

Optimal Patrolling of Fragmented Boundaries

Andrew Collins, Jurek Czyzowicz, Leszek Gasieniec, Adrian Kosowski,
Evangelos Kranakis, Danny Krizanc, Russell Martin, Oscar Morales Ponce

► **To cite this version:**

Andrew Collins, Jurek Czyzowicz, Leszek Gasieniec, Adrian Kosowski, Evangelos Kranakis, et al..
Optimal Patrolling of Fragmented Boundaries. SPAA - 25th Symposium on Parallelism in Algorithms
and Architectures, 2013, Montreal, Canada. ACM, pp.241-250, 2013, <10.1145/2486159.2486176>.
<hal-00686270v3>

HAL Id: hal-00686270

<https://hal.inria.fr/hal-00686270v3>

Submitted on 9 Oct 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimal Patrolling of Fragmented Boundaries

Andrew Collins
University of Liverpool

Jurek Czyzowicz
Université du Québec en Outaouais

Leszek Gąsieniec
University of Liverpool

Adrian Kosowski
INRIA Bordeaux Sud-Ouest

Evangelos Kranakis
Carleton University

Danny Krizanc
Wesleyan University

Russell Martin
University of Liverpool

Oscar Morales Ponce
Chalmers University of Technology

Abstract

A set of mobile robots is deployed on a simple curve of finite length, composed of a finite set of *vital segments* separated by *neutral segments*. The robots have to patrol the vital segments by perpetually moving on the curve, without exceeding their maximum speed. The quality of patrolling is measured by the *idleness*, i.e., the longest time period during which any vital point on the curve is not visited by any robot. Given a configuration of vital segments, our goal is to provide algorithms describing the movement of the robots along the curve so as to minimize the idleness.

Our main contribution is a proof that the optimal solution to the patrolling problem is attained either by the *cyclic strategy*, in which all the robots move in one direction around the curve, or by the *partition strategy*, in which the curve is partitioned into sections which are patrolled separately by individual robots. These two fundamental types of strategies were studied in the past in the robotics community in different theoretical and experimental settings. However, to our knowledge, this is the first theoretical analysis proving optimality in such a general scenario. Throughout the paper we assume that all robots have the same maximum speed. In fact, the claim is known to be invalid when this assumption does not hold, cf. [Czyzowicz et al., *Proc. ESA* 2011].

Keywords: Algorithms, Mobile robots, Boundary patrolling, Idleness, Cycle

1 Introduction

Protecting an environment by a set of stationary or mobile point-guards has been studied before in many various scenarios. The problem of patrolling a one-dimensional boundary using mobile robots has many real-world applications, and is extensively studied under the names of *boundary patrolling* and *fence patrolling* in the robotics literature. In order to prevent an intruder from penetrating into a protected region, the boundary of the region must be patrolled. Some parts of the boundary may be monitored with stationary devices like sensors or cameras (or they do not need to be monitored at all), while other portions require the aid of moving robots such as walking guards, illumination rays, mobile robotic devices, etc. Since the feasibility of an intrusion likely depends on the time during which the intruder remains undiscovered, it is important to design patrolling protocols which minimize the time during which boundary points are unprotected.

Some portions of the boundary may be unpenetrable at all, or they may be monitored with stationary devices like sensors or cameras. This paper is devoted to the scenario in which only a finite set of boundary segments, referred to as *vital regions*, need to be patrolled by a set of k mobile robots. The remaining part of the boundary, called *neutral regions*, do not have to be monitored by the mobile robots, but may nevertheless be traversed by a robot since this may be the way to reach one vital region from

another. We study the problem of patrolling with the goal of minimizing the idleness of points located in the vital regions, i.e., the longest time during which such a point remains unvisited by a robot. We assume that at any time during the traversal the speed of each robot cannot exceed a certain maximum value, identical for all robots. Our goal is to define a set of functions describing the trajectories of all the robots in time.

The most common heuristics adopted in the past to solve a variety of patrolling problems include the *cyclic strategy*, where robots move in one direction around the cycle covering the environment, and the *partition strategy*, in which the environment is partitioned into sections patrolled separately by individual robots (or subsets of robots), using the terminology introduced in [6]. However, to the best of our knowledge, no theoretical studies formally proving the optimality of such approaches in this setting were done in the past.

It is worth noting, that in the more heterogeneous scenario where robots have different maximum speeds, neither the cyclic strategy nor the partition strategy leads to the optimal performance. In fact, it has been shown in [11] that for the case of 3 mobile robots with different maximal speeds patrolling a cycle (forming a single vital region), neither a partition strategy nor a cyclic strategy is optimal. It turns out that a specific hybrid strategy is better than each of these two fundamental approaches.

1.1 Model, Preliminaries, and Notation

We consider $k \geq 1$ mobile robots traversing a continuous rectifiable curve C .

Definition 1.1 (Traversal strategy (or schedule) on a curve) Consider a rectifiable curve C . A traversal strategy for a single mobile robot is a continuous function $f : [0, +\infty) \rightarrow C$ such that $t \mapsto f(t)$, whereby $f(t)$ is the position of the robot on the curve at time $t \geq 0$. More generally, for k mobile robots, a traversal strategy consists of k such continuous functions, one for each mobile robot.

Without loss of generality, we will assume that C is either a segment of unit length (in the case when C is an open curve), or a cycle of unit perimeter (when C is a closed curve). In both cases, we will use the notation $[0, 1]$ to parameterize the curve, noting that in the case of the cycle, points 0 and 1 are identified with each other. All robots may move along C at speeds not exceeding a given maximum value. For simplicity of notation, we set a unit of speed for which this maximum speed is equal to 1.

Definition 1.2 (Unit maximum speed) Let $dist(p, q)$ denote the distance between points $p, q \in C$ along the curve C . We say that the traversal strategy for k robots respects the robots' unit maximum speed if for any $1 \leq r \leq k$ and any $t_1, t_2 \in [0, +\infty)$ we have $\frac{dist(f_r(t_1), f_r(t_2))}{|t_1 - t_2|} \leq 1$.

We remark that since the distance metric $dist(p, q)$ only needs to capture the time required for a robot to travel between p and q , and does not have to correspond to the physical distance between points in a real-world setting, the above definition takes into account non-uniform terrains (boundaries with corridors, stair-cases, etc.), in which robots move along some parts of the terrain more quickly than along others. However, we do not take into account scenarios in which some robots move faster through a given section of the terrain than other robots.

The task of the robots is to patrol, so called, vital regions, located in the unit-length curve C (along which the robots move), so as to minimize the idleness of the points in the vital regions.

Definition 1.3 (Vital and neutral regions) The curve C contains n disjoint, consecutive, closed vital intervals V_1, V_2, \dots, V_n , separated by neutral intervals. Throughout the paper, we will write $C = (V, N)$, where $[0, 1] \supseteq V = \bigcup_{i=1}^n V_i$, and $N = [0, 1] \setminus V$. In the case when C is the unit segment $[0, 1]$, we have $V_i = [b_i, e_i]$ and $e_i < b_{i+1}$ for $1 \leq i \leq n - 1$, with $b_1 = 0$ and $e_n \leq 1$.

Definition 1.4 (Idleness) Let \mathcal{A} be a traversal strategy consisting of k continuous functions f_i for a system of k mobile robots, respectively, and traversing the given curve.

1. The idleness induced by \mathcal{A} at a point x of the curve, denoted by $I_x(\mathcal{A})$, is the supremum of the length of time intervals during which point x remains unvisited by any agent:

$$I_x(\mathcal{A}) = \sup_{\{0 \leq T_1 < T_2 : \forall t \in [T_1, T_2] f_i(t) \neq x\}} (T_2 - T_1).$$

2. The idleness of the system of k mobile robots induced by \mathcal{A} is defined by $I(\mathcal{A}) = \sup_{x \in V} I_x(\mathcal{A})$, the supremum taken over all vital points of the curve.
3. Finally, the idleness of the system of k mobile robots is defined by $I_{opt} = \inf_{\mathcal{A}} I(\mathcal{A})$, the infimum taken over all traversal strategies \mathcal{A} .

We can now formulate the main question we will address in this paper.

Question. Suppose that we have k mobile robots traversing a given rectifiable open curve (resp., closed curve) $C = (V, N)$, represented w.l.o.g. as the unit-length segment (resp., the unit-perimeter cycle). What traversal strategy should the mobile robots follow so as to minimize the idleness of the system?

In order to describe the region patrolled by a single robot in the partition strategies we propose, we will frequently refer to the concept of a lid.

Definition 1.5 ((d, k)-Lid cover) A d -lid is a contiguous interval on the curve of length d . We say that a curve $C = (V, E)$ has a (d, k) -lid cover if all of its vital regions can be covered by some set of k (not necessarily disjoint) d -lids.

1.2 Related work

The patrolling problem may be viewed as a version of an art gallery question, in which a set of stationary or mobile guards have to protect a given geometric environment (see [27, 31, 34]). In the setting with stationary guards, in most research papers the number of guards, needed to view the entire environment, has to be minimized. The problem is NP-hard and many approximation and inapproximability results were obtained (cf. [15, 22]). For the case of mobile guards, often known as the *watchman route problem*, the question of a single watchman was most often addressed. The optimization criterion is the path length traversed by the watchman, so that every point of the environment is seen from some position on the path. This is closely related to the traveling salesman problem. Unsurprisingly, many general watchman route problems are NP-hard (e.g. watchman tours of simple polygon with holes, suggested in [7] and corrected in [13]), touring a sequence of non-convex polygons [12], or link-distance watchman tours of simple polygon with holes [4]). However, for many specific cases, polynomial algorithmic solutions are available. The solution for simple polygons was proposed by [8], while [9], [30] and [10] solve, respectively, the “zookeeper route”, “safari route” and “aquarium keeper” problems.

In the *m-watchmen routes* problem, the sum S of m path-lengths must be minimized, so that each point of the environment must be seen from some position of one of the watchmen (cf., e.g. [28, 32]). Clearly S decreases with increasing m . Hence at one extremity, we obtain the case when m is large enough to have $S = 0$ (m stationary guards for art gallery are sufficient) and on the other extremity the single watchman question arises. As *m-watchmen routes* problem is NP-hard for simple polygons (cf. [1]) some restricted classes of polygons were considered in [5, 29].

Central to the watchman route problem is the notion of *visibility*. Some papers (e.g. [24, 30]), considered limited visibility of the mobile (e.g. [30]), or stationary (e.g. [24]) guards. Our work corresponds to the case of *zero visibility*, in which the robot sees only the point of the environment at which it is currently

present, giving rise to so-called patrolling problems. Clearly, in finite time, the mobile robots may only patrol a geometric environment given as a union of paths of bounded total length.

The act of *patrolling* is defined as the perpetual process of surveillance consisting of walking around a terrain in order to protect or supervise it; it is performed either in a static or in a dynamically changing environment. It has been studied extensively in robotics literature (cf. [3, 6, 14, 16, 17, 23, 25, 36]) and it is sometimes viewed as a variant of *coverage* - a central task in robotics. Patrolling can be useful in settings where objects or humans need to be rescued from a disaster environment, but also network administrators may use mobile robot patrols to detect network failures or to discover web pages which need to be indexed by search engines, cf. [25].

Similarly, *boundary patrolling* may be motivated by the task of detecting intruders (from the exterior) in a two dimensional terrain by patrolling its boundary. There exist several studies on boundary patrolling (cf. [2, 16, 17, 33]); often the approach followed is ad hoc, emphasizing either experimental results (e.g. [25]), or uncertainty of the model and robustness of the solutions when failures are possible (e.g., [16, 17, 23]), or non-deterministic solutions (e.g. [2]).

The fundamental measure for evaluating the efficiency of patrolling is the criterion of *idleness*, first introduced in [25]. The general idea is to measure frequency of visits of the points of the environment by incoming robots (cf. [3, 6, 16, 17, 25]). As such the idleness is sometimes viewed as the average (cf. [16]), worst-case (cf. [36]), probabilistic (cf. [2]) or experimentally verified (cf. [25]) time elapsed since the last visit of a node (cf. [3, 6]). Also, in some papers the terms of *blanket time* (cf. [36]) or *refresh time* (cf. [33]) are being used instead, so as to indicate a similar measure of algorithm efficiency. Several approaches to patrolling based on idleness criteria were surveyed in [3], including machine learning, negotiation mechanisms for generating paths, heuristics based on local idleness, as well as approximation to the Traveling Salesmen Problem (TSP). Some papers study patrolling based on swarm or ant-based algorithms (cf. [18, 26, 36]) and explore various robot capabilities (sensing, memory, locomotion, etc). The *skeletonization technique*, where a terrain is first partitioned into cells is often applied in geometric environments prior to employing graph-theoretic methods in discrete time. In graph environments, cyclic strategies often rely on either TSP-related solutions or spanning tree-based approaches ([21, 33]). For the case of boundary patrolling where the robots maintain distinct maximal speeds, only partial solutions for small numbers of robots were proposed (cf. [11]).

One may also consider as a variant of patrolling the problem of searching a graph or polygon by teams of mobile robots (cf., e.g. [19, 35]), which are looking for a stationary or mobile intruder. This falls into the vastly investigated domain of cops and robbers (see [20]).

1.3 Outline and results of the paper

We start by recalling the *partition strategy* of patrolling in Section 2, in which each robot traverses some sub-interval of the curve back and forth. This strategy is always optimal on open curves (i.e., on the unit segment), but it need not be optimal for closed curves (i.e., for the unit cycle).

Our main results concern closed curves, and are given in Section 3. We prove that the optimal idleness for patrolling the boundary is always attained by the better of two strategies: the before-mentioned partition strategy, and the *cyclic strategy*, in which equally-spaced robots patrol the cycle, moving in the same direction. The choice of the strategy and the robots' responsibilities depends on the arrangement of the vital regions around the boundary. Our approach consists in showing that finding the optimal idleness for k robots and for any set of n vital intervals may be reduced to finding the idleness for some *critical* set of $2k + 1$ vital points (always resulting in the cyclic strategy) or of a critical set of $k + 1$ vital points (resulting in either the cyclic or the partition strategy).

Finally, in Section 4, we propose an $O(kn \log n)$ algorithm for designing traversal strategies with optimal idleness for robots on both open and closed curves.

2 Optimal Patrolling Strategy for the Segment

We start by studying the patrolling problem for k mobile robots on a curve $C = (V, N)$ consisting of n vital regions on the unit interval $[0, 1]$.

A natural approach to patrolling the segment is based on the partition strategy, in which each of the robots patrols exactly one of the k lids of the lid cover of minimum lid size L . The robot moves back and forth between its endpoints at maximum speed.

Partition strategy

1. Compute a (L, k) -lid cover of C , where L is chosen as the minimum lid length for which C admits such a lid cover. Let the i -th lid, $1 \leq i \leq k$, be a segment of the form $[c_i, c_i + L]$.
2. Deploy the i -th mobile robot so that at time $t = 2Lj + \tau$, where j is a non-negative integer and $-L \leq \tau < L$, the position $f_i(t)$ of this robot on the lid is $f_i(t) = c_i + |\tau|$.

We observe that each of the points of every lid, and consequently every vital point of the segment, is visited at least once during each time interval of size $2L$. So, for this strategy we have an idleness of $I \leq 2L$. The idleness of the partition strategy is, in fact, optimal.

Theorem 2.1 (Idleness for k mobile robots on a segment) *The optimal idleness for k mobile robots moving at speed at most 1 on a unit segment is given by $I = 2L$, where L is the minimum value such that the terrain admits a (L, k) -lid cover.*

The proof of Theorem 2.1 relies on the observation that, due to the fact that robots are indistinguishable, there exists an optimal solution in which the robots do not change their relative order on the segment, throughout the traversal (details are provided in the Appendix).

The complexity of computing the optimal lid cover for the partition strategy is discussed in detail in Section 4.

3 Optimal Patrolling Strategy for the Cycle

We are interested in computing the optimal idleness for k mobile robots traversing a unit-perimeter cycle $C = (V, N)$ with vital and neutral regions. The class of strategies under consideration in a cycle is larger than in the case of a segment due to the ability of the robots to traverse the perimeter of the cycle. In particular, the robots on the cycle can also apply a *cyclic strategy*, performing clockwise rotations around the cycle with even time spacing.

Cyclic strategy (on the cycle)

1. Deploy the i -th mobile robot at time 0 at position i/k along the circumference of the cycle.
2. Release all robots at their maximum speed to perform a clockwise traversal of the cycle.

Observation 3.1 *The idleness of the cyclic strategy on the cycle is $I = 1/k$, for any (non-empty) set of vital regions.* □

At the same time, we observe that the partition strategy introduced in the previous section is also applicable in the cycle, achieving an idleness of $I = 2L$, where L is the size of the minimum lid cover of the vital regions of the cycle with k lids. Depending on the configuration of the vital regions, one or the other of these two strategies may prove superior. In one extremal case when the cycle has no neutral

regions, the cyclic strategy achieves an idleness of $1/k$, while the partition strategy has an idleness of $2/k$. At the other extreme, for vital regions consisting of k discrete points, the idleness of the cyclic strategy is still $1/k$, but the partition strategy has an idleness of 0. This leads us naturally to a strategy which selects the better of the two approaches.

Combined strategy (on the cycle)

1. Let L be the lid size of the minimum (with respect to lid size) lid cover of the vital regions of the cycle with k lids.
2. If $2L < 1/k$, apply *Partition strategy*.
3. Else, apply *Cyclic strategy*.

Observation 3.2 *The idleness of the combined strategy on the cycle is $I = \min\{1/k, 2L\}$, where L is the minimum possible lid size of a (L, k) -lid cover of the cycle. \square*

This claim gives rise to the following natural question: Does there exist any other strategy which can achieve better idleness than both the partition and cyclic approaches? Such a question admits a positive answer for the cycle in the scenario where robots have different speeds [11], even when neutral regions are not present. In our scenario, with neutral regions but for robots with equal maximal speeds, the combined strategy turns out to be optimal. The proof of this fact is surprisingly involved.

Theorem 3.3 *The idleness $I(\mathcal{A})$ of any traversal strategy \mathcal{A} in a cycle with neutral regions satisfies $I \geq \min\{1/k, 2L\}$, where L is the minimum possible lid size of a (L, k) -lid cover of the cycle.*

The rest of this section is devoted to the proof of Theorem 3.3, which proceeds in three technical lemmas. First, we show that for any cycle with neutral regions we can find a subset of either exactly $k + 1$ or exactly $2k + 1$ (discrete) vital points that satisfy specific properties. Then, we show that the lower bound can be proved simply by considering the patrolling problem on the selected subset of points. The proofs of all the Lemmas are omitted or provided as sketches. For complete proofs, the reader is referred to the Appendix.

Lemma 3.4 (Critical point lemma) *Let $C = (V, N)$ be a cycle with set of vital regions V . Let L be the minimum size of the lid cover of the vital regions of C with k lids, and let $B = \sup\{\text{dist}(e, b) : b, e \in [0, 1], [b, e] \subseteq N\}$. Then:*

- (1) *If $B \geq 1/(2k)$, then there exists a set of $k + 1$ vital points $\{v_0, \dots, v_k\} \subseteq V$, ordered clockwise, s.t. $\min_{0 \leq i \leq k} \text{dist}(v_i, v_{(i+1) \bmod (k+1)}) \geq \min\{1/(2k), L\}$.*
- (2) *If $B < 1/(2k)$, then there exists a set of $2k + 1$ vital points $\{v_0, \dots, v_{2k}\} \subseteq V$, ordered clockwise, s.t. $\min_{0 \leq i \leq 2k} \text{dist}(v_i, v_{(i+2) \bmod (2k+1)}) > 1/(2k)$.*

Lemma 3.5 ($k + 1$ point lemma) *Let (v_0, v_1, \dots, v_k) be a set of $k + 1$ points chosen from vital regions of the terrain, arranged in clockwise order along the cycle. The idleness $I(\mathcal{A})$ of any traversal strategy \mathcal{A} for k mobile robots in this terrain satisfies $I(\mathcal{A}) \geq \min\{\frac{1}{k}, 2s\}$, where: $s = \min_{0 \leq i \neq j \leq k} \text{dist}(v_i, v_j)$.*

Proof. (sketch) The intuition of the proof is the following. First, we show that w.l.o.g. one can restrict the class of considered strategies in which the robots do not perform unnecessary movements within neutral regions. Subject to this assumption and supposing that for some strategy \mathcal{A} , $I(\mathcal{A}) < 2s$, we show that if we fix a sufficiently long time interval, then the number of times the robots traverse each arc of the cycle between two adjacent vital points is roughly the same for all such arcs. Finally, by applying a global counting argument, we prove that we must then have $I(\mathcal{A}) \geq 1/k$. In the argument, we capture

the intuition that each robot has to perform a $(1/k)$ -th part of the patrolling task. In particular, waiting at vital points cannot improve the performance of the strategy, since then each robot only protects 1 out of $k + 1$ vital points.

We will show the claim holds even if points $\{v_0, v_1, \dots, v_k\}$, satisfying $\text{dist}(v_i, v_{i+1}) \geq s$, for $0 \leq i \leq k$, are the only vital points of the cycle. Indices of points and robots are understood modulo $k + 1$.

Fix $\epsilon > 0$. Consider a strategy \mathcal{A} such that $I(\mathcal{A}) < 2s - \frac{\epsilon}{2}$. To prove the lemma, we will show there exists a point v_i , $0 \leq i \leq k$, such that the time between some two consecutive visits of a robot to point v_i is greater than $\tau = \frac{1}{k} - \epsilon$. By appropriately modifying the trajectories of the robots, we can convert a strategy \mathcal{A} into another strategy \mathcal{A}' so that the following properties are satisfied by \mathcal{A}' :

- (i) if a robot following \mathcal{A}' leaves some vital point v_i , then it does not reenter v_i before reaching some other vital point first (namely, v_{i-1} or v_{i+1}),
- (ii) no two robots following \mathcal{A}' ever meet,
- (iii) if a vital point is visited by a robot following strategy \mathcal{A} at time t , then it is visited by a robot following strategy \mathcal{A}' within the interval $[t - \frac{\epsilon}{4}, t + \frac{\epsilon}{4}]$.

By property (iii), if a point is unvisited by \mathcal{A}' in time interval $[t_1, t_2]$, then it is unvisited by \mathcal{A} in the time interval $[t_1 + \frac{\epsilon}{4}, t_2 - \frac{\epsilon}{4}]$. So, it now suffices to show that the time between some two consecutive visits of a robot following strategy \mathcal{A}' to point v_i is greater than $\tau + \frac{\epsilon}{2} = \frac{1}{k} - \frac{\epsilon}{2}$. Moreover, $I(\mathcal{A}') \leq I(\mathcal{A}) + \frac{\epsilon}{2} < 2s$. From now on we consider robots following \mathcal{A}' , only.

Since no two robots following \mathcal{A}' ever meet by (ii), we can denote an arbitrarily chosen robot by r_1 , and the other robots by r_2, \dots, r_k in clockwise order; this order never changes throughout the traversal. Suppose that at some time t , some robot r_j leaves point v_i on the arc towards point v_{i+1} . By (i), the next vital point it reaches has to be point v_{i+1} . Therefore, robot r_j cannot reenter point v_i before time $t + 2\text{dist}(v_i, v_{i+1}) \geq t + 2s > t + I(\mathcal{A}')$. It follows that within the time interval $[t, t + I(\mathcal{A}')$, robot r_{j-1} entered node v_i . Before this visit, the previous vital point visited by r_{j-1} must have been v_{i-1} . It follows that to each traversal of the arc (v_i, v_{i+1}) by robot r_j that starts at some time t , we can assign a distinct traversal of the arc (v_{i-1}, v_i) by robot r_{j-1} that ends within the time interval $[t, t + I(\mathcal{A}')$. Fix two values of time T_1 and T_2 , $0 \leq T_1 < T_2$. From now on, we will apply certain counting arguments within the time interval $[T_1, T_2]$; for the sake of intuition, one can see this interval as sufficiently long ($T_1 = 0$ and $T_2 \rightarrow \infty$). Let $C_j(i, i + 1)$ denote the number of traversals of arc (v_i, v_{i+1}) by robot r_j starting in the time interval $[T_1, T_2]$. Since only the first and last traversals of (v_i, v_{i+1}) by robot r_j within this time interval may be unmatched by corresponding traversals of (v_{i-1}, v_i) by robot r_{j-1} within the same time interval, we have $C_j(i, i + 1) - C_{j-1}(i - 1, i) \leq 2$. Let $C(i, i + 1) = \sum_{j=1}^k C_j(i, i + 1)$ be the total number of traversals of the arc (v_i, v_{i+1}) by all robots starting within the time interval $[T_1, T_2]$. Summing the above inequalities, we have: $C(i, i + 1) - C(i - 1, i) \leq 2k$, and analogously for the counter-clockwise direction, $C(i, i - 1) - C(i + 1, i) \leq 2k$. The above inequalities hold for all i , hence the difference between $C(i_1, i_1 + 1)$ and $C(i_2, i_2 + 1)$ (respectively, $C(i_1, i_1 - 1)$ and $C(i_2, i_2 - 1)$) can be bounded for any i_1 and i_2 on the cycle. Denoting $C^{cw} = \min_{0 \leq i \leq k} C(i, i + 1)$ and $C^{cc} = \min_{0 \leq i \leq k} C(i + 1, i)$, we have for all i :

$$C(i, i + 1) \leq C^{cw} + 2k^2, \quad C(i + 1, i) \leq C^{cc} + 2k^2.$$

Now, denote by $W_j(i) \geq 0$ the total time spent by robot r_j at point v_i within the time interval $[T_1, T_2]$, and let $W(i) = \sum_{j=1}^k W_j(i)$. Without loss of generality, let v_0 be a vital point with the minimal total waiting time, i.e., $W(0) = \min_{0 \leq i \leq k} W(i)$. With respect to point v_0 , the trajectory of each robot r_j within the time interval $[T_1, T_2]$ can be described by an ordered sequence of time moments $(e_j^1, l_j^1, e_j^2, l_j^2, \dots, l_j^{n_j})$, where e_j^p is the time at which robot r_j enters point v_0 for the p -th time, whereas l_j^p is the time at which robot r_j leaves point v_0 for the p -th time. We assume that $T_1 \leq e_j^1 \leq l_j^1 < e_j^2 \leq l_j^2 < \dots < e_j^{n_j} \leq l_j^{n_j} \leq T_2$,

where we put $e_j^1 = T_1$ if robot r_j was located at node v_0 at time T_1 , and $l_j^{n_j} = T_2$ if robot r_j was located at node v_0 at time T_2 . For the sake of notation, let $l_j^0 = T_1$ and $e_j^{n_j+1} = T_2$.

During the time interval $[T_1, T_2]$, point v_0 is covered by a robot during the set of moments X given as: $X = \bigcup_{j=1}^k \bigcup_{p=1}^{n_j} [e_j^p, l_j^p]$, such that $|X| = W(0)$. During the remaining time, i.e. $\bar{X} = [T_1, T_2] \setminus X$, no robot is located at v_0 . We observe that \bar{X} is a union of at most $1 + \sum_{j=1}^k n_j$ intervals. Hence, there exists some time interval of length τ , during which v_0 remains unvisited, where:

$$\tau \geq \frac{|\bar{X}|}{1 + \sum_{j=1}^k n_j} = \frac{(T_2 - T_1) - |X|}{1 + \sum_{j=1}^k n_j} = \frac{(T_2 - T_1) - W(0)}{1 + \sum_{j=1}^k n_j} \quad (1)$$

Notice that each robot r_j leaves point v_0 at least $n_j - 1$ times in the time interval $[T_1, T_2]$, going towards either point v_1 or point v_k . Thus, we have:

$$C(0, 1) + C(0, k) \geq \sum_{j=1}^k (n_j - 1) = \sum_{j=1}^k n_j - k.$$

Taking into account that $C(0, 1) \leq C^{cw} + 2k^2$ and $C(0, k) \leq C^{cc} + 2k^2$, we have:

$$\sum_{j=1}^k n_j \leq C^{cw} + C^{cc} + 4k^2 + k. \quad (2)$$

Moreover, since each arc of the cycle is traversed in either direction a total of at least $C^{cw} + C^{cc}$ times, the total distance covered by all the robots is at least $C^{cw} + C^{cc}$. Thus, the total time of movement of all k robots within the time interval $[T_1, T_2]$ is at least $C^{cw} + C^{cc}$, and we have the inequality $C^{cw} + C^{cc} + \sum_{i=0}^k W(i) \leq k(T_2 - T_1)$, which yields:

$$C^{cw} + C^{cc} \leq k(T_2 - T_1) - (k + 1)W(0). \quad (3)$$

Combining inequalities (1), (2), and (3), after several transformations we obtain $\tau > \frac{1}{k} - \frac{4k^2 + k + 1}{(T_2 - T_1) - (4k^2 + k + 1)}$.

Now, suppose that T_2 is chosen to be sufficiently large so that $\frac{4k^2 + k + 1}{(T_2 - T_1) - (4k^2 + k + 1)} < \frac{\epsilon}{2}$. We then have $\tau \geq \frac{1}{k} - \frac{\epsilon}{2}$, and so there exists a vital point on the cycle such that the time between some two successive visits of robots following \mathcal{A}' to this point is greater than $\frac{1}{k} - \frac{\epsilon}{2}$, which is what we needed to prove. \square

Lemma 3.6 ($2k + 1$ point lemma) *Let $(v_0, v_1, \dots, v_{2k})$ be a set of $2k + 1$ points chosen from vital regions of the terrain, arranged in clockwise order along the cycle, s.t. $\text{dist}(v_i, v_{(i+2) \bmod (2k+1)}) > \frac{1}{2k}$. The idleness $I(\mathcal{A})$ of any traversal strategy \mathcal{A} for k mobile robots in this terrain satisfies $I(\mathcal{A}) \geq \frac{1}{k}$.*

Proof. (sketch) Let $(v_0, v_1, \dots, v_{2k})$ be $2k + 1$ vital points chosen in accordance with the assumptions of the lemma. For the proof of the lower bound, we introduce the concept of a *shadow robot*, which can be seen as an auxiliary robot which temporarily appears in the system and assists robots in their patrolling task. More precisely, given a strategy \mathcal{A} , consider the trajectory of a robot r_j . Suppose that the robot leaves a vital point v_i at some time t_a , moves to an adjacent vital point $v_{i_1} \in \{v_{i-1}, v_{i+1}\}$ and then returns to point v_i at time t_b , without encountering any other vital points within the interval $[t_a, t_b]$. We say that a shadow robot r_j^{i*} is created at time t_a at point v_i , waits at v_i protecting it until time t_b , and then disappears. The addition of such a shadow robot, obviously, cannot increase the idleness of the strategy. The proof now proceeds by similar (though slightly more technical) arguments as the proof of Lemma 3.5, subject to the inclusion of shadow robots in the team of patrolling robots. The details are provided in the Appendix. \square

To complete the proof of Theorem 3.3, we consider an arbitrary terrain $C = (V, N)$. Let B be the length of the longest neutral interval of C , as defined in Lemma 3.3. We have two cases to consider.

- If $B \geq 1/(2k)$, then by clause (1) of Lemma 3.4, there exists a subset of $k + 1$ vital points $\{v_0, \dots, v_k\} \subseteq V$ such that for these points, in Lemma 3.5 we have $s = \min\{1/(2k), L\}$. Now, by Lemma 3.5 we obtain that for any strategy \mathcal{A} , the idleness is lower bounded by $I(\mathcal{A}) \geq \min\{1/k, 2s\} = \min\{1/k, 2L\}$.

- If $B < 1/(2k)$, then by clause (2) of Lemma 3.4, there exists a subset of $2k + 1$ vital points $\{v_0, \dots, v_{2k}\} \subseteq V$ that satisfies the assumptions of Lemma 3.6. Thus, by Lemma 3.6 we obtain that for any strategy \mathcal{A} , the idleness is lower bounded by $I(\mathcal{A}) \geq 1/k$.

In either case, we obtain that the idleness of any strategy patrolling C is at least $\min\{1/k, 2L\}$, which proves the claim of the Theorem.

4 Computing Optimal Robot Trajectories

Let $C = (V, N)$ be the unit segment $[0, 1]$ with vital and neutral regions. Assume w.l.o.g. that the vital intervals in C are arranged in a data structure from left to right as $V_i = [b_i, e_i]$, for $i = 1, 2, \dots, n$ where $b_1 = 0, b_i \leq e_i < b_{i+1}$. We assume that arithmetic operations involving these values can be performed in unit time.

Recall that in this case the solution is based on the use of lids, where with each lid we associate a different robot. We show first that one can test in time $O(\min\{n, k \log n\})$ whether for a collection of k lids each of length d can cover all vital points in $[0, 1]$.

We propose a recursive procedure $TestLidSize(k, d, p)$ that operates on sub-intervals of the form $[p, 1]$ of C , where k stands for the number of available lids and $d \leq 1$ refers to the uniform length of the lids. The procedure returns value *true* if all vital points in C can be covered with the collection of k lids. Otherwise the returned value is *false*.

Procedure $TestLidSize(k, d, p)$: {true,false};

1. Use the next available lid to cover segment $[p, p + d]$;
2. **if** $(p + d) \geq e_n$ **then** $return(true)$; /* all vital points are covered */
3. $p^* = \inf\{p' \in V : p' > p + d\}$; /* p^* exists since $p + d < e_n$. */
4. **if** $(k > 1)$ **then** $return(TestLidSize(k - 1, d, p^*))$
else $return(false)$;

Lemma 4.1 For any positive integer $k, d > 0$, and $p \in C = (V, N)$, such that p is vital, procedure $TestLidSize(k, d, p)$ verifies in time $O(\min\{n, k \log n\})$ whether all vital points in $[p, e_n]$ can be covered by k lids of length d .

We now show how to efficiently compute the optimal (minimal) size of the lid. We will need the following lemma.

Lemma 4.2 If L is the optimal (minimal) size of lids, there must be some integer $1 \leq l \leq k$, s.t. $L = \frac{e_j - b_i}{l}$, for some $1 \leq i \leq j \leq n$.

We now present the algorithm that computes the optimal size of lids. Using Lemma 4.2, one can observe that we need to test at most $O(kn^2)$ values in search for the optimal size of lids. We can sort these values in time $O(kn^2 \log n)$ and later use binary search to find the optimal value. The number of tests during the binary search is $O(\log n)$ and the cost of each test is $O(\min\{n, k \log n\})$, see Lemma 4.1. Thus the total complexity is dominated by sorting performed in time $O(kn^2 \log n)$.

Observation 4.3 The optimal size of lids can be computed in time $O(kn^2 \log n)$. □

The complexity of this algorithm can be further improved if we use an implicit representation of $O(kn^2)$ candidates based on values $\frac{e_j - b_i}{l}$, for $1 \leq i \leq j \leq n$ and $1 \leq l \leq k$, and perform search for the

optimal size of lids in a more sophisticated fashion. We perform search among values based on each l separately.

Let M_i^l represent implicitly the list of values $(\frac{e_i-b_i}{l}, \frac{e_{i+1}-b_i}{l}, \dots, \frac{e_n-b_i}{l})$, for $1 \leq i \leq n$. Each list M_i^l contains at most n values. Any value from this list can be calculated on the basis of the sequence e_1, \dots, e_n , where values e_1 through e_n are stored in an array of length n . In particular, using this representation one can calculate the value of any requested element in M_i^l in constant time.

The search algorithm operates in rounds, processing all M_i^l s, with $1 \leq i \leq n$ and $1 \leq l \leq k$, simultaneously. At the beginning all entries in M_i^l are potential candidates for being the optimal length of the lids. During each round the list of candidates in half of M_i^l s is reduced by half but the remaining candidates in M_i^l always form a sublist of consecutive elements in M_i^l . The reduction process in each round is performed as follows.

Procedure FastLidSearch (C): $\{true, false\}$;
repeat in consecutive rounds

1. Find the medians m_i^l among the remaining candidates in each non-empty list M_i^l .
2. Find the median m^* among all m_i^l , ranging over indices $1 \leq i \leq n, 1 \leq l \leq k$ corresponding to non-empty M_i^l s.
3. Use procedure *TestLidSize* to test whether m^* is long enough.
4. And if yes, reduce by half the content of lists M_i^l with $m_i^l > m^*$.
5. Otherwise, reduce by half the content of lists M_i^l with $m_i^l < m^*$.

until only one candidate value is left.

Note that the cost of each round, in which there are $a \leq kn$ non-empty lists M_i^l , can be bounded by $O(a + \min\{n, k \log n\})$. We will now compute the bound on the total number of rounds. At the start we associate with each list M_i^l a potential of $\log n$. This means that the combined potential of all lists is $kn \log n$. During each round the potential of half of the lists is reduced by 1. Eventually some lists M_i^l become empty which is reflected in the null potential. Note that until at least $\frac{kn}{2}$ lists are non-empty the combined potential is reduced by $\frac{kn}{4}$ during each round. Thus the number of rounds with at least $\frac{kn}{2}$ non-empty lists is limited by $(kn \log n) / (\frac{kn}{4}) = 4 \log n$, and the total duration of these rounds is $O(\log n \cdot (kn + \min\{n, k \log n\}))$. Furthermore, we reduce the number of non-empty lists from $\frac{kn}{2}$ to $\frac{nk}{4}$ also in at most $(\frac{kn}{2} \log n) / (\frac{kn}{8}) = 4 \log n$ rounds, and the total duration of these rounds is $O(\log n \cdot (\frac{kn}{2} + \min\{n, k \log n\}))$. Thus, if we continue this process until only one element in one list is left, the total time of execution can be bounded as:

$$O\left(\sum_{j=0}^{\log(kn)} (\log n \cdot (\frac{kn}{2^j} + \min\{n, k \log n\}))\right) = O(kn \log n + \log n (\log n + \log k) \min\{n, k \log n\}) = O(kn \log n + k \log^3 n + n \log k \log n) = O(kn \log n).$$

Corollary 4.4 *The optimal size of lids can be computed in time $O(kn \log n)$.* □

This approach is also applicable to the combined strategy on the cycle, since, in fact, the optimal lid size only needs to be computed in the case when the cycle contains some neutral region N_i of length at least $1/(2k)$. Then, the problem on the cycle C reduces to that on the closed segment $C \setminus N_i$. We have the following:

Theorem 4.5 *Consider k robots patrolling a boundary cycle (resp., segment) with n vital regions. The robot trajectories which result in minimal idleness can be described using the combined strategy (resp., the partition strategy). Such a description can be computed using an $O(kn \log n)$ algorithm.* □

References

- [1] A. Aggarwal. *The Art Gallery Theorem and Algorithm*. PhD thesis, Johns Hopkins University, 1984.
- [2] N. Agmon, S. Kraus, and G. A. Kaminka. Multi-robot perimeter patrol in adversarial settings. In *ICRA*, pages 2339–2345, 2008.
- [3] A. Almeida, G. Ramalho, H. Santana, P. A. Tedesco, T. Menezes, V. Corruble, and Y. Chevaleyre. Recent advances on multi-agent patrolling. In *SBIA*, pages 474–483, 2004.
- [4] E. M. Arkin, J. S. B. Mitchell, and C. D. Piatko. Minimum-link watchman tours. *IPL*, 86:203–207, 2003.
- [5] S. Carlsson, B. J. Nilsson, and S. C. Ntafos. Optimum guard covers and m -watchmen routes for restricted polygons. *Int. J. Comp. Geom. Appl.*, 3:85–105, 1993.
- [6] Y. Chevaleyre. Theoretical analysis of the multi-agent patrolling problem. In *IAT*, pages 302–308, 2004.
- [7] W. Chin and S. C. Ntafos. Optimal watchman routes. *Inf. Proc. Lett.*, 28:39–44, 1988.
- [8] W. Chin and S. C. Ntafos. Shortest watchman routes in simple polygons. *Discr. Comp. Geom.*, 6:9–31, 1991.
- [9] W. Chin and S. C. Ntafos. The zookeeper route problem. *Inf. Sci.*, 63:245–259, 1992.
- [10] J. Czyzowicz, P. Egyed, H. Everett, D. Rappaport, T. C. Shermer, D. L. Souvaine, G. T. Toussaint, and J. Urrutia. The aquarium keeper’s problem. In *SODA*, pages 459–464, 1991.
- [11] J. Czyzowicz, L. Gąsieniec, A. Kosowski, and E. Kranakis. Boundary patrolling by mobile agents with distinct maximal speeds. In *ESA 2011*, pages 701–712. Springer, 2011.
- [12] M. Dror, A. Efrat, A. Lubiw, and J. S. B. Mitchell. Touring a sequence of polygons. In *STOC*, pages 473–482, 2003.
- [13] A. Dumitrescu and C. Tóth. Watchman tours for polygons with holes. *Comput. Geom. Theory Appl.*, 45:326–333, 2012.
- [14] K. Easton and J. W. Burdick. A coverage algorithm for multi-robot boundary inspection. In *ICRA*, pages 727–734. IEEE, 2005.
- [15] S. Eidenbenz, C. Stamm, and P. Widmayer. Inapproximability results for guarding polygons and terrains. *Algorithmica*, 31(1):79–113, 2001.
- [16] Y. Elmaliach, N. Agmon, and G. A. Kaminka. Multi-robot area patrol under frequency constraints. *Ann. Math. Artif. Intell.*, 57(3-4):293–320, 2009.
- [17] Y. Elmaliach, A. Shiloni, and G. A. Kaminka. A realistic model of frequency-based multi-robot polyline patrolling. In *AAMAS (1)*, pages 63–70, 2008.
- [18] Y. Elor and A. M. Bruckstein. Autonomous multi-agent cycle based patrolling. In *ANTS*, pages 119–130, 2010.
- [19] F. V. Fomin, P. A. Golovach, A. Hall, M. Mihalák, E. Vicari, and P. Widmayer. How to guard a graph? *Algorithmica*, 61(4):839–856, 2011.

- [20] F. V. Fomin and D. M. Thilikos. An annotated bibliography on guaranteed graph searching. *Theor. Comput. Sci.*, 399(3):236–245, 2008.
- [21] Y. Gabriely and E. Rimon. Spanning-tree based coverage of continuous areas by a mobile robot. In *ICRA*, pages 1927–1933, 2001.
- [22] S. K. Ghosh. Approximation algorithms for art gallery problems in polygons and terrains. In *WALCOM*, pages 21–34, 2010.
- [23] N. Hazon and G. A. Kaminka. On redundancy, efficiency, and robustness in coverage for multiple robots. *Rob. and Autonom. Syst.*, 56:1102–1114, 2008.
- [24] G. D. Kazazakis and A. A. Argyros. Fast positioning of limited-visibility guards for the inspection of 2d workspaces. In *IEEE Int. Conf. on Intelligent Robots and Systems, Vol. 3*, pages 2843–2848, 2002.
- [25] A. Machado, G. Ramalho, J.-D. Zucker, and A. Drogoul. Multi-agent patrolling: An empirical analysis of alternative architectures. In *MABS*, pages 155–170, 2002.
- [26] A. Marino, L. E. Parker, G. Antonelli, and F. Caccavale. Behavioral control for multi-robot perimeter patrol: A finite state automata approach. In *ICRA*, pages 831–836, 2009.
- [27] J. Mitchell. Geometric shortest paths and network optimization. In *Handbook of Computational Geometry, Chapter 15*, pages 633–701, 2000.
- [28] B. J. Nilsson. *Guarding art galleries; Methods for mobile guards*. PhD thesis, Lund U., Sweden, 1995.
- [29] B. J. Nilsson and D. Wood. Optimum watchmen route in spiral polygons. In *2nd Canadian Conf. Comput. Geometry*, pages 269–272, 1990.
- [30] S. C. Ntafos. Watchman routes under limited visibility. *Comput. Geom.*, 1:149–170, 1991.
- [31] J. O’Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, 1987.
- [32] E. Packer. Computing multiple watchman routes. In *WEA’08 Proc. 7th Workshop on Experimental Algorithms*, pages 114–128, 2008.
- [33] F. Pasqualetti, A. Franchi, and F. Bullo. On optimal cooperative patrolling. In *CDC*, 2010.
- [34] T. C. Shermer. Recent results in art galleries. In *Proc. of the IEEE*, volume 80, pages 1384–1399, 1992.
- [35] M. Yamashita, H. Umemoto, I. Suzuki, and T. Kameda. Searching for mobile intruders in a polygonal region by a group of mobile searchers. *Algorithmica*, 31:208–236, 2001.
- [36] V. Yanovski, I. A. Wagner, and A. M. Bruckstein. A distributed ant algorithm for efficiently patrolling a network. *Algorithmica*, 37(3):165–186, 2003.

APPENDIX

Proof of Theorem 2.1

To prove the theorem, we start by showing a property of a greedy cover of the segment with lids.

Lemma .6 *Suppose that L is the minimum value such that a given terrain C admits a (L, k) -lid cover; and let $L' < L$. Then, the region $C \setminus [0, L']$ does not admit a $(L', k - 1)$ -lid cover.*

Proof. Indeed, if the region $C \setminus [0, L']$ admitted a $(L', k - 1)$ -lid cover, then one could obtain a (L', k) -lid cover of C simply by adding lid $[0, L']$ to the lid cover obtained for the region $C \setminus [0, L']$. This contradicts the minimality of L . \square

Lemma .7 *Any patrolling strategy may be converted to a strategy, achieving the same idleness, for which the relative order of the mobile robots on the segment is maintained throughout the traversal.*

Proof. The proof is based on the simple observation that when two mobile robots meet while moving in opposite directions they can “exchange” roles, so that the coverage of the points on the segment by one robot is the same as coverage by the other. Since after this change of roles the set of visited nodes at any time remains the same, this does not affect the idleness of visited nodes. \square

The proof of the claim $I \geq 2L$ now proceeds by induction on the number of robots. It is clearly true for $k = 1$, since the idleness of the strategy cannot be smaller than twice the distance between the extremal vital points C , which corresponds precisely to the size of the smallest lid cover.

Suppose, by contradiction, that there exists a value of k and some terrain C such that the idleness of some patrolling strategy \mathcal{A} is $I = 2L' < 2L$. By Lemma .7, w.l.o.g. we can assume that the robots never change places along the segment. Consider the trajectory of the leftmost robot r_1 following strategy \mathcal{A} . Let c be the supremum of all points along the segment reached by robot r_1 . If point c is reached by robot r_1 at some time t , its last visit to point 0 must have been no later than at time $t - c$, and the next visit to point 0 will take place at time not earlier than $t + c$. From Lemma .7, we have that point 0 is never visited by any robot when it is not visited by r_1 . Consequently, we must have $2c \leq 2L'$, and so $c \leq L'$. It follows that the region $C \setminus [0, L']$ must be patrolled solely by the set of $k - 1$ robots, without the help of r_1 , with idleness at most $2L'$. From the inductive assumption, we have that $C \setminus [0, L']$ admits a L' -lid cover. This is a contradiction, by Lemma .6. This completes the proof of Theorem 2.1.

Proof of Lemma 3.4

Proof. To prove clause (1), let v_0 be the first vital point located at the clockwise endpoint of a neutral region of length B . Fix $\epsilon > 0$ and consider the set of points chosen iteratively as follows: let v_{i+1} be the first vital point located at arc distance not less than $L - \epsilon$ from v_i , moving in the clockwise direction. We claim that point v_k is reached before completing one full rotation around the cycle, starting from v_0 . Indeed, if this were not the case, then there would exist a set of k lids: $[v_0, v_0 + L - \epsilon], \dots, [v_{k-1}, v_{k-1} + L - \epsilon]$, covering the whole of V , a contradiction with the minimality of lid cover size L . Finally, note that the distance between points v_0 and v_k is at least $B \geq 1/(2k)$. Parametrizing each of the points v_i as $v_i(\epsilon)$, it follows that for any $\epsilon > 0$, we can find a set of $k + 1$ points $(v_0(\epsilon), v_1(\epsilon), \dots, v_k(\epsilon))$ such that $\min_{0 \leq i \leq k} \text{dist}(v_i(\epsilon), v_{(i+1) \bmod (k+1)}(\epsilon)) \geq \min\{1/(2k), L\} - \epsilon$. By taking into account that the set of vital points V is a closed set and $v_i(\epsilon)$ is non-decreasing and bounded (w.r.t. shifts in the clockwise direction) for any sequence $\epsilon \searrow 0$, we converge to a sequence of vital points $(v_0(0), \dots, v_k(0))$ satisfying clause (1).

To prove clause (2), we will show a slightly stronger version “(2’)” of this clause in which we replace the assumption “ $B < 1/(2k)$ ” by “ $B \leq 1/(2k)$ ”. Suppose that terrain (V, N) is a counterexample

to the claim of (2'), such that for any other terrain (V', N') which violates clause (2') it holds that $V' \subsetneq V$. (Such an inclusion-wise minimal counterexample always exists, since the set of vital points is by assumption a closed set.) By the minimality of V , the set of its vital points must be discrete, say, $V = \{u_0, \dots, u_{n-1}\}$. Suppose that for some $0 \leq i \leq n$, we have $\text{dist}(u_i, u_{(i+2) \bmod n}) \leq 1/(2k)$. The terrain $(V \setminus \{u_{i+1}\}, N \cup \{u_{i+1}\})$ has no neutral intervals of length greater than $1/(2k)$, and thus is a smaller counterexample to our claim, a contradiction. It follows that $\min_{0 \leq i < n} \text{dist}(u_i, u_{(i+2) \bmod n}) > 1/(2k)$. Since for all $0 \leq i < n$, $\text{dist}(u_i, u_{(i+1) \bmod n}) < 1/(2k)$, we must have $n \geq 2k + 1$. So, choosing points $\{v_0, \dots, v_{2k}\}$ as $v_i = u_i$, for all $0 \leq i \leq 2k$, we find in V the subset of vital points satisfying clause (2'). So, V cannot be a counter-example to the claim. \square

Proof of Lemma 3.5

Proof. Let (v_0, v_1, \dots, v_k) be $k + 1$ vital points chosen so that $\text{dist}(v_i, v_{(i+1)}) \geq s$, for all $0 \leq i \leq k$. Throughout the proof, indices of points and robots are understood modulo $k + 1$, i.e. $v_i \equiv v_{i \bmod (k+1)}$ and $r_j \equiv r_{j \bmod (k+1)}$. We will show that the claim holds even if $\{v_0, v_1, \dots, v_k\}$ are the only vital points of the cycle.

Fix $\epsilon > 0$. Consider a strategy \mathcal{A} such that $I(\mathcal{A}) < 2s - \frac{\epsilon}{2}$. To prove the lemma, we will show that there exists a point v_i , $0 \leq i \leq k$, such that the time between some two consecutive visits of a robot to point v_i is greater than $\tau = \frac{1}{k} - \epsilon$.

Without loss of generality, by modifying the trajectories of the robots, we can convert a strategy \mathcal{A} into another strategy \mathcal{A}' so that the following properties are satisfied by \mathcal{A}' :

- (i) if a robot following \mathcal{A}' leaves some vital point v_i , then it does not reenter this vital point before reaching some other vital point first (namely, v_{i-1} or v_{i+1}),
- (ii) no two robots following \mathcal{A}' ever meet,
- (iii) if a vital point is visited by a robot following strategy \mathcal{A} at time t , then it is visited by a robot following strategy \mathcal{A}' within the interval $[t - \frac{\epsilon}{4}, t + \frac{\epsilon}{4}]$.

For completeness, we outline the technical steps which are required to perform the above conversion. First, property (i) is achieved by modifying the trajectories of the robots in neutral regions, only. Next, properties (ii) and (iii) can be ensured by first converting the strategy to one which preserves the ordering of the robots as in Lemma .6, and then delaying the movements of some of the robots to avoid meetings, without changing the time intervals during which a vital point is occupied by more than $\frac{\epsilon}{4}$.

By property (iii), if a point is unvisited by \mathcal{A}' in time interval $[t_1, t_2]$, then it is unvisited by \mathcal{A} in the time interval $[t_1 + \frac{\epsilon}{4}, t_2 - \frac{\epsilon}{4}]$. It now suffices to show that the time between some two consecutive visits of a robot following strategy \mathcal{A}' to point v_i is greater than $\tau + \frac{\epsilon}{2} = \frac{1}{k} - \frac{\epsilon}{2}$. Moreover, $I(\mathcal{A}') \leq I(\mathcal{A}) + \frac{\epsilon}{2} < 2s$. From now on we consider robots following \mathcal{A}' , only.

Since no two robots following \mathcal{A}' ever meet by (ii), we can denote an arbitrarily chosen robots by r_1 , and the other robots by r_2, \dots, r_k in clockwise order; this order never changes throughout the traversal.

Suppose that at some time t , some robot r_j leaves point v_i on the arc towards point v_{i+1} . By (i), the next vital point it reaches has to be point v_{i+1} . Therefore, robot r_j cannot reenter point v_i before time $t + 2\text{dist}(v_i, v_{i+1}) \geq t + 2s > t + I(\mathcal{A}')$. So, some other robot must visit point v_i in between the two visits by robot r_j . Since the robots never meet, it follows that within the time interval $[t, t + I(\mathcal{A}')$, robot r_{j-1} entered node v_i . Before this visit, the previous vital point visited by r_{j-1} must have been v_{i-1} . It follows that to each traversal of the arc (v_i, v_{i+1}) by robot r_j that starts at some time t , we can assign a distinct traversal of the arc (v_{i-1}, v_i) by robot r_{j-1} that ends within the time interval $[t, t + I(\mathcal{A}')$. Fix two values of time T_1 and T_2 , $0 \leq T_1 < T_2$. From now on, we will apply certain counting arguments within the time interval $[T_1, T_2]$. Let us denote by $C_j(i, i + 1)$ the number of traversals of arc (v_i, v_{i+1}) by robot r_j starting in the time interval $[T_1, T_2]$. Since only the first and last traversals of (v_i, v_{i+1}) by

robot r_j within this time interval may be unmatched by corresponding traversals of (v_{i-1}, v_i) by robot r_{j-1} within the same time interval, we have:

$$C_j(i, i+1) - C_{j-1}(i-1, i) \leq 2.$$

Let $C(i, i+1) = \sum_{j=1}^k C_j(i, i+1)$ be the total number of traversals of the arc (v_i, v_{i+1}) by all robots starting within the time interval $[T_1, T_2]$. Summing the above inequalities, we have:

$$C(i, i+1) - C(i-1, i) \leq 2k.$$

An analogous analysis can be performed for the counter-clockwise direction, i.e., considering values of the form $C(i+1, i)$, corresponding to traversal of the arc from v_{i+1} to v_i . We obtain:

$$C(i, i-1) - C(i+1, i) \leq 2k.$$

In general, by iterating the above around the cycle, for any two vital points v_{i_1} and v_{i_2} we obtain:

$$C(i_1, i_1+1) - C(i_2, i_2+1) \leq 2k^2.$$

Denoting by $C^{cw} = \min_{0 \leq i \leq k} C(i, i+1)$, we have for any i :

$$C(i, i+1) \leq C^{cw} + 2k^2.$$

An analogous analysis can be performed for the counter-clockwise direction, i.e., considering values of the form $C(i+1, i)$, corresponding to traversal of the arc from v_{i+1} to v_i . Consequently, denoting $C^{cc} = \min_{0 \leq i \leq k} C(i+1, i)$, we have for any i :

$$C(i+1, i) \leq C^{cc} + 2k^2.$$

Now, denote by $W_j(i) \geq 0$ the total time spent by robot r_j at point v_i within the time interval $[T_1, T_2]$, and let $W(i) = \sum_{j=1}^k W_j(i)$. Without loss of generality, let v_0 be a vital point with the minimal total waiting time, i.e., $W(0) = \min_{0 \leq i \leq k} W(i)$.

With respect to point v_0 , the trajectory of each robot r_j within the time interval $[T_1, T_2]$ can be described by an ordered sequence of time moments $(e_j^1, l_j^1, e_j^2, l_j^2, \dots, e_j^{n_j}, l_j^{n_j})$, where e_j^p is the time at which robot r_j enters point v_0 for the p -th time, whereas l_j^p is the time at which robot r_j leaves point v_0 for the p -th time. We assume that $T_1 \leq e_j^1 \leq l_j^1 < e_j^2 \leq l_j^2 < \dots < e_j^{n_j} \leq l_j^{n_j} \leq T_2$, where we put $e_j^1 = T_1$ if robot r_j was located at node v_0 at time T_1 , and $l_j^{n_j} = T_2$ if robot r_j was located at node v_0 at time T_2 . For the sake of notation, let $l_j^0 = T_1$ and $e_j^{n_j+1} = T_2$.

During the time interval $[T_1, T_2]$, point v_0 is covered by a robot during the set of moments X given as:

$$X = \bigcup_{j=1}^k \bigcup_{p=1}^{n_j} [e_j^p, l_j^p],$$

such that $|X| = W(0)$. During the remaining time, i.e. $\bar{X} = [T_1, T_2] \setminus X$, no robot is located at v_0 . We observe that \bar{X} is a union of at most $1 + \sum_{j=1}^k n_j$ intervals. Hence, there exists some time interval of length τ :

$$\tau \geq \frac{|\bar{X}|}{1 + \sum_{j=1}^k n_j} = \frac{(T_2 - T_1) - |X|}{1 + \sum_{j=1}^k n_j} = \frac{(T_2 - T_1) - W(0)}{1 + \sum_{j=1}^k n_j} \quad (4)$$

during which v_0 remains unvisited.

Notice that each robot r_j leaves point v_0 at least $n_j - 1$ times in the time interval $[T_1, T_2]$, going towards either point v_1 or point v_k . Thus, we have:

$$C(0, 1) + C(0, k) \geq \sum_{j=1}^k (n_j - 1) = \sum_{j=1}^k n_j - k.$$

Taking into account that $C(0, 1) \leq C^{cw} + 2k^2$ and $C(0, k) \leq C^{cc} + 2k^2$, we have:

$$\sum_{j=1}^k n_j \leq C^{cw} + C^{cc} + 4k^2 + k. \quad (5)$$

Moreover, since each arc of the cycle is traversed in either direction a total of at least $C^{cw} + C^{cc}$ times, the total distance covered by all the robots is at least $C^{cw} + C^{cc}$. Thus, the total time of movement of all k robots within the time interval $[T_1, T_2]$ is at least $C^{cw} + C^{cc}$, and we have the following inequality:

$$\begin{aligned} C^{cw} + C^{cc} + \sum_{i=0}^k W(i) &\leq k(T_2 - T_1) \\ C^{cw} + C^{cc} &\leq k(T_2 - T_1) - (k + 1)W(0). \end{aligned} \quad (6)$$

Combining inequalities (4), (5), and (6), we obtain:

$$\begin{aligned} \tau &\geq \frac{(T_2 - T_1) - W(0)}{k(T_2 - T_1) - (k + 1)W(0) + 4k^2 + k + 1} \geq \frac{1}{k + \frac{4k^2 + k + 1 - W(0)}{(T_2 - T_1) - W(0)}} \geq \\ &\geq \frac{1}{k + \frac{4k^2 + k + 1}{(T_2 - T_1) - (4k^2 + k + 1)}} > \frac{1}{k} - \frac{4k^2 + k + 1}{(T_2 - T_1) - (4k^2 + k + 1)}. \end{aligned}$$

In the above, we assumed that $(T_2 - T_1) - W(0) > 0$, i.e., there cannot be a robot covering v_0 throughout the time interval $[T_1, T_2]$. This is true, since otherwise, taking into account that $W(i) \geq W(0)$ for all $1 \leq i \leq k$, all $k + 1$ points would have to be covered by a robot throughout $[T_1, T_2]$, and there are only k robots, a contradiction.

Now, suppose that T_2 is chosen to be sufficiently large so that $\frac{4k^2 + k + 1}{(T_2 - T_1) - (4k^2 + k + 1)} < \frac{\epsilon}{2}$. We then have $\tau \geq \frac{1}{k} - \frac{\epsilon}{2}$, and so there exists a vital point on the cycle such that the time between some two successive visits of robots following \mathcal{A}' to this point is greater than $\frac{1}{k} - \frac{\epsilon}{2}$. This completes the proof of Lemma 3.5. \square

Proof of Lemma 3.6

Proof. Let $(v_0, v_1, \dots, v_{2k})$ be $2k + 1$ vital points chosen in accordance with the assumptions of the lemma. For the proof of the lower bound, we introduce the concept of a *shadow robot*, which can be seen as an auxiliary robot which temporarily appears in the system and assists robots in their patrolling task. More precisely, given a strategy \mathcal{A} , consider the trajectory of a robot r_j . Suppose that the robot leaves a vital point v_i at some time t_a , moves to an adjacent vital point $v_{i_1} \in \{v_{i-1}, v_{i+1}\}$ and then returns to point v_i at time t_b , without encountering any other vital points within the interval $[t_a, t_b]$. We say that a shadow robot r_j^{i*} is created at time t_a at point v_i , waits at v_i protecting it until time t_b , and then disappears. The addition of such a shadow robot, obviously, cannot increase the idleness of the strategy.

Observe that one robot can create at most two shadow robots at a time: when r_j is located anywhere within a closed arc $[v_i, v_{i+1}]$, then it may only have the shadow robots r_j^{i*} and $r_j^{(i+1)*}$. Robot r_j and its shadow robots can wait at not more than two vital points simultaneously.

Will show that the claim holds even if $\{v_0, v_1, \dots, v_{2k}\}$ are the only vital points of the cycle. The rest of the proof proceeds analogously to the proof of Lemma 3.5, subject to the inclusion of shadow robots in the team of robots patrolling the terrain. Once again, for a fixed $\epsilon > 0$, we modify the trajectories of the robots, converting any strategy \mathcal{A} into another strategy \mathcal{A}' fulfilling the following properties:

- (i) if a robot following \mathcal{A}' leaves some vital point v_i , then it does not reenter this vital point before reaching some other vital point first (namely, v_{i-1} or v_{i+1}),
- (ii) no two robots following \mathcal{A}' ever meet each other or the shadow robots of other robots,
- (iii) if a point P is visited by a robot following strategy \mathcal{A} at time t , then it is visited by a robot or shadow robot following strategy \mathcal{A}' within the interval $[t - \frac{\epsilon}{4}, t + \frac{\epsilon}{4}]$.

From now on we consider robots (r_1, r_2, \dots, r_k) and their shadow robots following \mathcal{A}' , only, and proceed to perform a modification of the proof of Lemma 3.5 which takes shadow robots into account.

Suppose that $I(\mathcal{A}') < 1/k - \epsilon$. We will call a traversal of the directed arc (v_i, v_{i+1}) by robot r_j *shadowless* if after arriving at v_{i+1} , the next vital point visited by robot r_j is v_{i+2} (not v_i). Equivalently, a traversal of arc (v_i, v_{i+1}) is shadowless if r_j does not leave its shadow robot r_j^{i*} at v_i during this traversal.

Fix a time interval $[T_1, T_2]$. For all $0 \leq i \leq 2k$, we will denote by $C_j(i, i+2)$ the number of shadowless traversals of the directed arc (v_i, v_{i+1}) by robot r_j starting in this time interval. Suppose that robot r_j initiates a shadowless traversal at some time t . Since $\text{dist}(v_i, v_{i+2}) > \frac{1}{2k}$ by assumption, we have that the next visit of this robot to v_i takes place after time $t + 1/k > t + I(\mathcal{A}')$. Since v_i is not occupied by a shadow robot, the robot r_{j-1} must arrive at point v_i at some time $t' \in (t, t + 1/k)$. The previous vital point occupied by robot r_{j-1} before t' must have been v_{i-1} . Before that, the robot cannot have occupied vital point v_i , since then, during its traversal from v_i to v_{i-1} and back to v_i , the shadow robot r_{j-1}^{i*} would have existed at v_i . This shadow robot must have met robot r_i at point v_i at time t , which contradicts the assumption that robots and shadow robots do not meet. It follows that before arriving at v_{i-1} robot r_{j-1} must have been located at v_{i-2} . Thus, robot r_{j-1} was performing a shadowless traversal of (v_{i-2}, v_{i-1}) . This traversal counts towards $C_{j-1}(i-2, i)$ if robot r_j left v_{i-2} within the interval $[T_1, T_2]$. Following the reasoning from Lemma 3.5, we obtain the following bound:

$$C_j(i, i+2) - C_{j-1}(i-2, i) \leq 2.$$

Summing the above inequalities over all robots, and performing the analysis for the counter-clockwise direction we have:

$$C(i, i+2) - C(i-2, i) \leq 2k.$$

$$C(i, i-2) - C(i+2, i) \leq 2k.$$

Since the number of points $2k+1$ is odd, by iterating the above around the cycle at most $2k$ times in one direction, for any two vital points v_{i_1} and v_{i_2} we obtain:

$$C(i_1, i_1+2) - C(i_2, i_2+2) \leq 2k^2.$$

$$C(i_1, i_1-2) - C(i_2, i_2-2) \leq 2k^2.$$

Denoting $C^{cw} = \min_{0 \leq i \leq 2k} C(i, i+2)$ and $C^{cc} = \min_{0 \leq i \leq 2k} C(i+2, i)$, we have for any i :

$$C(i, i+2) \leq C^{cw} + 2k^2.$$

$$C(i+2, i) \leq C^{cc} + 2k^2.$$

Now, denote by $W_j(i) \geq 0$ the total time spent by robot r_j at point v_i and by its shadow r_j^{i*} within the time interval $[T_1, T_2]$, and let $W(i) = \sum_{j=1}^k W_j(i)$. Without loss of generality, let v_0 be a vital point with the minimal total waiting time, i.e., $W(0) = \min_{0 \leq i \leq 2k} W(i)$.

With respect to point v_0 , the trajectory of each robot r_j within the time interval $[T_1, T_2]$ can be described by an ordered sequence of time moments $(e_j^1, l_j^1, e_j^2, l_j^2, \dots, l_j^{n_j})$, where e_j^p is the time at which robot r_j enters point v_0 for the p -th time after a shadowless traversal (of arc (v_{2k-1}, v_{2k}) or (v_2, v_1)), whereas l_j^p is the time at which robot r_j leaves point v_0 starting its p -th shadowless traversal (of arc (v_0, v_1) or (v_0, v_{2k})). We assume that $T_1 \leq e_j^1 \leq l_j^1 < e_j^2 \leq l_j^2 < \dots < e_j^{n_j} \leq l_j^{n_j} \leq T_2$, and make the same boundary assumptions as in the proof of Lemma 3.5.

During the time interval $[T_1, T_2]$, point v_0 is covered by some robot or some shadow robot during the set of moments X given as:

$$X = \bigcup_{j=1}^k \bigcup_{p=1}^{n_j} [e_j^p, l_j^p],$$

such that $|X| = W(0)$. During the remaining time, i.e. $\bar{X} = [T_1, T_2] \setminus X$, no robot and no shadow robot is located at v_0 . We observe that \bar{X} is a union of at most $1 + \sum_{j=1}^k n_j$ intervals. Hence, there exists some time interval of length τ during which point v_0 remains unvisited, lower-bounded by an inequality of the same form as (1):

$$\tau \geq \frac{|\bar{X}|}{1 + \sum_{j=1}^k n_j} = \frac{(T_2 - T_1) - |X|}{1 + \sum_{j=1}^k n_j} = \frac{(T_2 - T_1) - W(0)}{1 + \sum_{j=1}^k n_j} \quad (7)$$

Notice that each robot r_j leaves point v_0 at least $n_j - 1$ times in the time interval $[T_1, T_2]$, embarking on a shadowless traversal of the arc either to point v_1 (and then to v_2) or to point v_{2k} (and then to v_{2k-1}). Thus, we have:

$$C(0, 2) + C(0, 2k - 1) \geq \sum_{j=1}^k (n_j - 1) = \sum_{j=1}^k n_j - k.$$

Taking into account that $C(0, 2) \leq C^{cw} + 2k^2$ and $C(0, 2k - 1) \leq C^{cc} + 2k^2$, we have:

$$\sum_{j=1}^k n_j \leq C^{cw} + C^{cc} + 4k^2 + k. \quad (8)$$

For any robot r_j , we trace its trajectory within the time interval $[T_1, T_2]$, looking at the number of shadow robots in time. At any time, r_j and its shadow robots may be waiting at at most two vital points in total. Moreover, suppose that r_j embarks on a shadowless traversal of some arc (v_i, v_{i+1}) , leaving v_i at some moment of time t , arriving at v_{i+1} not earlier than at time $t + \text{dist}(v_i, v_{i+1})$. Then, throughout the time interval $[t, t + \text{dist}(v_i, v_{i+1})]$, robot r_j can have at most one shadow located at v_{i+1} . Suppose this shadow robot $r_j^{(i+1)*}$ exists. Then, the last traversal of r_j preceding time t must have been one of the arc (v_{i+1}, v_i) , and not shadowless. Tracing back in time the zig-zags of robot r_j between points v_i, v_{i+1} , during which it had shadows at both v_i and v_{i+1} , we come back to the earliest traversal of arc (v_i, v_{i+1}) (or possibly arc (v_{i+1}, v_i) , directly after the robot's arrival from v_{i-1} (respectively, from v_{i+2}). During this traversal, of duration at least $\text{dist}(v_i, v_{i+1})$, robot r_j had precisely one shadow located at v_i (respectively, at v_{i+1}). In summary, we have shown that during every shadowless traversal of arc (v_i, v_{i+1}) , robot r_j either has no shadow, or it has exactly 1 shadow and we can associate with this traversal another time period of length $\text{dist}(v_i, v_{i+1})$ during which it has exactly 1 shadow (with no overlap of time periods). The same argument applies for the counter-clockwise direction. Thus, we obtain a bound on the total waiting time of robot r_j and its shadows:

$$\sum_{i=0}^{2k} W_j(i) \leq 2(T_2 - T_1) - 2 \sum_{i=0}^{2k} (C_j(i, i+2) + C_j(i+1, i-1) - 2) \text{dist}(v_i, v_{i+1}),$$

where the constant is subtracted from C_j to account for boundary conditions around times T_1 and T_2 . Summing over all k robots we obtain:

$$\sum_{i=0}^{2k} W(i) \leq 2k(T_2 - T_1) - 2 \sum_{i=0}^{2k} (C(i, i+2) + C(i+1, i-1) - 2) \text{dist}(v_i, v_{i+1}).$$

Taking into account that the circumference of the cycle is 1 and that $W(0)$ is the minimum of all $W(i)$, we have:

$$W(0) \leq \frac{1}{2k+1} \sum_{i=0}^{2k} W(i) \leq \frac{2k(T_2 - T_1) - 2(C^{cc} + C^{cw}) + 4}{2k+1}.$$

and finally:

$$C^{cc} + C^{cw} < k(T_2 - T_1) - (k + \frac{1}{2})W(0) + 2. \quad (9)$$

(compare inequality (3)). Combining inequalities (7), (8), and (9), we obtain:

$$\begin{aligned} \tau &\geq \frac{(T_2 - T_1) - W(0)}{k(T_2 - T_1) - (k + \frac{1}{2})W(0) + 4k^2 + k + 3} \geq \frac{1}{k + \frac{1}{2} \frac{2(4k^2+k+3)-W(0)}{(T_2-T_1)-W(0)}} \geq \\ &\geq \frac{1}{k + \frac{4k^2+k+3}{(T_2-T_1)-2(4k^2+k+3)}} > \frac{1}{k} - \frac{4k^2+k+3}{(T_2 - T_1) - 2(4k^2 + k + 3)}. \end{aligned}$$

In the above, we assumed that $(T_2 - T_1) - W(0) > 0$, i.e., there cannot be a robot covering v_0 throughout the time interval $[T_1, T_2]$. This is true, since otherwise, taking into account that $W(i) \geq W(0)$ for all $1 \leq i \leq k$, all $2k+1$ points would have to be covered by a robot or its shadow robot throughout $[T_1, T_2]$, and there are at most $2k$ robots and shadow robots in total at any time, a contradiction.

Now, suppose that T_2 is chosen to be sufficiently large so that $\frac{4k^2+k+3}{(T_2-T_1)-2(4k^2+k+3)} < \frac{\epsilon}{2}$. We then have $\tau \geq \frac{1}{k} - \frac{\epsilon}{2}$, and so there exists a vital point on the cycle such that the time between some two successive visits of robots following \mathcal{A}' to this point is greater than $\frac{1}{k} - \frac{\epsilon}{2}$. This completes the proof of the lemma. \square

Proof of Lemma 4.1

Proof. We first show that our recursive procedure performs verification correctly. In the proof we use induction on k . More precisely, we assume inductively that for any $1 \leq l < k$ and $q \in C$ the call $TestLidSize(l, d, q)$ verifies whether one can cover all valid points in the interval $[q, e_n]$ using l lids of length d .

Consider the call $TestLidSize(k, d, p)$ in which the first lid is chosen to cover all vital points in $[p, p+d]$. A further recursive call $TestLidSize(k-1, d, p^*)$ verifies whether the remaining $k-1$ lids suffice to cover all valid points in $[p^*, e_n]$, where $p^* = \inf\{\text{vital } p' \in C = [p+d, e_n] : p' > p+d\}$. By the inductive assumption on k , we know that this call provides the correct answer. And if this answer is positive, i.e., the vital points in $[p^*, e_n]$ can be covered by $k-1$ lids we conclude that all vital points in $[p, e_n]$ (formed of vital points in $[p, p+d]$ and $[p^*, e_n]$) can be covered by k lids. Alternatively, if $k-1$ lids are insufficient to cover vital points $[p^*, e_n]$ the extra lid that covers vital points in $[p, p+d]$ is of no use for valid points in $[p^*, e_n]$ since the left endpoint of this lid must be aligned with p . Thus in this case, as expected, the answer computed by $TestLidSize(k, d, p)$ is also negative.

The time complexity $O(\min\{n, k \log n\})$ is dominated by computation of p^* at most k times, see line 3. If $p+d$ is vital and $(p+d) \neq e_j$, for any $1 \leq j < n$, p^* can be computed in constant time. Otherwise, we either use binary search on points b_1, \dots, b_n to find p^* imposing complexity $O(k \log n)$ or we search through this list of points in time $O(n)$. \square

Proof of Lemma 4.2

Proof. Consider any cover based on lids with the minimal size. In such a cover one can arrange the lids so that they touch but do not overlap with each other. If such an arrangement is not possible, one could decrease the length of the lids, contradicting the minimality of their length. Thus, we can assume that in the cover all the lids are partitioned into maximal sequences, such that in each sequence the lids are placed tightly one after another, but different sequences do not share their endpoints. Consider any such sequence based on m lids. The left endpoint of the leftmost lid in this sequence must coincide with some b_i . Otherwise, this would not be the leftmost lid in the sequence. If the right endpoint of the rightmost (m -th) lid in this sequence coincides with some e_j , the claim of the lemma follows. Assume, to the contrary, that this is not the case for any maximal sequence of lids. This means that the last lid in each maximal sequence overlaps with some neutral region, and consequently, that the length of the lids could be decreased. \square