

When Patrolmen Become Corrupted: Monitoring a Graph using Faulty Mobile Robots

Jurek Czyzowicz¹, Leszek Gasieniec², Adrian Kosowski³,
Evangelos Kranakis⁴, Danny Krizanc⁵, and Najmeh Taleb⁴

¹ Dépt. d'informatique, Univ. du Québec en Outaouais, Gatineau, Québec, Canada.

² Dep. of Computer Science, University of Liverpool, Liverpool, UK.

³ LIAFA, Inria and Université Paris Diderot, Paris, France.

⁴ School of Computer Science, Carleton University, Ottawa, Canada.

⁵ Dep. of Math. and CS, Wesleyan University, Middletown CT, USA.

Abstract. A team of k mobile robots is deployed on a weighted graph whose edge weights represent distances. The robots perpetually move along the domain, represented by all points belonging to the graph edges, not exceeding their maximal speed. The robots need to patrol the graph by regularly visiting all points of the domain. In this paper, we consider a team of robots (patrolmen), at most f of which may be unreliable, i.e. they fail to comply with their patrolling duties.

What algorithm should be followed so as to minimize the maximum time between successive visits of every edge point by a reliable patrolmen? The corresponding measure of efficiency of patrolling called *idleness* has been widely accepted in the robotics literature. We extend it to the case of untrusted patrolmen; we denote by $\mathfrak{I}_k^f(G)$ the maximum time that a point of the domain may remain unvisited by reliable patrolmen. The objective is to find patrolling strategies minimizing $\mathfrak{I}_k^f(G)$.

We investigate this problem for various classes of graphs. We design optimal algorithms for line segments, which turn out to be surprisingly different from strategies for related patrolling problems proposed in the literature. We then use these results to study the case of general graphs. For Eulerian graphs G , we give an optimal patrolling strategy with idleness $\mathfrak{I}_k^f(G) = (f + 1)|E|/k$, where $|E|$ is the sum of the lengths of the edges of G . Further, we show the hardness of the problem of computing the idle time for three robots, at most one of which is faulty, by reduction from 3-edge-coloring of cubic graphs — a known NP-hard problem. A byproduct of our proof is the investigation of classes of graphs minimizing idle time (with respect to the total length of edges); an example of such a class is known in the literature under the name of Kotzig graphs.

Key Words and Phrases: Fault Tolerant, Idleness, Kotzig Graphs, Patrolling.

1 Introduction

Patrolling occurs in many activities of everyday life whenever it is required to monitor a specific region, for example, the perimeter of a piece of land or a building, so as to investigate a feature of interest for purposes of surveillance. Typically, in such a setting patrolmen are assigned to monitor specified regions by moving perpetually at regular intervals through areas assigned to them.

In this paper, we are interested in patrolling when some of the patrolmen may be unreliable (faulty) in that they fail to report their monitoring activities. More specifically, we model and study the following problem: We are given a team of robot patrolmen and a domain to be monitored. Assume that some of the patrolmen may be unreliable. We want to design a strategy constructing perpetual patrolmen trajectories, so that, independently of which subset of them (of a given size) will turn out to be faulty, no point of the environment will ever be left unvisited by some reliable robot longer than the allowed *idle time*.

Preliminaries and Notation. We are given a connected topological graph $G = (V, E)$ with V being its set of vertices and E its set of edges. In the sequel we define several useful concepts.

The Jordan arc representing each edge $e \in E$ of the graph $G = (V, E)$ is modeled as a smooth continuous and rectifiable curve of arbitrary positive length represented by its edge weight $w(e)$. We may suppose that the graph is embedded in 3D space, with no edge crossings. By $|E|$ we denote the sum of the lengths of the edges of G .

At any time a robot may occupy any point belonging to edge e (so the sum of its distances from both endpoints of e sums up to $w(e)$). We denote by \mathcal{D}_G the domain (the union of edges) along which the robots walk. We assume a continuous traversal model, whereby the movement of the i -th robot within \mathcal{D}_G follows a continuous function of time $\pi_i : [0, \infty) \rightarrow \mathcal{D}_G$, for each $i = 1, 2, \dots, k$. Hence, $\pi_i(t)$ denotes the position in \mathcal{D}_G of the i -th robot at time t . Each robot may move in any direction along an edge not exceeding the maximum (unit) speed so within time interval $[t_1, t_2]$ each robot may travel a distance of at most $t_2 - t_1$. We also suppose that when walking at maximum speed, a robot travels the unit distance in unit time, so that time and distance travelled are commensurable. By *patrolling strategy* we understand the set $\mathcal{P} = \{\pi_1, \pi_2, \dots, \pi_k\}$ of infinite trajectories of k robots in \mathcal{D}_G , where $\pi_i(t)$ is the point of \mathcal{D}_G occupied by the i -th robot at time t .

The performance of the patrolling strategy is evaluated by using a measure of *idleness*, widely used in robotics literature. Suppose that we design the patrolling strategy $\mathcal{P} = \{\pi_1, \pi_2, \dots, \pi_k\}$ for k robots moving in the domain of a geometric graph G when each robot is reliable. Then the idleness of strategy \mathcal{P} for graph G (or its *idle time*), denoted by $\mathfrak{S}_k^f(G, \mathcal{P})$ is the supremum of the lengths of time intervals between two consecutive visits to the same point of \mathcal{D}_G (supremum taken over time and all points of \mathcal{D}_G). When up to f robots may be faulty, we assume that the adversary, knowing our strategy, may choose a set F of f faulty robots, a point p of the domain and a time moment $t \geq 0$. The idleness of the strategy is the supremum (taken over all such adversarial choices) of time intervals T such that point p is not visited during the time interval $[t, t + T]$ by any reliable robot. Finally, the idleness of a graph G for k robots, at most f of which may be faulty, is denoted by $\mathfrak{S}_k^f(G) := \inf_{\mathcal{P}} \mathfrak{S}_k^f(G, \mathcal{P})$. Hence $\mathfrak{S}_k^f(G)$ is the lower bound of idleness over all possible patrolling strategies. When there are no faulty robots (i.e., $f = 0$) we use the notation $\mathfrak{S}_k(G, \mathcal{P}) := \mathfrak{S}_k^0(G, \mathcal{P})$.

Consider a walk of a robot within the segment, which starts at one of its endpoints, walks to the other endpoint and returns to the initial one. Such a cyclic path around the segment has length equal to twice its size. By an Eulerian tour of the segment by r robots we mean a perpetual movement of these robots, which are equally spaced around such a cyclic path, and walking in the same cyclic direction with the same speed.

The Chinese Postman Problem consists in finding the shortest closed tour of a graph G which visits each edge at least once. A tour of shortest length is also called a Chinese Postman Tour. By $CPT(G)$ we denote the length of a Chinese Postman Tour on the graph G . Recall that for an Eulerian graph, an Eulerian cycle is the optimal solution to the Chinese Postman Tour, while in a tree, such a tour crosses each edge twice (see [6]).

Related work. Patrolling has been defined as the act of surveillance consisting in walking perpetually around an area in order to protect or supervise it. It is useful in monitoring and locating objects or humans that need to be rescued from a disaster, in ecological monitoring or detecting intrusion. Network administrators may use mobile agent patrols to detect network failures or to discover web pages which need to be indexed by search engines, cf. [23]. Patrolling has been re-

cently intensively studied in robotics (cf. [7, 14, 15, 18, 23, 31]) where it is often viewed as a version of terrain *coverage*, a central task in robotics.

Boundary and area patrolling have been studied in [1, 14, 15, 27] with approaches placing more emphasis on experimental results. The accepted measure of the algorithmic efficiency of patrolling is called *idleness* and it is related to the frequency with which the points of the environment are visited (cf. [7, 14, 15, 23]); this criterion was first introduced in [23]. Depending on the requirements, idleness may sometimes be viewed as the average [14], worst-case [31], probabilistic [1] or experimentally verified [23] time elapsed since the last visit of a node (cf. also [7]). In some papers the terms of *blanket time* [31] or *refresh time* [27] have been used instead.

A survey of diverse approaches to patrolling based on the idleness criteria can be found in [28]. In [3–5] patrolling is studied as a game between patrollers and the intruder. Some papers consider the patrolling problem based on swarm or ant-based algorithms [16, 25, 31]. In these approaches robots are memoryless (or having small memory), decentralized [25] with no explicit communication permitted either with other robots or the central station, with local sensing capabilities (e.g., [16]). Ant-like algorithms usually mark the visited nodes of the graph. [31] presents an evolutionary process and shows that a team of memoryless robots, by leaving marks at the nodes while walking through them, after relatively short time stabilizes to the patrolling scheme in which the frequency of the traversed edges is uniform to a factor of two (i.e., the number of traversals of the most often visited edge is at most twice the number of traversals of the least visited one).

A theoretical analysis of approaches to patrolling in graph-based models can be found in [7]. The two basic methods are referred to as *cyclic strategies*, where a single cycle spanning the entire graph is constructed with the robots assigned to consecutively traverse this cycle in the same direction, and as *partition-based strategies*, where the region is split into a number of either disjoint or overlapping portions to be patrolled by subsets of robots assigned to these regions. The environment and the time considered in the models studied are usually discrete in an underlying graph environment. In [27], polynomial-time patrolling solutions for lines and trees are proposed. For the case of cyclic graphs, [27] proves the NP-hardness of the problem and a constant-factor approximation is proposed.

Patrolling with robots that do not necessarily have identical speeds has been initiated in [10]. As shown in [13, 21] it offers several surprises both in terms of the difficulty of the problem as well as in terms of the algorithmic results obtained. In particular, no optimal patrolling strategy involving more than three robots has yet been proposed.

Fault tolerance related to mobile robots has been considered for several problems in distributed computing with failures occurring either to the environment (nodes or links) or to the robots themselves. The cases of faulty robots were often studied for the robot gathering problem under various assumptions of faults (crash and Byzantine), e.g., [2, 11, 12]. Other studies concerned the problem of convergence, e.g. [8], flocking, e.g., [30] and many other ones. Several papers, e.g., [9, 20, 29] concerned unreliable or inaccurate robot sensing devices, rather than the robots themselves. Experimental papers related to unreliable robots performing patrolling were considered in the robotics literature [14, 15, 18, 24]. To the best of our knowledge, the theoretical study, considered in our paper, concerning optimally patrolling a connected graph in the presence of faulty robots has not been investigated in the past.

Outline and results of the paper. In Section 2, we provide optimal patrolling strategies for line segments. These non-intuitive strategies rely on a decomposition of the set of robots into three groups with different patrolling tours, in a way dependent on k and f . Next we employ these

results in Section 3 as building blocks to provide strategies for general graphs. In particular, for any Eulerian graph G we show that the idleness satisfies $\mathfrak{S}_k^f(G) = (f + 1)|E|/k$. In Section 3.2, we analyze the hardness of the problem of computing the idle time on a specific class of graphs (derived from the class of Kotzig graphs) by showing that if the idle time could be computed optimally then we could solve 3-edge-coloring of cubic graphs, a well known NP-hard problem (see [17]). Finally, in Section 4 we conclude with a summary of our results and mention additional work and various related open problems. All missing proofs can be found in the Appendix.

2 Idleness of Line Segments

In this section we study exclusively the idleness of the line segment and provide upper and lower bounds for idleness. Without loss of generality we assume that we have to patrol the unit-length segment, represented by the interval $I = [0, 1]$. However, the results can be easily reformulated for segments of any given length. Throughout the main part of this section we assume that most of the robots are reliable, more precisely that $f < \frac{k-2}{2}$. We first give a patrolling strategy and analyze its performance. Then we analyze the lower bound for segment idleness showing that our strategy is optimal for odd f and almost optimal for even f .

2.1 The upper bound

The idea of the strategy is the following. We partition the segment I into three subsegments I_L (left), I_R (right), and I_M (middle), where I_M does not contain any endpoint of I . Two subsets of robots will follow Eulerian tours of I_L and I_R and the remaining robots are assigned to do the Eulerian tour of the entire I . We show that by choosing sizes of the segments of the partition as well as the number of robots assigned to each Eulerian tour we obtain an efficient strategy. We have the following theorem:

Theorem 1. *Consider k robots patrolling segment $I = [0, 1]$, with at most f of them faulty where $k > 2$ and $f < \frac{k}{2} - 1$. There exists a patrolling strategy \mathcal{P} of I whose idleness satisfies*

$$\mathfrak{S}_k^f(I, \mathcal{P}) \leq \frac{2 \lfloor f/2 \rfloor + 2}{k - 2 \lceil f/2 \rceil}$$

Proof. First we give explicitly the patrolling strategy.

1. Decompose the unit interval I into three segments I_L, I_M and I_R with pairwise disjoint interiors:

$$I_L := \left[0, \frac{\lfloor f/2 \rfloor}{k - 2 \lceil f/2 \rceil}\right], \quad I_M := \left[\frac{\lfloor f/2 \rfloor}{k - 2 \lceil f/2 \rceil}, 1 - \frac{\lfloor f/2 \rfloor}{k - 2 \lceil f/2 \rceil}\right], \quad I_R := \left[1 - \frac{\lfloor f/2 \rfloor}{k - 2 \lceil f/2 \rceil}, 1\right]$$

2. For each of the segments I_L, I_R assign $\lceil f/2 \rceil$ equally spaced robots to perform an Eulerian tour of this segment.
3. The remaining $k - 2 \lceil f/2 \rceil$ robots perform an Eulerian tour of the entire segment I . These robots are also equally spaced around I .

Observe first that the subsegments I_L, I_R, I_M are well defined. Indeed, as f is an integer $f < \frac{k}{2} - 1$ implies $f \leq \frac{k-3}{2}$. Hence $2 \lceil f/2 \rceil \leq f + 1 \leq \frac{k-3}{2} + 1 = \frac{k-1}{2}$. However, for any integer $k > 2$ we have $\frac{k-1}{2} \leq k - 2$ which implies $2 \lceil f/2 \rceil \leq k - 2$, so the denominator of the fractions in the

definitions of the segments I_L, I_M, I_R is not zero. Moreover, the point $1/2$ belongs to I_M , hence all these segments are well defined.

Denote by S_L (respectively S_R, S_I) the set of robots executing an Eulerian tour of I_L (respectively I_R, I). Observe that the distance d between two consecutive robots of S_I , computed around this Eulerian tour, equals $d = \frac{2}{k-2\lceil f/2 \rceil}$. We now prove the correctness of the upper bound on the idleness of any point $p \in I$. We consider two cases: when p belongs to an extremal subsegment I_L or I_R and when p is in the middle segment I_M .

Case 1: Point p is in extremal subsegment (by symmetry we may assume without loss of generality that $p \in I_L$). Suppose first that at least one robot $r_i \in S_L$ is not faulty. Then r_i revisits every point of I_L at time intervals of at most $2|I_L|$. Hence the idleness of $p \in I_L$ (maximized at endpoints of I_L) is bound by

$$\mathfrak{S}_k^f(I, \mathcal{P}) \leq 2|I_L| = \frac{2 \lceil f/2 \rceil}{k - 2 \lceil f/2 \rceil} < \frac{2 \lfloor f/2 \rfloor + 2}{k - 2 \lceil f/2 \rceil}$$

When all $\lceil f/2 \rceil$ robots of S_L are faulty, the idle time is maximized for $p = 0$, while the adversary chooses the remaining $\lfloor f/2 \rfloor$ faulty robots to form a stream of consecutive robots of S_M . Then the time between visits of point $p = 0$ by two reliable robots (i.e. one preceding and one following such a stream) equals $d(\lfloor f/2 \rfloor + 1)$ and we have

$$\mathfrak{S}_k^f(I, \mathcal{P}) \leq d(\lfloor f/2 \rfloor + 1) = \frac{2(\lfloor f/2 \rfloor + 1)}{k - 2 \lceil f/2 \rceil} \quad (1)$$

again verifying the claim of the theorem. The argument is entirely symmetric when $p \in I_R$ and is therefore omitted.

Case 2: $p \in I_M$. The visits to this point are made exclusively by the robots from S_I . Point p is being visited by two streams of robots executing the Eulerian tour of I , one walking over p from left to right and the other one from right to left (clearly, in the Eulerian cycle the robots are moving in the same direction, but from the “point of view of the point p ” the traversal is in opposite directions). Each of these streams may have several faulty robots, and the idle time at p depends on the distance between the two reliable robots starting and ending such faulty streams.

Consider first the case when the two faulty-robots streams, visiting p at the same time, are disjoint, i.e. separated by at least one reliable robot (cf. Fig. 1 (a), where \circ denotes a reliable robot and \bullet a faulty robot).




Fig. 1.

To maximize the time while point p remains unvisited by reliable robots, the adversary has to make faulty two sequences of consecutive robots (i.e. those belonging to both left-to-right and right-to-left streams) arriving at p at the same time. The idle time is then determined by the length of the

shorter of the two sequences of consecutive faulty robots, which in the worst-case contains $\lfloor f/2 \rfloor$ robots. Then the claim of the theorem is again satisfied by Equation (1).

Consider now the case when there is a single faulty-robot stream visiting p in both directions (see Fig. 1 (b) which depicts a time moment t when this happens). In the worst case this stream may contain f robots. Let t_1 be the time since the last visit of p by a reliable robot y and t_2 - the time when the next reliable robot z visits p . As $x > |I_L| = \frac{\lfloor f/2 \rfloor}{k-2\lfloor f/2 \rfloor}$ and all robots present within I_L at time t are faulty, as well as the distance between y and z around the Eulerian cycle is at most $d(f+1)$, we have

$$\mathfrak{S}_k^f(I, \mathcal{P}) = t_1 + t_2 \leq d(f+1) - 2|I_L| = \frac{2(f+1)}{k-2\lfloor f/2 \rfloor} - \frac{2\lfloor f/2 \rfloor}{k-2\lfloor f/2 \rfloor} = \frac{2\lfloor f/2 \rfloor + 2}{k-2\lfloor f/2 \rfloor}.$$

This completes the proof of Theorem 1. ■

2.2 The lower bound

We first show the following lemma, which applies for general graphs.

Lemma 1. *Consider a patrolling strategy \mathcal{P} of graph G . Let E' be a subset of segments of edges of G , such that starting from some time moment of the strategy, in the union of the interiors of all elements of E' there are always at most r robots. Then $\mathfrak{S}_k^f(\mathcal{I}, \mathcal{P}) \geq \frac{(f+1)|E'|}{r}$, where $|E'|$ denotes the sum of lengths of segments of E' .*

The next theorem proves that the patrolling strategy from the previous section is optimal for odd f and almost optimal for even f .

Theorem 2. *For any k and f such that $f < k/2 - 1$ we have $\mathfrak{S}_k^f(\mathcal{I}) \geq \frac{f+1}{k-f-1}$.*

Proof. Partition the unit interval into the following three segments

$$I_L := \left[0, \frac{f+1}{2(k-f-1)}\right], \quad I_M := \left(\frac{f+1}{2(k-f-1)}, 1 - \frac{f+1}{2(k-f-1)}\right), \quad I_R := \left[1 - \frac{f+1}{2(k-f-1)}, 1\right].$$

By the condition $f < k/2 - 1$ in the hypothesis of the theorem the three sub-segments should not have a non-trivial overlap. Before proving the theorem we derive a crucial claim.

Claim. If $\mathfrak{S}_k^f(\mathcal{I}) < \frac{f+1}{k-f-1}$ then at each time moment during the patrolling there must be at least $f+1$ robots in each of the segments I_L and I_R .

By symmetry it is sufficient to prove this claim for the segment I_L . Suppose that $\mathfrak{S}_k^f(\mathcal{I}) < \frac{f+1}{k-f-1}$ and assume on the contrary that at some time, say t_0 , we have at most f robots in the segment I_L . If an adversary makes all of these robots faulty then it would follow that no reliable robot could visit the endpoint 0 during the entire time interval $[t_0 - |I_L|, t_0 + |I_L|]$, where $|I_L|$ denotes the length of the interval I_L . Therefore the idle time at the endpoint 0 would be larger than $2|I_L| = \frac{f+1}{k-f-1}$, which contradicts the hypothesis of the claim.

From the Claim above we see that at all times each of the two intervals I_L and I_R contains at least $f+1$ robots. Since at each time moment at least $2f+2$ robots must be visiting I_L , I_R , the open interval I_M must always contain at most $k-2f-2$ robots. Applying Lemma 1 to the set E' consisting of the segment I_M , since $|I_M| = \left(1 - \frac{f+1}{k-f-1}\right)$ we have

$$\mathfrak{S}_k^f(\mathcal{I}) \geq \frac{(f+1)|I_M|}{(k-2f-2)} = \frac{(f+1)\left(1 - \frac{f+1}{k-f-1}\right)}{(k-2f-2)} = \frac{f+1}{k-f-1}.$$

■

3 Idleness of Arbitrary Graphs

In this section we study upper and lower bounds for patrolling times on general graphs. First we prove a theorem associating the patrolling time to the length of a Chinese Postman Tour on the graph. Next we use the results in Section 2 concerning line segments so as to determine asymptotic bounds on the patrolling time for arbitrary graphs. The efficiency of the proposed strategy is arbitrarily close to the optimal one when k is sufficiently large.

3.1 A general result and algorithm

First we prove the following theorem and approximation patrolling strategy on arbitrary graphs.

Theorem 3. *For any connected graph G , and $k \geq 2$ robots, at most f of which are faulty, ($f \leq k - 1$), we have that*

$$\frac{(f + 1)|E|}{k} \leq \mathfrak{S}_k^f(G) \leq \frac{(f + 1)CPT(G)}{k}.$$

Proof. The upper bound is implied from the following patrolling algorithm:

1. Select any Chinese Postman Tour of G .
2. Have the robots patrol the graph by placing them equidistant along the Chinese Postman Tour.

It is clear that the respective distances between consecutive robots will be $\frac{CPT(G)}{k}$. The worst case idle time occurs when we have f consecutive faulty robots. In this case the resulting idle time will never exceed $\frac{(f+1)CPT(G)}{k}$, which proves the upper bound.

The lower bound follows directly from Lemma 1 applied to E' being the set of all edges of G . ■

As a corollary of Theorem 3 we obtain the following tight (and simple) expression for the value of the idleness for Eulerian graphs.

Corollary 1 (Idleness for Connected Eulerian Graphs). *For any connected Eulerian graph G , and $k \geq 2$ robots, at most f of which are faulty ($f \leq k - 1$), we have that $\mathfrak{S}_k^f(G) = \frac{(f+1)|E|}{k}$.*

The claim is immediate since in this case $|E| = CPT(G)$.

3.2 Hardness of Computing the Idleness

To show the hardness of our problem in general graphs, we restrict ourselves to the special case of $k = 3$ robots with exactly $f = 1$ fault. We will now prove that the problem of computing the idleness $\mathfrak{S}_3^1(G)$ is NP-hard for general graphs with unit-length edges. The proof proceeds by reduction from 3-Edge-Coloring in Cubic Graphs (3ECC), a well-known NP-complete problem (see [19]).

First we show the following auxiliary result which partially characterizes graphs having minimum possible idleness \mathfrak{S}_3^1 with respect to the total length of their edges. For a graph $H = (V, E)$ and $E' \subseteq E$, denote by $H[E'] \subseteq H$ the connected subgraph of H with edge set E' .

Lemma 2. *Let $H = (V, E)$ be a graph with unit-length edges.*

- (i) *If $\mathfrak{S}_3^1(H) = \frac{2}{3}E(H)$, then there exists a partition of the edge set $E = E_1 \cup E_2 \cup E_3$ such that each of the graphs $H[E \setminus E_i]$, $1 \leq i \leq 3$, is semi-Eulerian (i.e., connected and with at most two vertices of odd degree).*

(ii) Conversely, if there exist a decomposition of the edge set: $E = E_1 \cup E_2 \cup E_3$, such that $|E_1| = |E_2| = |E_3| = \frac{1}{3}E(H)$ and each of the graphs $H[E \setminus E_i]$, $1 \leq i \leq 3$, is Eulerian, then $\mathfrak{S}_3^1(H) = \frac{2}{3}E(H)$.

We use Lemma 2 to show the following theorem.

Theorem 4 (Hardness of Computing the Idleness). *It is NP-hard to decide whether for a given graph H with unit-length edges we have $\mathfrak{S}_3^1(H) = \frac{2}{3}E(H)$.*

The above theorem shows that the problem of computing the optimal idle time for patrolling with unreliable robots is NP-hard in general. For an unbounded number of robots k (i.e, when k is treated as part of the input) and graphs with edges of integer length, the decision problem belongs to PSPACE, but we do not know whether it belongs to NP. We leave this as an open problem.

3.3 Characterizing graphs with minimum idle time

We close this section by considering some properties which hold for graphs with small idle time in fault-tolerant patrolling, and giving examples of classes of such graphs.

For the case of 3 robots, some classes of graphs minimizing idle time \mathfrak{S}_3^1 (with respect to the total length of edges) are given by Lemma 2. An example of such a class is known in the literature under the name of Kotzig graphs [22]. A graph is *Kotzig* if it is 3-regular and admits a decomposition into three matchings M_1, M_2, M_3 such that $E = M_1 \cup M_2 \cup M_3$ and for each pair $i \neq j$, the union $M_i \cup M_j$ forms a Hamiltonian cycle of the graph. By Lemma 2(ii), we immediately obtain that Kotzig graphs have the minimum possible idleness \mathfrak{S}_3^1 in the class of cubic graphs, i.e., $\mathfrak{S}_3^1 = \frac{2}{3}E = \frac{2}{3} \cdot \frac{3}{2}n = n$. Interestingly, this idleness \mathfrak{S}_3^1 is also best possible in the sense that we cannot obtain better idle time if we know beforehand which of the three robots is faulty, and attempt to solve the problem only for two non-faulty robots: for Kotzig graphs, we have $\mathfrak{S}_2^0 = \frac{CPT}{2} = \frac{2n}{2} = n$.

Corollary 2 (Idleness of Kotzig Graphs). $\mathfrak{S}_3^1(G) = \mathfrak{S}_2^0(G) = n$, for n -vertex Kotzig graphs G .

We remark that we do not know of a complete structural characterization of all graphs having minimum idleness $\mathfrak{S}_3^1 = \frac{2}{3}E$. The characterization from Lemma 2 is only partial, and the distinction between semi-Eulerian and Eulerian graphs in claims (i) and (ii) of the Lemma 2 is important. For example, when the patrolled graph is a cycle with 3 unit-length edges, we have $\mathfrak{S}_3^1 = \frac{2}{3}E = 2$, whereas this graph does not admit a decomposition $E = E_1 \cup E_2 \cup E_3$ into non-empty sets E_1, E_2, E_3 such that each of the graphs induced by $E \setminus E_i$ is Eulerian. On the other hand, when the patrolled graph is a star with 3 unit-length edges, this graph admits a decomposition $E = E_1 \cup E_2 \cup E_3$ into single edges such that each of the graphs induced by $E \setminus E_i$ is semi-Eulerian, but this graph does not minimize idle time: by Theorem 3, $\mathfrak{S}_3^1 \geq \frac{CPT}{2} = 3 > \frac{2}{3}E$.

Variants of Lemma 2 can also be obtained for a larger number of robots. For k even, classes of graphs having minimum possible idle time $\mathfrak{S}_k^1 = \frac{2}{k}E$ include *Hamiltonian Decomposable Graphs*, i.e., k -regular graphs whose edge set can be partitioned into $k/2$ edge-disjoint Hamiltonian cycles [26].

4 Conclusion and Open Problems

We gave optimal fault-tolerant patrolling strategy for segments (for odd f) and Eulerian graphs. In all proposed strategies the collection of patrolmen is divided into sub-collections, each of the

sub-collections, forming a “cycle” of equally spaced robots walking around a portion of the graph (with some portions being covered by more than one sub-collection). Somewhat surprisingly, for a graph as simple as a segment, the optimal strategy consists of two sub-collections patrolling small sub-segments and the third sub-collection patrolling the entire segment (hence the points close to the endpoints being visited by the robots belonging to two sub-collections). We also proved that for some graphs finding an optimal patrolling strategy is NP-hard.

While optimal strategies for Eulerian graphs work for any ratio of faulty patrolmen, the strategies for segments assume the maximal faulty robots ratio to be slightly smaller than half the total of all robots. One open question is to give optimal patrolling strategies for segments when the faulty robot ratio is high. There are plenty of open questions concerning different models of patrolling: robot failures may be dynamic, failures may happen with given probability, robots may have non-zero visibility radii, or may be allowed to communicate. Some questions, like robots with distinct patrolling speeds and two-dimensional domains may be hard.

Acknowledgements

This work was partially supported by NSERC grants. Research on this problem was initiated at the MITACS “International Problem Solving Workshop” held on July 16-20, 2012, in Vancouver, BC, Canada. The authors would like to express their deepest appreciation for the generous support of MITACS.

References

1. N. Agmon, S. Kraus, and G. A. Kaminka. Multi-robot perimeter patrol in adversarial settings. In *ICRA*, pages 2339–2345, 2008.
2. N. Agmon and D. Peleg. Fault-tolerant gathering algorithms for autonomous mobile robots. *SIAM Journal on Computing*, 36(1):56–82, 2006.
3. S. Alpern, A. Morton, and K. Papadaki. Optimizing randomized patrols. *Operational Research Group, London School of Economics and Political Science*, 2009.
4. S. Alpern, A. Morton, and K. Papadaki. Patrolling games. *Operations research*, 59(5):1246–1257, 2011.
5. F. Amigoni, N. Basilico, N. Gatti, A. Saporiti, and S. Troiani. Moving game theoretical patrolling strategies from theory to practice: An USARSim simulation. In *ICRA*, pages 426–431, 2010.
6. J. Bondy and U. Murty. *Graph Theory (GTM vol 244)*. New York: Springer, 2008.
7. Y. Chevaleyre. Theoretical analysis of the multi-agent patrolling problem. In *IAT*, pages 302–308, 2004.
8. R. Cohen and D. Peleg. Convergence properties of the gravitational algorithm in asynchronous robot systems. *SIAM J. Comput.*, 41(1):1516–1528, 2005.
9. R. Cohen and D. Peleg. Convergence of autonomous mobile robots with inaccurate sensors and movements. *SIAM Journal on Computing*, 38(1):276–302, 2008.
10. J. Czyzowicz, L. Gasieniec, A. Kosowski, and E. Kranakis. Boundary patrolling by mobile agents with distinct maximal speeds. *ESA 2011*, pages 701–712, 2011.
11. X. Défago, M. Gradinariu, S. Messika, and P. Raïpin-Parvédy. Fault-tolerant and self-stabilizing mobile robots gathering. *Distributed Computing*, pages 46–60, 2006.
12. Y. Dieudonné, A. Pelc, and D. Peleg. Gathering despite mischief. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 527–540. SIAM, 2012.
13. A. Dumitrescu, A. Ghosh, and C. D. Tóth. On fence patrolling by mobile agents. *Electr. J. Comb.*, 21(3):P3.4, 2014.
14. 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.
15. 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.

16. Y. Elor and A. M. Bruckstein. Autonomous multi-agent cycle based patrolling. In *ANTS Conference*, pages 119–130, 2010.
17. M. Garey and D. Johnson. *Computers and intractability*, volume 174. Freeman San Francisco, CA, 1979.
18. N. Hazon and G. A. Kaminka. On redundancy, efficiency, and robustness in coverage for multiple robots. *Robotics and Autonomous Systems*, 56:1102–1114, 2008.
19. I. Holyer. The NP-completeness of edge-coloring. *SIAM J. Comput.*, 10(4):718–720, 1981.
20. T. Izumi, S. Souissi, Y. Katayama, N. Inuzuka, X. Défago, K. Wada, and M. Yamashita. The gathering problem for two oblivious robots with unreliable compasses. *SIAM Journal on Computing*, 41(1):26–46, 2012.
21. A. Kawamura and Y. Kobayashi. Fence patrolling by mobile agents with distinct speeds. In *Proceedings of ISAAC 2012*, December 19–21, 2012.
22. A. Kotzig. Hamilton graphs and hamilton circuits. In *Theory of Graphs and its Applications, Proceedings of the Symposium of Smolenice*, pages 63–82. Publ. House Czechoslovak Acad. Sci., 1964.
23. 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.
24. A. Marino, L. Parker, G. Antonelli, F. Caccavale, and S. Chiaverini. A fault-tolerant modular control approach to multi-robot perimeter patrol. In *Robotics and Biomimetics (ROBIO)*, pages 735–740, 2009.
25. 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.
26. J. Park and H. Kim. Dihamiltonian decomposition of regular graphs with degree three. In *WG*, pages 240–249. Springer, 1999.
27. F. Pasqualetti, A. Franchi, and F. Bullo. On optimal cooperative patrolling. In *CDC*, pages 7153–7158, 2010.
28. D. Portugal and R. P. Rocha. A survey on multi-robot patrolling algorithms. In *Technological Innovation for Sustainability - Second IFIP WG 5.5, Costa de Caparica, Portugal, February 21-23, 2011. Proceedings*, pages 139–146, 2011.
29. S. Souissi, X. Défago, and M. Yamashita. Gathering asynchronous mobile robots with inaccurate compasses. *Principles of Distributed Systems*, pages 333–349, 2006.
30. Y. Yang, S. Souissi, X. Défago, and M. Takizawa. Fault-tolerant flocking for a group of autonomous mobile robots. *Journal of Systems and Software*, 84(1):29–36, 2011.
31. V. Yanovski, I. A. Wagner, and A. M. Bruckstein. A distributed ant algorithm for efficiently patrolling a network. *Algorithmica*, 37(3):165–186, 2003.

A Appendix

Proof. (Lemma 1) Fix any time interval of length $\mathfrak{S}^* = \mathfrak{S}_k^f(\mathcal{I}, \mathcal{P})$ and consider the walk of robots during this interval. The sum σ of lengths of trajectories of r robots within E' is at most $r\mathfrak{S}^*$. Consider first the case when $\sigma < (f+1)|E'|$. Then there exists some point p of E' which is visited by at most f robots. Consequently, the adversary can make all these robots faulty and point p is not visited during a time interval of length \mathfrak{S}^* by any reliable robot proving the claim of the lemma. In the remaining case we have $\sigma \geq (f+1)|E'|$. Since $\sigma \leq r\mathfrak{S}^*$ we have $\mathfrak{S}^* \geq \frac{(f+1)|E'|}{r}$. This proves Lemma 1. ■

Proof. (Lemma 2) First, note that $\mathfrak{S}_3^1(H) \geq \frac{2}{3}E(H)$ by Theorem 3.

(i) Suppose that $\mathfrak{S}_3^1(H) = \frac{2}{3}E(H)$. Consider the time interval $I = [0, \frac{2}{3}E(H))$. For simplicity of notation only, we assume that time 0 is chosen so that at time 0 all three robots are in “general positions”, i.e., not located on vertices. Observe that there cannot exist a point on an edge of the graph which was only visited by one robot during time interval I . Indeed, suppose that point x was only visited by robot i , $1 \leq i \leq 3$; then, if robot i turns out to be faulty, point x would be uncovered throughout I , hence $\mathfrak{S}_3^1(H) > \frac{2}{3}E(H)$, a contradiction. Now, since the total distance covered by all robots in time interval I is $2E(H)$, it follows from the pigeon-hole principle that each point is visited by precisely two distinct robots during interval I , and moreover, no point can be visited by any robot more than once during I . Since a robot which changes direction at a point other than a vertex during I traverses some point twice, it follows that a robot entering an edge during I has to continue along this edge until it reaches its other endpoint. Thus, for each robot i there can exist at most two edges for which this robot explored a proper and non-empty subset of points during interval I : the edge $e_i(0)$ on which robot i was located at time 0, and the edge $e_i(\frac{2}{3}E(H))$ on which i was located at time $\frac{2}{3}E(H)$.

With every edge $e \in E$ we now associate an integer $c(e) \in \{1, 2, 3\}$ as follows. If edge e was explored completely by exactly two robots during interval I , then $c(e)$ is defined as the label of the only robot which did *not* traverse e during interval I . Otherwise, $c(e)$ is set as the label of either of the two robots which did not traverse e in full during interval I . Define $E_i = \{e \in E : c(e) = i\}$, $1 \leq i \leq 3$. Then, set $E \setminus E_i$ contains all edges such that all of their points were traversed in full by robot i during interval I , and does not contain any edges which were not visited at all by robot i during interval I . The edges $e_i(0)$ and $e_i(\frac{2}{3}E(H))$, which may have been traversed in part by robot i during interval I , may either belong to E_i or to $E \setminus E_i$. Recalling that no point on an edge is visited more than once by robot i during interval I , we obtain that the traversal performed by robot i during time interval I , appropriately truncated (or extended) to exclude (or include) the whole of its starting edge $e_i(0)$ and final edge $e_i(\frac{2}{3}E(H))$ determines a semi-Eulerian traversal of graph $H[E \setminus E_i]$. Hence, all of the graphs $H[E \setminus E_i]$ are semi-Eulerian.

(ii) Given a decomposition $E = E_1 \cup E_2 \cup E_3$ satisfying the assumptions of the claim, we define the route of robot i to be a perpetual traversal of a fixed Eulerian circuit of graph $H[E \setminus E_i]$ at maximal speed. The length of this circuit is $\frac{2}{3}E(H)$. Consider an arbitrary point x on any edge of the graph. The circuits of exactly two robots pass through point x . Hence, even if one of these robots is faulty, point x will be visited by the other robot with a time spacing of exactly $\frac{2}{3}E(H)$ between visits. It follows that $\mathfrak{S}_3^1(H) \leq \frac{2}{3}E(H)$, which completes the proof of Lemma 2. ■

Proof. (Theorem 4) We show that the posed question is NP-hard by reduction from 3-Edge-Coloring in Cubic Graphs (3ECC). Let the 3ECC instance $G = (V, E)$ be a given cubic graph

with $V = \{v_1, \dots, v_n\}$. We will define the graph $H = (V_H, E_H)$, in which we ask about \mathfrak{S}_3^1 , by taking 7 copies of graph G and inserting auxiliary subgraphs (gadgets) in place of their edges, so as to connect all these copies to one auxiliary vertex r used for garbage collection. Initially, let $V_H = V^{(1)} \cup \dots \cup V^{(7)} \cup \{r\}$, where each $V^{(j)} = \{v_1^{(j)}, \dots, v_n^{(j)}\}$, $1 \leq j \leq 7$, is a copy of V , and let E_H be empty. Next, for each edge $e = \{u, v\} \in E$, for all copies $1 \leq j \leq 7$, we insert in between vertices $u^{(j)}$, $v^{(j)}$, and r of H the gadget X having 21 edges and 12 additional internal vertices, presented in Figure 2. We will denote this copy of gadget X by $X^{(j)}(e)$. To complete the proof, we




Fig. 2. Gadget X used in the construction of graph H in the hardness proof.

now show that G is 3-edge-colorable if and only if $\mathfrak{S}_3^1(H) = \frac{2}{3}E(H)$.

(\Rightarrow) Suppose that G is 3-edge-colorable and fix a coloring of its edges $c: E \rightarrow \{1, 2, 3\}$. Based on coloring c , we define a partition of the edges of H , $E(H) = E_1 \cup E_2 \cup E_3$, as follows. For each edge $e \in |E|$ having color $c(e) = i$, we consider the 7 corresponding gadgets $X^{(j)}(e)$ in H , and for each such gadget assign its edges to sets E_1, E_2, E_3 as shown in Figure 3. The obtained decomposition




Fig. 3. Assignment of edges of gadget $X^{(j)}(e) \subseteq H$ to sets E_1, E_2, E_3 for an edge e having color i . Edges marked i in the figure belong to E_i , edges marked $i+1$ belong to $E_{(i \bmod 3)+1}$, and edges marked $i+2$ belong to $E_{((i+1) \bmod 3)+1}$.

$E(H) = E_1 \cup E_2 \cup E_3$ has the following properties:

- $|E_1| = |E_2| = |E_3| = \frac{1}{3}E(H)$. Indeed, for any gadget $X^{(j)}(e)$, regardless of the color of e , we assign 7 edges of the gadget to E_1 , 7 edges to E_2 , and 7 edges to E_3 . Hence, the sets E_1, E_2, E_3 have the same number of elements.
- For $1 \leq i \leq 3$, graph $H[E(H) \setminus E_i]$ has all nodes of even degree. Indeed, for the internal nodes of each gadget $X^{(j)}(e)$, their degree in $H[E(H) \setminus E_i]$ is 2. Since c is a coloring, each node from $V^{(j)}$ is adjacent to exactly one edge from each of the sets E_1, E_2, E_3 , hence its degree in $H[E(H) \setminus E_i]$ is also 2. The degree of the last remaining node r must be even as well.
- For $1 \leq i \leq 3$, graph $H[E(H) \setminus E_i]$ is connected. Observe that in each gadget $X^{(j)}(e)$, the subgraph $X^{(j)}(e)[E(H) \setminus E_i]$ is connected and spans all vertices of $X^{(j)}(e)$, including r . Since all these gadgets share node r , $H[E(H) \setminus E_i]$ is also connected.

So, we have that $|E_1| = |E_2| = |E_3| = \frac{1}{3}E(H)$ and that graph $H[E(H) \setminus E_i]$ is Eulerian, thus, applying Lemma 2(ii), we obtain $\mathfrak{S}_3^1(H) = \frac{2}{3}E(H)$.

(\Leftarrow) Suppose that $\mathfrak{S}_3^1(H) = \frac{2}{3}E(H)$. It follows from Lemma 2(i) that there exists a decomposition $E(H) = E_1 \cup E_2 \cup E_3$ such that each of the graphs $H[E(H) \setminus E_i]$, $1 \leq i \leq 3$, is semi-Eulerian. Let $V_H^{(j)}$ be the set of nodes of the connected component of $H \setminus \{r\}$ containing $V^{(j)}$, $1 \leq j \leq 7$, and let $H_+^{(j)} = H[V_H^{(j)} \cup \{r\}]$. Since $7 > 3 \cdot 2$, it follows from the pigeon-hole principle that there must exist some $1 \leq j \leq 7$, such that for all $1 \leq i \leq 3$, all vertices from $V_H^{(j)}$, have even degree in $H[E(H) \setminus E_i]$. Thus, for the considered value of j , each of the graphs $H_+^{(j)}[E(H_+^{(j)}) \setminus E_i]$, $1 \leq i \leq 3$, is Eulerian. Hence, the following properties hold:

- Each vertex $v^{(j)} \in V^{(j)}$ is adjacent in $H_+^{(j)}$ to exactly one edge from each of the sets E_1, E_2, E_3 . Indeed, each vertex $v^{(j)} \in V^{(j)}$ is of degree 3, and since $H_+^{(j)}[E(H_+^{(j)}) \setminus E_i]$ is Eulerian, $v^{(j)}$ may be adjacent to either 1 or 3 edges from E_i , for all $1 \leq i \leq 3$. It follows that $v^{(j)}$ must be adjacent to exactly one edge from each of these three sets.
- For each gadget $X^{(j)}(e) \subseteq H_+^{(j)}$, where $e = \{u, v\}$, we have that the edge $\{u^{(j)}, \bar{u}^{(j)}\}$ in this gadget belongs to E_i if and only if edge $\{\bar{v}^{(j)}, v^{(j)}\}$ belongs to the same set E_i , $1 \leq i \leq 3$. This property can be shown for each gadget separately, by exhaustively testing all possible assignments of edges from $X^{(j)}(e)$ to disjoint sets E_1, E_2, E_3 so that $H_+^{(j)}[E(H_+^{(j)}) \setminus E_i]$ has all nodes of even degree and does not contain small connected components (cycles of 4 or 6 vertices).

We now define a labeling of edges of G with labels $\{1, 2, 3\}$ as follows. Edge $e = \{u, v\} \in E$ gets label i if in the gadget $X^{(j)}(e) \subseteq H$, edges $\{u^{(j)}, \bar{u}^{(j)}\}$ and $\{\bar{v}^{(j)}, v^{(j)}\}$ belong to set E_i . By the above properties, this assignment yields a well-defined 3-edge-coloring of G , which completes the proof. \blacksquare