Optimal Strategies for Patrolling Fences

Bernhard Haeupler

Carnegie Mellon University, Pittsburgh, PA, USA haeupler@cs.cmu.edu

Fabian Kuhn

University of Freiburg, Germany kuhn@cs.uni-freiburg.de

Anders Martinsson

ETH Zurich, Switzerland anders.martinsson@inf.ethz.ch

Kalina Petrova

ETH Zurich, Switzerland kpetrova@student.ethz.ch

Pascal Pfister

ETH Zurich, Switzerland ppfister@inf.ethz.ch

Abstract

A classical multi-agent fence patrolling problem asks: What is the maximum length L of a line fence that k agents with maximum speeds v_1, \ldots, v_k can patrol if each point on the line needs to be visited at least once every unit of time. It is easy to see that $L = \alpha \sum_{i=1}^{k} v_i$ for some efficiency $\alpha \in [\frac{1}{2}, 1)$. After a series of works [3, 8, 9, 10] giving better and better efficiencies, it was conjectured by Kawamura and Soejima [10] that the best possible efficiency approaches $\frac{2}{3}$. No upper bounds on the efficiency below 1 were known.

We prove the first such upper bounds and tightly bound the optimal efficiency in terms of the minimum speed ratio $s = \frac{v_{\text{max}}}{v_{\text{min}}}$ and the number of agents k. Our bounds of $\alpha \leq \frac{1}{1+\frac{1}{s}}$ and $\alpha \leq 1 - \frac{1}{\sqrt{k+1}}$ imply that in order to achieve efficiency $1 - \epsilon$, at least $k \geq \Omega(\epsilon^{-2})$ agents with a speed ratio of $s \geq \Omega(\epsilon^{-1})$ are necessary. Guided by our upper bounds, we construct a scheme whose efficiency approaches 1, disproving the conjecture stated above. Our scheme asymptotically matches our upper bounds in terms of the maximal speed difference and the number of agents used.

A variation of the fence patrolling problem considers a circular fence instead and asks for its circumference to be maximized. We consider the unidirectional case of this variation, where all agents are only allowed to move in one direction, say clockwise. At first, a strategy yielding $L = \max_{r \in [k]} r \cdot v_r$ where $v_1 \ge v_2 \ge \cdots \ge v_k$ was conjectured to be optimal by Czyzowicz et al. [3] This was proven not to be the case by giving constructions for only specific numbers of agents with marginal improvements of L. We give a general construction that yields $L = \frac{1}{33 \log_e \log_2(k)} \sum_{i=1}^k v_i$ for any set of agents, which in particular for the case $1, 1/2, \ldots, 1/k$ diverges as $k \to \infty$, thus resolving a conjecture by Kawamura and Soejima [10] affirmatively.

2012 ACM Subject Classification Networks \rightarrow Network algorithms; Theory of computation \rightarrow Distributed algorithms

Keywords and phrases multi-agent systems, patrolling algorithms

Digital Object Identifier 10.4230/LIPIcs.ICALP.2019.144

Category Track C: Foundations of Networks and Multi-Agent Systems: Models, Algorithms and Information Management

Related Version A full version of the paper is available at https://arxiv.org/abs/1809.06727.

Funding Bernhard Haeupler: Part of this work was done as a visiting professor at ETH Zürich supported by Mohsen Ghaffari. Supported also in part by NSF grants CCF-1527110, CCF-1618280, CCF-1814603, CCF-1910588, NSF CAREER award CCF-1750808 and a Sloan Research Fellowship. Kalina Petrova: Supported in part by Excellence Scholarship & Opportunity Programme.



 $\ensuremath{\mathbb O}$ Bernhard Haeupler, Fabian Kuhn, Anders Martinsson, Kalina Petrova, and Pascal Pfister;

licensed under Creative Commons License CC-BY 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019). Editors: Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi; Article No. 144; pp. 144:1–144:13



Leibniz International Proceedings in Informatics LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

144:2 Optimal Strategies for Patrolling Fences

1 Introduction

Patrolling is a fundamental task in robotics, multi-agent systems, and security settings. Given some environment of interest, and a collection of mobile agents, the aim is to coordinate the movements of the agents in order to, for example, guard an area from intrusion by an enemy, prevent accidents or failure of equipment, maintain up-to-date information of the environment, etc. For each of these tasks, ensuring that certain points in the environment get visited/monitored frequently is crucial. Performance of patrolling algorithms is consequently often measured in terms of *idleness* – roughly speaking, the time between two consecutive visits to a point in the environment.

Multi-agent patrolling has been extensively studied in the robotics literature since the early 2000s, e.g., see [1, 11] and the survey [12]. However, even for extremely clean and very simple models, determining optimal patrolling schemes poses many natural mathematical questions with interesting and surprisingly sophisticated answers [2, 3, 4, 5, 6, 7, 8, 9, 10].

1.1 Fence Patrolling

This paper studies a classical fence patrolling problem introduced by Czyzowicz et al. [3], which might be one of the cleanest and most natural patrolling problems: What is the maximum length L of a fence that k agents a_1, \ldots, a_k with maximum speeds v_1, \ldots, v_k can patrol if each point needs to be visited at least once every unit of time. Czyzowicz et al. introduce two variations of this question – the fence could be either an open curve, or a closed curve. For simplicity, we assume the open curve is a line segment and the closed curve is a circle.

For the line segment, it is easy to see that for any speeds the maximum length L satisfies $L = \alpha \sum_{i=1}^{k} v_i$ for some efficiency $\alpha \in [\frac{1}{2}, 1)$. In particular, in one unit of time an agent a_i can cover a length of at most v_i and all agents can cover at most a total length of $\sum_{i=1}^{k} v_i$. An efficiency of exactly $\alpha = 1$ is furthermore never possible because agents have to turn around eventually. On the other hand, an efficiency of $\alpha = \frac{1}{2}$ can easily be achieved by the following strategy:

PARTITION-BASED STRATEGY, \mathcal{A}_1 : For all $i \in [k]$, agent a_i patrols a subsegment of length $\frac{1}{2}v_i$ by going back and forth on this segment once every unit of time. This patrols a segment of length $L = \frac{1}{2} \sum_{i=1}^{k} v_i$ with idle time 1.

Considering patrol schedules on a circle, the picture is quite different than for a line segment. Again, the length L of any circle that can be patrolled by a set of agents is upper-bounded by the sum of the maximum speeds of the agents, since agent a_i cannot cover a length of more than v_i . Here, however, it is easy to find collections of agents and a corresponding patrol schedule that achieves this exactly – imagine k identical agents starting equidistantly along the circle and moving in unison in the same direction, say counter-clockwise.

1.2 Prior Work on Fence Patrolling

1.2.1 Prior Work on the Line Segment

Czyzowicz et al. [3] observed that the trivial scheme \mathcal{A}_1 with efficiency $\frac{1}{2}$ is optimal if the paths of the agents never cross. To see this, note that the leftmost agent a_i cannot walk away further than $\frac{1}{2}v_i$ from the leftmost point of the fence as it would take more than one unit of time between two visits of this point. By the same argument the agent a_i to the

right of agent a_i cannot ever be further away than $\frac{1}{2}(v_i + v_j)$ from the leftmost point of the fence and induction shows that a total fence length of $\frac{1}{2}\sum_{i=1}^{k}v_i$ is best possible. For the special case of all agents having the same speed the assumption that the paths of the agents never cross is furthermore without loss of generality as one can equally well switch identities of agents at a crossing, making the agents bounce off each other instead of crossing. In the worst case an efficiency of $\frac{1}{2}$ is thus optimal and Czyzowicz et al. posited [3] that indeed no better efficiency can be achieved for any speeds.

Surprisingly, Kawamura and Kobayashi [9] disproved this by providing an explicit fence patrolling schedule for 6 agents with speeds 1, 1, 1, 1, $\frac{7}{3}$, and $\frac{1}{2}$ for a fence of length $\frac{7}{2}$, thus achieving an efficiency of $\frac{21}{41} > \frac{1}{2}$. This was improved by Dumitrescu, Ghosh and Tóth [8], who proposed a family of patrolling schedules with efficiency approaching $\frac{25}{48}$, and finally by Kawamura and Soejima [10] who achieved an efficiency approaching $\frac{2}{3}$. Kawamura and Soejima furthermore explicitly conjectured that no efficiency better than $\frac{2}{3}$ is possible for any set of speeds [10, Conjecture 6, page 9].

On the other hand, except for the setting of equal speeds discussed above, no upper bounds on the efficiency below 1 have been provided in the literature [3, 9, 8, 10].

1.2.2 **Prior work on the Circle**

For a general set of agents, Czyzowicz et al. [3] proposed the following universal scheme that generalizes the schedule above for equal speeds:

RUNNERS STRATEGY, \mathcal{A}_2 : Assume $v_1 \geq v_2 \geq \cdots \geq v_k$. Find the $r \in [k]$ that maximizes $r \cdot v_r$, and let the *r* fastest agents move equidistantly along the circle at speed v_r . This patrols a circle of length $L = \max_{r \in [k]} r \cdot v_r$ with idle time 1.

Suppose for a collection of agents with maximum speeds $v_1 \ge v_2 \ge \cdots \ge v_k$, \mathcal{A}_2 produces a schedule on a circle with length L. Without loss of generality, we can assume L = 1. Then $\max_{r \in [k]} r \cdot v_r = 1$, and by possibly increasing the maximum speed of some agents we may assume $v_i = 1/i$ for each $i \in [k]$. Note that increasing speeds in this way can only increase the maximum circumference that can be patrolled with idle time 1 using these agents, but will not increase the length produced by \mathcal{A}_2 . Thus, if there is any collection of agents where there is a patrol schedule that performs better than \mathcal{A}_2 , there must be such a schedule in the case of *harmonic* maximum speeds $1, 1/2, \ldots, 1/k$.

To analyse the performance of patrol schemes on the circle, Czyzowicz et al. considered two different cases: *unidirectional* patrol schedules, where agents are only allowed to move in one direction, and general (or *bidirectional*) patrol schedules, where agents are allowed to go in both directions. Clearly, any patrol schedule obtained through \mathcal{A}_2 is unidirectional.

In the bidirectional case, it is not too hard to see that there are situations where \mathcal{A}_2 is not optimal. Indeed, in the case of harmonic maximum speeds, the partition-based strategy \mathcal{A}_1 , which works in the same way for a circular fence as for a line segment fence, would give $L = (1 + 1/2 + \cdots + 1/k)/2$, which is bigger than 1 as given by \mathcal{A}_2 for any $k \ge 4$. In fact, an example with three agents was given in [3] where neither \mathcal{A}_1 nor \mathcal{A}_2 are optimal. This was strengthened further by Dumitrescu et al. [8], who showed for any $k \ge 4$ there exists a collection of k agents where what they call the *train strategy* \mathcal{A}_3 performs strictly better than both \mathcal{A}_1 and \mathcal{A}_2 . To the authors' knowledge, no universal scheme has been proposed to always produce an optimal patrol schedule in this setting.

For the unidirectional case, it was initially conjectured by Czyzowicz et al. that \mathcal{A}_2 is optimal for any set of agents. This was proved to be true for up to four agents. However, it was shown incorrect by parallel results by Dumitrescu et al. [8] and Kawamura and Soejima [10], who gave explicit examples of patrol schemes for 32 and 122 agents (with harmonic

speeds) with $L = 1 + \epsilon$ (for a small unspecified $\epsilon > 0$) and L = 1.05 respectively. Kawamura and Soejima further conjectured that the maximum length of a unidirectional circle that can be patrolled by agents with speeds $1, 1/2, \ldots 1/k$ diverges as $k \to \infty$.

2 Our Results

This paper advances the understanding of the fence patrolling problem by giving tight upper and lower bounds on the optimal efficiency for the line segment, and a construction for the circle with efficiency of $\Theta(\frac{1}{\log_e \log_2 k})$ for any set of k agents. To a large extent it concludes the main line of inquiry put forward in the works discussed above [3, 9, 8, 10].

2.1 Results for the Line Segment

We provide the first technique to prove general impossibility results for the fence patrolling problem. We explain our ideas in more detail in Section 4 and merely state our main upper bound here:

▶ **Theorem 2.1.** Any fence patrol schedule with k agents with maximum speeds v_1, \ldots, v_k patrols a fence of length at most

$$L \le \sum_{i=1}^{k} \frac{v_i}{1 + \frac{v_i}{\max_j v_j}}.$$

One way to interpret Theorem 2.1 is that the contribution of an agent a_i depends not only on his/her own speed v_i but also on how much slower he/she is than the fastest agent. In particular, instead of always contributing v_i , as in the trivial upper bound, an agent contributes at most $\frac{1}{1+\frac{1}{s_i}} \cdot v_i$ given that the fastest agent patrolling is a factor of s_i faster than a_i . That is, the "relative efficiency" of an agent a_i ranges anywhere between 1/2and 1 depending on s_i , which always constitutes an improvement over the trivial upper bound of $\sum_i v_i$.

We also show that Theorem 2.1 can be used to prove an upper bound on the efficiency of a schedule solely in terms of the number of agents:

▶ Lemma 2.2. Any fence patrolling schedule with k agents has an efficiency of at most $1 - \frac{1}{\sqrt{k+1}}$.

We note that our upper bounds are tight in several interesting special cases. Specifically, for the case of agents having identical speeds, Theorem 2.1 shows that the efficiency of the schedule (and indeed each agent) is at most $\frac{1}{2}$, reproving the result of [3]. In contrast to the symmetry argument about non-crossing agents explained above, our arguments and upper bounds easily extend to near-identical speeds as well. Lastly, it is easy to check that Theorem 2.1 is tight when applied to the configuration of agents used by Dumitrescu et al. [8] and Kawamura and Soejima [10] for their construction to obtain efficiency ratios of 25/48 - o(1) and 2/3 - o(1), respectively.

Our upper bounds do not exclude schedules with efficiency close to 1. They do however give important restrictions and clues about what an extremely efficient schedule, if it exists, has to look like. In particular, Lemma 2.2 implies that any schedule with efficiency $1 - \epsilon$ has to have at least $\left(\frac{1}{2\epsilon}\right)^2$, i.e., quadratically in $\frac{1}{\epsilon}$ many agents. In the same manner, Theorem 2.1 implies that, with $\epsilon \to 0$, the ratio between the fastest and slowest agent has to be at least $\Omega(\frac{1}{\epsilon})$, i.e. grow unboundedly. Even more interestingly, the way the upper bound in Theorem 2.1 depends on $\max_i v_i$ seems to indicate that even just a single very fast agent can raise the "relative efficiency" of slower agents from 1/2 to almost 1.

Equipped with this better understanding and guidance from our impossibility results we were, to our surprise, able to design schedules which achieve an efficiency arbitrarily close to 1, thus disproving the conjecture of [10]:

▶ **Theorem 2.3.** For any sufficiently large k, there exists a fence patrolling schedule with efficiency $1 - \frac{3.5}{\sqrt{k}}$. Such a schedule uses k - 1 agents of speed one and one agent with maximum speed $\Theta(\sqrt{k})$.

Note that this theorem implies that for any $\epsilon > 0$ there exists a fence patrolling schedule with efficiency $1 - \epsilon$ using $O(\frac{1}{\epsilon^2})$ agents – one with speed $\Theta(\frac{1}{\epsilon})$ and all others with speed 1. In other words, the efficiency can be made arbitrarily close to 1 by choosing the appropriate number and maximum speeds of agents.

We remark that Theorem 2.3 also shows that both our upper bounds are asymptotically tight. In particular, the optimal efficiency for any schedule with k agents is indeed $1 - \Theta(\frac{1}{\sqrt{k}})$. Furthermore, for any $s \ge 1$, there is a configuration (with $k = \Theta(s^2)$ agents), where the maximum speeds of the agents differ by a factor s and for which the optimal efficiency is $\frac{1}{1+\frac{1}{\Theta(s)}} = 1 - \Theta(\frac{1}{s})$.

2.2 Results for the Circle

We resolve the conjecture by Kawamura and Soejima affirmatively. Namely, for any large enough k, we can construct a patrol schedule with idle time 1 using agents with maximum speeds $1, 1/2, \ldots 1/k$ that patrols a unidirectional circle of length $L = \Theta\left(\frac{\log_2 k}{\log_e \log_2 k}\right)$. In fact, our construction extends to a new universal scheme for the unidirectional circle. This is captured in the following theorem.

▶ **Theorem 2.4.** For k sufficiently large and for any k agents with maximum speeds $v_1, \ldots v_k$ there exists a patrol scheme with idle time 1 that patrols a unidirectional circle of length

$$L = \frac{1}{33 \log_e \log_2 k} \sum_{i=1}^k v_i.$$

The construction of our schedule has two steps: we first divide the agents into $\Theta(\log_2 k)$ groups, reducing the speed of some and discarding others so that each group consists of a power of 2 number of agents that move with the same speed, which is also a power of 2 times the sum of speeds. This allows us to use a randomized construction, in which the agents from each group are placed equidistantly around the circle with a random offset from some fixed "beginning" of the circle, and move around it with the same speed. We show that with this patrol schedule, most points are visited as frequently as required by our theorem. Then as a second phase, we cut out the bad points – that is, the ones that are not visited as frequently as necessary. We move the patrol schedule to a smaller circle, intuitively only consisting of the good bits. Agents move as if they were on the larger circle, but whenever moving though a cut-out segment, they just stand still instead. The details of this scheme together with a proof sketch will be given in Section 6.

2.3 Organization

The rest of the paper is organized as follows: We first give a more formal model description of the fence patrolling problem as well as discuss some related models and works in Section 3. In Section 4 we explain and prove our upper bounds for the line segment. Section 5 explains

ICALP 2019

144:6 Optimal Strategies for Patrolling Fences

and gives a proof sketch of our optimal fence patrolling schedule for the line segment. Finally, Section 6 presents and and gives a proof sketch our schedule for a circle with length $\Theta(\frac{1}{\log_e \log_2 k} \sum_{i=1}^k v_i)$. Formal and complete proofs for the two schedules can be found at https://arxiv.org/abs/1809.06727.

3 Fence Patrolling and Related Models

In this section we give a more detailed formal definition for the fence patrolling model/problem and briefly discuss related models and results. The fence patrolling model as given by [3] is defined as follows:

- The environment \mathcal{E} to be partolled is 1-dimensional and consists of a line segment of length L or a circle of circumference L. This line segment or circle is also referred to as a fence.
- The fence patrolling problem consists of some finite number $k \in \mathbb{N}$ of mobile agents a_1, a_2, \ldots, a_k to patrol the fence, each having a possibly distinct positive maximum speed $v_1, v_2, \ldots, v_k \in \mathbb{R}_+$.
- A schedule for the fence patrolling problem consists of a k-tuple of functions $a_1, a_2, \ldots, a_k : [0, \infty) \to \mathcal{E}$ such that, for all $i \in [k], t \ge 0$ and $\epsilon > 0$,

 $\operatorname{dist}(a_i(t+\epsilon), a_i(t)) \le \epsilon \cdot v_i.$

That is, we assume patrolling starts at t = 0 and goes on indefinitely. Each agent follows a predetermined trajectory, in which he/she moves along \mathcal{E} with at most his/her maximum speed. In the case of a circular fence, the function $\operatorname{dist}(x, y)$ refers to the length of the shorter circle arc between $x, y \in \mathcal{E}$. In the case of the unidirectional circle, we have the additional requirement that $\forall i \in [k], t \geq 0$ and $0 < \epsilon < \frac{L}{2v_i}$, the shorter arc between $a_i(t)$ and $a_i(t + \epsilon)$ is the one that spans clockwise from $a_i(t)$.

- We say that a patrol schedule has *idle time* T for some fixed positive parameter T if for all $t \ge T$ and for all $x \in \mathcal{E}$, there is some agent that visits x during [t T, t]. Intuitively, this condition means that an intruder cannot remain undetected at a point for more than T time.
- Given a patrol schedule, we say that a point $(x,t) \in \mathcal{E} \times [T,\infty)$ is *T*-covered if some agent a_i visits the point $x \in \mathcal{E}$ on the fence during the time interval [t T, t]. Note that in this model an agent patrols/monitors a point $x \in \mathcal{E}$ by visiting it. On the one hand, this means the agents are limited to zero line of sight. On the other hand, no additional operation (e.g. stop and look around) is necessary to patrol a point.

One can see that a schedule has idle time T if and only if every point $(x,t) \in \mathcal{E} \times [T,\infty)$ is T-covered. It is easy to observe that any patrol schedule of a fence of length L with idle time T can be rescaled to a schedule of a fence of length $\alpha \cdot L$ with idle time $\frac{1}{\alpha} \cdot T$ for any $\alpha > 0$. Thus, to simplify terminology, we assume henceforth that T = 1 and we refer to 1-covered simply as covered.

Related models have been considered in the literature: where agents have positive line of sight [7], where agents have distinct walking and patrolling speeds [5], where some agents may be faulty [4], where only some regions of the environment need to be patrolled [2], or where the environment is a geometric tree [6]. However, all of these models feature identical agents and in particular do not allow for varying maximum speeds. Overall, the model given above is likely the cleanest and most natural model in which agents with different speeds can and have been studied. Despite the extreme simplicity of this model, this paper and prior

works on the fence patrolling problem [3, 9, 8, 10] show that very surprising and intricate phenomena occur when agents have different speeds and that these nontrivial consequences can be studied in the model defined above.

4 Impossibility Results for the Line Segment: Proof of Theorem 2.1 and Lemma 2.2

In this section, we prove two upper bounds on the length of a straight line fence (i.e. $\mathcal{E} = [0, L]$) patrolled by agents of maximum speeds v_1, \ldots, v_k .

Proof of Theorem 2.1. The main idea of the proof is to consider the two-dimensional spacetime continuum $S := [0, L] \times [0, \infty)$ and the trajectories of the agents along with the points they cover. Since, as noted in Section 3, a patrol schedule with agents a_1, \ldots, a_k has idle time 1 if and only if $\forall x \in [0, L], \forall t \geq 1, (x, t) \in S$ is covered by at least one agent a_j , the theorem can be equivalently stated as that, for any patrol schedule such that all points $(x, t) \in S$ with $t \geq 1$ are covered by some agent, we have

$$L \le \sum_{i=1}^{k} \frac{1}{\frac{1}{v_i} + \frac{1}{v_{\max}}}.$$
(4.1)

In fact, we will show (4.1) under the weaker assumption that only points $(x,t) \in S$ with $t \in [1, 2k]$ are covered by some agent.

Given a patrol schedule of [0, L] with agents a_1, \ldots, a_k and a non-empty subset $A \subseteq \{a_1, \ldots, a_k\}$, we define the *right border* (see Figure 2) of A as the function $B^A : [1, \infty) \to [0, L]$ given by $B^A(t) := \max\{x \in [0, L] : (x, t) \text{ is covered by some agent in } A\}$. We show (4.1) by considering consecutively the collections of agents A_1, \ldots, A_q , where $A_1 = \{a_{i_1}\}, \forall j \in [q] \setminus \{1\}, A_j = A_{j-1} \cup \{a_{i_j}\}, i_1, \ldots, i_q$ is a sequence of distinct integers in [k] to be specified later, and $1 \leq q \leq k$ as we might not need to consider all agents. The intuition behind this is that we are starting with an empty set and adding more agents in a specific order until some termination condition is met. It is clear that $\forall t \geq 0, \forall j \in [q-1], B^{A_j}(t) \leq B^{A_{j+1}}(t)$. The key idea of the proof is to consider what happens to the right border of A_j and A_{j+1} for some j is shown in Figure 1.

At this point we prove a claim which will be useful in specifying the sequence i_1, \ldots, i_q .

 \triangleright Claim 4.1. For any patrol schedule of [0, L] with set of agents $A = \{a_1, \ldots, a_k\}$ and idle time 1, for any subset A' of $\{a_1, \ldots, a_k\}$, and for any point $(p_x, p_t) \in S$ on the right border of A' (that is, such that $B^{A'}(p_t) = p_x$), there exists an $\varepsilon > 0$ such that there is at least one agent $a_i \in A \setminus A'$ that covers all points $(p_x + \nu, p_t)$ for $\nu \in [0, \varepsilon]$.

Proof. Note that there could not exist three points $(x_1, t), (x_2, t), (x_3, t)$ such that $0 \leq x_1 < x_2 < x_3 \leq L$ and some agent a_j covers (x_1, t) and (x_3, t) but not (x_2, t) . This is because the trajectory of every agent can be considered as a continuous function $f_{a_j} : [0, \infty) \to [0, L]$, so it cannot be that $\exists t_1, t_3 \in [t-1, t]$ such that $f(t_1) = x_1$ and $f(t_3) = x_3$ but $\exists t_2 \in [t-1, t]$ with $f(t_2) = x_2$. It follows that the set of points C_j on the segment between (p_x, p_t) and (L, p_t) covered by some agent a_j must be either the empty set or a segment. Now consider the set of agents A_p in $A \setminus A'$ that cover (p_x, p_t) and note that $\exists \varepsilon > 0$ such that $\exists a \in A_p$ that covers $(p_x + \varepsilon, p_t)$ (otherwise choose the non-empty C_j with $a_j \in A \setminus A'$ with the leftmost left end (p_l, p_t) that is strictly to the right of p_x and notice that all points $(p_{x'}, p_t)$ with $p_{x'} \in (p_x, p_l)$ are not covered by any agent in A, which is impossible as all points in S should be covered). But now $a \in A_p$ covers both (p_x, p_t) and $(p_x + \varepsilon, p_t)$, therefore it also covers anything in between since the points it covers between (p_x, p_t) and (L, p_t) must form a segment.



Figure 1 If we have added the blue agent and the red agent so far, then the right border is shown in yellow. If we now also add the green agent, the new right border is shown in black.

Now we can continue with the proof of our main upper bound by specifying i_1, \ldots, i_q . Consider a fixed patrol schedule of [0, L] with agents a_1, \ldots, a_k and idle time 1. To pick the sequence i_1, \ldots, i_q , consider the following procedure: initially, put $l_0 = (x_0, t_0) = (0, k)$ and pick $i_1 \in [k]$ such that agent a_{i_1} covers l_0 . For each consecutive $j = 1, 2, \ldots$, we let $l_j = (x_j, t_j)$ be such that

$$t_j = \arg\min_{t \in [k-j,k+j]} B^{A_j}(t)$$

and $x_j = B^{A_j}(t_j)$. Intuitively, l_j is the leftmost point of B^{A_j} between times k - j and k + j. Now if $x_j = L$, we stop adding agents and we set q := j. Note that if j = k, then $x_j = L$ as agents a_1, \ldots, a_k cover all of [1, 2k]. If $x_j < L$, pick as agent $a_{i_{j+1}}$ an agent that covers all the points with coordinates $(x_j + \nu, t_j)$ for $\nu \in [0, \varepsilon]$ for some small enough $\varepsilon > 0$. Such an agent should exist by Claim 4.1. To make sure l_j is always defined for any $j \in \{0, 1, \ldots, q\}$, if q = k, set $l_k := (L, k)$.

We note that $x_0, x_1, ..., x_q$ is non-decreasing, $x_0 = 0$ and $x_q = L$ since we either stopped adding agents when q < k because $x_q = L$ or we stopped when q = k, in which case all of $[0, L] \times [1, 2k]$ should be covered. Hence the theorem follows if we can show that $\forall j \in \{0, 1, ..., q - 1\}, x_{j+1} - x_j \leq \frac{1}{\frac{1}{v_{i_{j+1}}} + \frac{1}{v_{\max}}}$.

In order to bound this difference, we investigate how the right border moves when agent $a_{i_{j+1}}$ is added. Note that $\forall t \in [0, \infty)$,

$$B^{A_{j+1}}(t) = \max(B^{A_j}(t), B^{\{a_{j+1}\}}(t)).$$

For any time $t > t_j$, the rightmost point at time t that could be covered by any agent in A_j is $(x_j + (t - t_j)v_{\max}, t)$ since the speed of any agent is at most v_{\max} . Similarly, for any time $t < t_j$, the rightmost point at time t that could be covered is $(x_j + (t_j - t)v_{\max}, t)$. Thus, $\forall t \in [0, \infty)$,

$$B^{A_j}(t) \le x_j + |t - t_j| v_{\max}.$$

Denote by u the ray $(x_j + (t_j - t)v_{\max}, t)$ where $t \le t_j$, and by w the ray $(x_j + (t - t_j)v_{\max}, t)$ where $t \ge t_j$.

Next, consider $B^{\{a_{j+1}\}}(t)$. Since agent a_{j+1} covers l_j , this means that he/she visits x_j at some time between $t_j - 1$ and t_j , say at point (x_j, t_{visit}) . Under the restriction that the trajectory of agent $a_{i_{j+1}}$ should go through (x_j, t_{visit}) , it is clear that $B^{\{a_{i_j+1}\}}(t)$ is maximized if agent $a_{i_{j+1}}$ comes from the right at maximum speed, hits (x_j, t_{visit}) and then turns around and moves to the right at maximum speed, in which case equality is achieved in

$$\forall t \in [0, \infty), B^{\{a_{i_{j+1}}\}}(t) \le \min(x' + |t - t'| v_{i_{j+1}}, L),$$

where $(x',t') = (x_j + \frac{v_{i_{j+1}}}{2}, t_{visit} + \frac{1}{2})$. Denote by *h* the ray $(x' + (t'-t)v_{i_{j+1}}, t)$ where $t \le t'$ and by *g* the ray $(x' + (t-t')v_{i_{j+1}}, t)$ where $t \ge t'$.



Figure 2 The rays R_1 and R_2 give a bound to the right of the current right border and are given in red and green respectively, and in black and gray we have a bound to the right of the trajectory of the newly added agent a_i and its shadow, i.e. the points that are covered by it.

We thus get $\forall t \in [0, \infty)$,

$$B^{A_{j+1}}(t) = \max(B^{A_j}(t), B^{\{a_{i_{j+1}}\}}(t)) \le f(t),$$

where $f(t) = \max(x_j + |t - t_j| v_{\max}, x' + |t - t'| v_{i_{j+1}})$. Consider $t_{new} = \arg\min_{t \in [1,\infty)} f(t)$. Let $L_{new} = (x_{new} := f(t_{new}), t_{new})$ be the leftmost point on the aforementioned upper bound on $B^{A_{j+1}}(t)$. Notice that L_{new} is either the intersection of h and w, or the intersection of g and u, or the intersection of g and h. These three cases are illustrated in Figure 2. We have u in red and w in green. Consider the upper bound on $B^{\{a_{i_{j+1}}\}}(t)$ mentioned above. The trajectory of agent $a_{i_{j+1}}$ that would correspond to matching this upper bound is given in black and the points $a_{i_{j+1}}$ would cover if this was his/her trajectory are given in gray. We have that $L_{old} = (x_{old}, t_{old}) := l_j$. It can be seen by inspection of the three cases in Figure 2 that $|t_{new} - t_j| \le 1$. Then $t_{new} \in [k - (j + 1), k + (j + 1)]$, which makes L_{new} a candidate for l_{j+1} , therefore $l_{j+1} = (x_{j+1}, t_{j+1})$ will have $x_{j+1} \le x_{new}$. Thus it is enough to show that $x_{new} - x_j \le \frac{1}{v_{i_{j+1}}^{i_{j+1}+v_{max}}}$.

We need an upper bound on $d = x_{new} - x_{old}$. We consider the points $M = (x_M, t_M)$ and $N = (x_N, t_N)$ as illustrated in Figure 2, such that in all three cases $x_M = x_{new}$ and $|t_M - t_{new}| = 1$. In Case 1, we consider the segment MN of slope $\frac{1}{v_{i_{j+1}}}$ and $x_M - x_N = d$, and the segment $L_{old}L_{new}$ of w of slope $-\frac{1}{v_{max}}$ and $x_{new} - x_{old} = d$. This gives us

$$\frac{d}{v_{i_{j+1}}} + \frac{d}{v_{\max}} \le 1 \Rightarrow d \le \frac{1}{\frac{1}{v_{i_{j+1}}} + \frac{1}{v_{\max}}}.$$
(4.2)

In Case 2, we consider the segment $L_{new}L_{old}$ of u of slope $\frac{1}{v_{max}}$ and $x_{new} - x_{old} = d$, and the segment NM of slope $-\frac{1}{v_{i_{j+1}}}$ and $x_M - x_N = d$. This implies Equation (4.2) for Case 2

144:10 Optimal Strategies for Patrolling Fences

as well. In Case 3, we consider the segment MN of slope $\frac{1}{v_{i_{j+1}}}$ and $x_M - x_N = d$, and the segment NL_{new} of slope $-\frac{1}{v_{i_{j+1}}}$ and $x_{new} - x_N = d$. This means that

$$\frac{2d}{v_{i_{j+1}}} = 1 \Rightarrow d = \frac{v_{i_{j+1}}}{2} \le \frac{1}{\frac{1}{v_{i_{j+1}}} + \frac{1}{v}}.$$

Therefore, $x_{new} - x_{old} \le \frac{1}{\frac{1}{v_{i_{j+1}}} + \frac{1}{v_{\max}}}$ as desired.

Proof of Lemma 2.2. We show how $L \leq \left(1 - \frac{1}{\sqrt{k+1}}\right) \sum_{i=1}^{k} v_i$ follows from Theorem 2.1. First note that $L \leq \sum_{i=1}^{k} v_i - \frac{v_{\max}}{2}$ as each agent a_i contributes at most $v_i \cdot \frac{1}{1 + \frac{v_i}{v_{\max}}} \leq v_i$ while the agent with maximum speed contributes exactly $\frac{v_{\max}}{2}$. Therefore, if $v_{\max} \geq \frac{1}{\sqrt{k}} \sum_{i=1}^{k} v_i$ the desired upper bound for L follows immediately. It remains to deal with the case $v_{\max} < \frac{1}{\sqrt{k}} \sum_{i=1}^{k} v_i$. For this we first note that $x \cdot \frac{1}{1 + \frac{w_i}{v_{\max}}} = \frac{1}{\frac{1}{2} + \frac{1}{v_{\max}}}$ is a concave function in x for $0 \leq x \leq v_{\max}$, since the second derivative $-\frac{2}{(1 + \frac{w_i}{v_{\max}})^3 v_{\max}}$ is always negative. This allows us to apply Jensen's inequality and thus we have

$$L \le \sum_{i=1}^{k} \frac{1}{\frac{1}{v_i} + \frac{1}{v_{\max}}} \le k \frac{1}{\frac{1}{v_{avg}} + \frac{1}{v_{\max}}} = \frac{1}{1 + \frac{\sum_{i=1}^{k} v_i}{k \cdot v_{\max}}} \sum_{i=1}^{k} v_i \le \left(1 - \frac{1}{\sqrt{k} + 1}\right) \sum_{i=1}^{k} v_i,$$

where $v_{avg} = \frac{1}{k} \sum_{i=1}^{k} v_i$. This concludes the proof of Lemma 2.2.

5 A Schedule with Efficiency $1 - \epsilon$ for the Line Segment: Proof of Theorem 2.3

In this section, we give proof sketch that for any k agents, there exist speeds v_1, \ldots, v_k and a scheme for these agents to patrol a fence of length

$$L = \left(1 - \frac{3.5}{\sqrt{k}} + O(1/k)\right) \sum_{i=1}^{k} v_i.$$

This improves the result from [10] and therewith falsifies the corresponding conjecture stated in that paper.

Proof sketch of Theorem 2.3. Assume k is sufficiently large, and, for ease of notation, define n := k - 2. Let $L = n - 3/2\sqrt{n}$. We construct a schedule that patrols $\mathcal{E} = [0, L]$ with idle time 1, using n + 1 agents with maximum speed 1 and 1 agent with maximum speed $2\sqrt{n} - 1$. Thus we have a total speed of $V = \sum_{i=1}^{k} v_i = n + 2\sqrt{n}$. As the ratio between L and V approaches $1 - \frac{3.5}{\sqrt{k}} + O(1/k)$, Theorem 2.3 follows.

To simplify the presentation of the patrol schedule, we will allow the agents to occasionally "step out of the fence [0, L]", i.e. we allow an agent a_i to assume positions $a_i(t) < 0$ and $a_i(t) > L$ (to avoid this, we could also modify the schedule so that they stay at the respective end of the fence for a while). To keep the notation as clean as possible, we henceforth assume that n is a square number. Our schedule works as follows (see Figure 3 for a graphical representation):

SLOW AGENTS a_1, \ldots, a_n : For each $i \in \{0, \ldots, n\}$, agent a_i starts at time 0 at position $x = i - i/\sqrt{n}$ and moves $i/(2\sqrt{n})$ time units to the right. Then he or she repeats:

- move to the left for \sqrt{n} time units.
- move to the right for \sqrt{n} time units.

FAST AGENT a_{n+1} : The fast agent a_{n+1} starts at time 0 and repeats the following four steps: (1) Move from position 0 to position L + 1/2 with speed $2\sqrt{n} - 1$.

- (2) Move from position L + 1/2 to position -1/2 during the next $\sqrt{n}/2 + 1$ time units (e.g. with constant speed $(L+1)/(\sqrt{n}/2+1) = 2\sqrt{n} 7 + (16)/(\sqrt{n}+2)$).
- (3) Move from position -1/2 to position L with speed $2\sqrt{n} 1$.
- (4) Move from position L to position 0 in the next $\sqrt{n}/2$ time units (e.g. with constant speed $(L)/(\sqrt{n}/2) = 2\sqrt{n} 3$).

The idea behind our patrol schedule is to initially place the agents with maximum speed 1 equidistantly along the fence with gaps of length slightly smaller than 1, similar to the schedule for the fast agents in [10]. In contrast to their schedule, this is performed slightly out of phase between the agents. This will cover most of the points on the fence. The only problem appears whenever the agents turn around, as then the points right next to these turning points are not visited for more than 1 time unit, hence creating uncovered triangles in the "spacetime" diagram (white triangles in Figure 3). By timing the turning times of the agents appropriately, we ensure that these uncovered triangles are placed such that they can all be cleaned up by the last fast agent. Figure 3 gives a complete illustration of our schedule for 18 agents and intuitively shows that each point x on the line L is visited at least once every unit of time.

6 A schedule for the unidirectonal circle: Proof of Theorem 2.4

In this section, we will present a schedule with which a group of agents a_1, \ldots, a_k with maximum speeds v_1, \ldots, v_k can patrol a circle with circumference

$$\frac{1}{33\log_e \log_2(k)} \sum_{i=1}^k v_i,$$

followed by a short sketch of why the proposed schedule behaves as claimed. Our schedule will be divided in two parts. First, we will give a randomized construction for a strategy which allows the agents a_1, \ldots, a_k to patrol a circle with circumference 1 such that with probability 1 - o(1) "most" of the points are visited at least every $\Theta\left(\log \log(k)/(\sum_{i=1}^k v_i)\right)$ time units. Then, we will argue that one can remove the "bad" points and then "blow up" the circle to obtain the desired result.

Proof sketch of Theorem 2.4. We are given a group a_1, \ldots, a_k of agents with maximum speeds $v_1 \ldots, v_k$ and define $V := \sum_{i=1}^k v_i$. We propose the following schedule:

Circle Schedule:

- (1) Round speeds down to the next power of 2 and omit too slow agents: For all $i \in [k]$ let $j_i \in \mathbb{N}$ be the non-negative integer such that $V \cdot 2^{-j_i} \leq v_i < V \cdot 2^{-j_i+1}$. We define $v'_i := V \cdot 2^{-j_i}$ and $I := \{i \in [k] : v'_i \geq \frac{V}{4 \cdot k}\}$.
- (2) Group remaining agents according to their speed: For all $i \in \{0, ..., \lceil \log_2(k) \rceil + 2\}$ define $G_i := \{j \in I : v'_j = V \cdot 2^{-i}\}.$
- (3) Reduce number of agents in each group to a power of 2: For all $i \in \{0, \ldots, \lceil \log_2(k) \rceil + 2\}$ let $h_i \in \mathbb{N}$ be the positive integer such that $2^{h_i} \leq |G_i| < 2^{h_i+1}$ and let $G'_i \subseteq G_i$ be an arbitrary subset of G_i of size 2^{h_i} . We denote by $m'_i = |G'_i| \cdot v'_a = V \cdot 2^{h_i-i}$ the mass of the group G'_i , where $a \in G'_i$ (that is, each agent in G'_i has maximum speed v'_a after the rounding down in step 1).

144:12 Optimal Strategies for Patrolling Fences



Figure 3 The described schedule for n = 16. The dark grey area describes the points (x, t) which are covered by the slow agents a_0, \ldots, a_n while the light grey shaded area describes the points (x, t)which are covered by the fast agent in steps (1) and (3) of his protocol.

(4) Omit groups with too small mass: Let J := {i ∈ {0,..., ⌈log₂(k)⌉ + 2} : m'_i ≥ V/(16(⌈log₂(k)⌉+3))}.
(5) Patrol schedule: For each j ∈ J, pick an independent uniform random number r_j in the second schedule: For each j ∈ J, pick an independent uniform random number r_j in the second schedule: For each j ∈ J, pick an independent uniform random number r_j in the second schedule: For each j ∈ J, pick an independent uniform random number r_j in the second schedule: For each j ∈ J, pick an independent uniform random number r_j in the second schedule: For each j ∈ J, pick an independent uniform random number r_j in the second schedule: For each j ∈ J, pick an independent uniform random number r_j in the schedule: For each j ∈ J, pick an independent uniform random number r_j in the schedule: For each j ∈ J, pick an independent uniform random number r_j in the schedule: For each j ∈ J, pick an independent uniform random number r_j in the schedule: For each j ∈ J, pick an independent uniform random number r_j in the schedule: For each j ∈ J, pick an independent uniform random number r_j in the schedule: For each j ∈ J, pick an independent uniform random number r_j in the schedule: For each j ∈ J, pick an independent uniform random number r_j in the schedule: For each j ∈ J, pick an independent uniform random number r_j in the schedule: For each j ∈ J, pick an independent uniform random number r_j in the schedule: For each j ∈ J, pick an independent uniform random number r_j in the schedule: For each j ∈ J, pick an independent uniform random number r_j in the schedule: For each j ∈ J, pick an independent uniform random number r_j in the schedule: For each j ∈ J, pick an independent uniform random number r_j = for each j ∈ J, pick an independent uniform random number r_j = for each j ∈ J, pick an independent uniform random number r_j = for each j ∈ J, pick an independent uniform random number r_j = for each j ∈ J, pick an independent uniform random number r_j = for each interval $[0, \frac{1}{|G'_j|})$. At time 0, we place the $|G'_j|$ agents from the group G'_j at positions

$$r_j, r_j + \frac{1}{|G'_j|}, \dots, r_j + \frac{|G'_j| - 1}{|G'_j|},$$

i.e. we place all agents from the same group equidistantly from each other along the circle with a random offset from the origin. Then the agents G'_j walk along the circle with speed v'_i at all times.

Note that the two rounding steps and the two omitting steps each reduced the total "available" speed by a factor of at most 2, therefore we have that $\sum_{j \in J} m'_j \ge V/16$. Defining

$$T = \frac{1}{\min_{j \in J} m'_j} \le 16(\lceil \log_2(k) \rceil + 3)/V,$$

we observe that after T time units the distribution of agents along the circle repeats itself, i.e. if an agent with speed s is located at position x at time t, then another (or the same) agent of speed s is located at position x at time t + T. Cutting the time interval into small pieces of size $\tau \approx 16 \log_e \log_2(k)/V$, one can show that with probability 1 - o(1), almost all points get visited at least once in each of these small time intervals, and hence all these points get visited at least every 2τ time units. Next, we can "cut away" the few "bad" points for which this is not true and obtain a circle of circumference $1 - \epsilon$. Adjusting the above schedule so that agents "stand still" on the smaller circle whenever they cross a "bad" point on the original circle gives us a schedule for which each point (on the small circle) is visited at least every 2τ time units. Rescaling this patrol schedule then gives the scheme to patrol a circle with circumference $(1 - \epsilon) \frac{1}{2\tau} \ge \frac{1}{33 \log_{\tau} \log_2(k)} \sum_{i=1}^k v_i$ as desired.

— References

- Yann Chevaleyre. Theoretical analysis of the multi-agent patrolling problem. In Proceedings. IEEE/WIC/ACM International Conference on Intelligent Agent Technology, 2004. (IAT 2004)., pages 302–308, 2004. doi:10.1109/IAT.2004.1342959.
- 2 Andrew Collins, Jurek Czyzowicz, Leszek Gasieniec, Adrian Kosowski, Evangelos Kranakis, Danny Krizanc, Russell Martin, and Oscar Morales Ponce. Optimal patrolling of fragmented boundaries. In SPAA, pages 241–250. ACM, 2013.
- 3 Jurek Czyzowicz, Leszek Gąsieniec, Adrian Kosowski, and Evangelos Kranakis. Boundary patrolling by mobile agents with distinct maximal speeds. In *European Symposium on Algorithms*, pages 701–712. Springer, 2011.
- 4 Jurek Czyzowicz, Leszek Gasieniec, Adrian Kosowski, Evangelos Kranakis, Danny Krizanc, and Najmeh Taleb. When Patrolmen Become Corrupted: Monitoring a Graph Using Faulty Mobile Robots. *Algorithmica*, 79(3):925–940, 2017. doi:10.1007/s00453-016-0233-9.
- 5 Jurek Czyzowicz, Konstantinos Georgiou, Evangelos Kranakis, Fraser MacQuarrie, and Dominik Pajak. Distributed Patrolling with Two-Speed Robots (and an Application to Transportation). In Begoña Vitoriano and Greg H. Parlier, editors, *Operations Research and Enterprise Systems*, pages 71–95, Cham, 2017. Springer International Publishing.
- 6 Jurek Czyzowicz, Adrian Kosowski, Evangelos Kranakis, and Najmeh Taleb. Patrolling Trees with Mobile Robots. In Frédéric Cuppens, Lingyu Wang, Nora Cuppens-Boulahia, Nadia Tawbi, and Joaquin Garcia-Alfaro, editors, *Foundations and Practice of Security*, pages 331–344, Cham, 2017. Springer International Publishing.
- 7 Jurek Czyzowicz, Evangelos Kranakis, Dominik Pajak, and Najmeh Taleb. Patrolling by Robots Equipped with Visibility. In Magnús M. Halldórsson, editor, *Structural Information and Communication Complexity*, pages 224–234, Cham, 2014. Springer International Publishing.
- 8 Adrian Dumitrescu, Anirban Ghosh, and Csaba D Tóth. On Fence Patrolling by Mobile Agents. The Electronic Journal of Combinatorics, 21(3):P3–4, 2014.
- 9 Akitoshi Kawamura and Yusuke Kobayashi. Fence patrolling by mobile agents with distinct speeds. Distributed Computing, 28(2):147–154, 2015.
- 10 Akitoshi Kawamura and Makoto Soejima. Simple strategies versus optimal schedules in multi-agent patrolling. In International Conference on Algorithms and Complexity, pages 261–273. Springer, 2015.
- 11 Aydano Machado, Geber Ramalho, Jean-Daniel Zucker, and Alexis Drogoul. Multi-agent Patrolling: An Empirical Analysis of Alternative Architectures. In Jaime Simão Sichman, Françcois Bousquet, and Paul Davidsson, editors, *Multi-Agent-Based Simulation II*, pages 155–170, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.
- 12 David Portugal and Rui Rocha. A Survey on Multi-robot Patrolling Algorithms. In Luis M. Camarinha-Matos, editor, *Technological Innovation for Sustainability*, pages 139–146, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

ICALP 2019