of frequency-based multi-robot polyline patrolling. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*, 63–70. International Foundation for Autonomous Agents and Multiagent Systems.
- Gao, F., and Han, L. 2012. Implementing the nelder-mead simplex algorithm with adaptive parameters. *Computational Optimization and Applications* 51(1):259–277.
- Hefferan, B.; Cliff, O. M.; and Fitch, R. 2016. Adversarial patrolling with reactive point processes. In *Proceedings of the ARAA Australasian Conference on Robotics and Automation (ARAA, 2016)*, 39–46.
- Huang, L.; Zhou, M.; Hao, K.; and Hou, E. 2019. A survey of multi-robot regular and adversarial patrolling. *IEEE/CAA Journal of Automatica Sinica*.
- Ivanová, M., and Surynek, P. 2017. Area protection in adversarial path-finding scenarios with multiple mobile agents on graphs: a theoretical and experimental study of target-allocation strategies for defense coordination. *arXiv preprint arXiv:1708.07285*.
- Kraft, D. 1988. A software package for sequential quadratic programming. Technical Report DFVLR-FB 88-28, DLR German Aerospace Center - Institute for Flight Mechanics, Cologne, Germany.
- Lin, K. Y. 2019. Optimal patrol on a perimeter. *arXiv preprint arXiv:1905.03600*.
- Nelder, J. A., and Mead, R. 1965. A simplex method for function minimization. *The computer journal* 7(4):308–313.
- Portugal, D., and Rocha, R. P. 2013. Multi-robot patrolling algorithms: examining performance and scalability. *Advanced Robotics* 27(5):325–336.
- Sless, E.; Agmon, N.; and Kraus, S. 2014. Multi-robot adversarial patrolling: facing coordinated attacks. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, 1093–1100. International Foundation for Autonomous Agents and Multiagent Systems.
- Talmor, N., and Agmon, N. 2017. On the power and limitations of deception in multi-robot adversarial patrolling. In *Proceedings of IJCAI*.
- Vorobeychik, Y.; An, B.; and Tambe, M. 2012. Adversarial patrolling games. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 3*, 1307–1308. International Foundation for Autonomous Agents and Multiagent Systems.
- Vorobeychik, Y.; An, B.; Tambe, M.; and Singh, S. P. 2014. Computing solutions in infinite-horizon discounted adversarial patrolling games. In *ICAPS*.