



Optimal placement of base stations in border surveillance using limited capacity drones

S. Bereg^{a,1}, J.M. Díaz-Báñez^{b,2,3}, M. Haghpanah^a, P. Horn^{c,4}, M.A. Lopez^{d,5},
N. Marín^{e,*,6}, A. Ramírez-Vigueras^{f,7}, F. Rodríguez^{b,3}, O. Solé-Pi^g, A. Stevens^c,
J. Urrutia^{f,8}

^a Department of Computer Science, University of Texas at Dallas, Richardson, TX 75080, United States of America

^b Departamento de Matemática Aplicada II, Universidad de Sevilla, Spain

^c Department of Mathematics, University of Denver, United States of America

^d Department of Computer Science, University of Denver, United States of America

^e Posgrado en Ciencia e Ingeniería de la Computación, Universidad Nacional Autónoma de México, México

^f Instituto de Matemáticas, Universidad Nacional Autónoma de México, México

^g Departamento de Matemáticas, Facultad de Ciencias, Universidad Nacional Autónoma de México, México

ARTICLE INFO

Article history:

Received 1 September 2021

Received in revised form 29 April 2022

Accepted 10 June 2022

Available online 16 June 2022

Communicated by T. Calamoneri

Keywords:

Border protection

Optimal location

Unmanned aerial vehicles

Algorithms

Polygons

ABSTRACT

Imagine an island modeled as a simple polygon \mathcal{P} with n vertices whose coastline we wish to monitor. We consider the problem of building the minimum number of refueling stations along the boundary of \mathcal{P} in such a way that a drone can follow a polygonal route enclosing the island without running out of fuel. A drone can fly a maximum distance d between consecutive stations and is restricted to move either along the boundary of \mathcal{P} or its exterior (i.e., over water). We present an algorithm that, given \mathcal{P} , finds the locations for a set of refueling stations whose cardinality is at most the optimal plus one. The time complexity of this algorithm is $O(n^2 + \frac{L}{d}n)$, where L is the length of \mathcal{P} . We also present an algorithm that returns an additive ϵ -approximation for the problem of minimizing the fuel capacity required for the drones when we are allowed to place k base stations around the boundary of the island; this algorithm also finds the locations of these refueling stations. Finally, we propose a practical discretization heuristic which, under certain conditions, can be used to certify optimality of the results.

© 2022 Elsevier B.V. All rights reserved.

* Corresponding author.

E-mail addresses: bsp@utdallas.edu (S. Bereg), dbanez@us.es (J.M. Díaz-Báñez), Mohammadreza.Haghpanah@utdallas.edu (M. Haghpanah), paul.horn@du.edu (P. Horn), mario.lopez@du.edu (M.A. Lopez), nestaly@ciencias.unam.mx (N. Marín), adriana.rv@im.unam.mx (A. Ramírez-Vigueras), frrodriguex@us.es (F. Rodríguez), oriolandreu@ciencias.unam.mx (O. Solé-Pi), alex.stevens@du.edu (A. Stevens), urrutia@matem.unam.mx (J. Urrutia).

¹ Supported in part by NSF award CCF-1718994.

² Partially supported by European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement #734922.

³ Partially supported Ministerio de Ciencia e Innovación CIN/AEI/10.13039/501100011033/(PID2020-114154RB-I00).

⁴ Supported in part by Simons Collaboration Grant #525039.

⁵ Supported in part by a University of Denver John Evans Research Award.

⁶ Supported by SEP-CONACyT of Mexico.

⁷ Partially supported by PAPIIT IN105221, Programa de Apoyo a la Investigación e Innovación Tecnológica UNAM.

⁸ Supported in part by PAPIIT IN102117 Programa de Apoyo a la Investigación e Innovación Tecnológica, UNAM.

<https://doi.org/10.1016/j.tcs.2022.06.024>

0304-3975/© 2022 Elsevier B.V. All rights reserved.




Fig. 1. Example. A polygon \mathcal{P} and d -hull \mathcal{C} (black dashed line) that uses an optimal set of base stations (solid circles) for the MinStation problem.

1. Introduction

The rapid development and use of *Unmanned Aerial Vehicles* (UAVs), commonly called drones, in many activities of our daily life has created a need for the development of new algorithms to optimize their use. A factor that is common to most types of drones is their relatively short flying range due mostly to their restricted energy capacity [23].

Border patrolling is one typical application of air surveillance systems where the deployment of drones has become a natural choice for providing monitoring, surveillance, and search and rescue services for the protection of human lives or natural resources [2,17,21,26,30]. In this context, a team of UAVs can be deployed along the boundaries of a region to collect useful information, such as images or videos, and send it to the nearest control center. The energy limitation of small UAVs prevents them from remaining in flight for long periods of time. Thus, recharging stations, platforms where the drone can autonomously land to recharge its battery before continuing its mission, have been recently introduced. However, the cost of those platforms remains a significant obstacle and, consequently, it is important to reduce their number.

Inspired by this type of applications, we study in this paper the following geometric optimization problems:

The MinStation Problem: Suppose that we want to guard the border of an island \mathcal{I} whose boundary is modeled by a simple polygon \mathcal{P} using a set of drones that can fly a distance d before they need refueling. Our drones can fly over the boundary of \mathcal{I} or over small sectors of the sea surrounding it, but not over the interior of \mathcal{I} . Our objective is to place a set $S = \{s_0, \dots, s_{k-1}\}$ of k refueling base stations with *minimum cardinality* k and located on the boundary of \mathcal{I} , such that when a drone visits all the refueling stations it travels a closed curve that encloses \mathcal{P} ; the flying distance between s_i and s_{i+1} is at most d , with addition taken mod k . See Fig. 1 for an example. We will refer to S as an *optimal solution*. A set $S' = \{s'_0, \dots, s'_k\}$ with $k+1$ refueling stations will be called a *quasi-optimal solution*.

The MinDistance Problem: Suppose that we have a budget that allows us to build k refueling stations. Find the smallest d such that we can build k refueling stations that allow a drone with flight capacity d to guard the border of the given island.

Solutions to these problems can be implemented using fleets of small drones with limited capacity, resulting in cheaper systems that use less resources and increase the frequency with which the drones patrol the border. We consider that a drone can replace the used battery with a fresh one, rather than charging the battery on site. The used batteries can be recharged when a charging station is not servicing drones. Thus, a drone will not occupy a charging station for a long time and there is at most one drone at each station at the same time. Our main contributions are as follows:

- We give an algorithm, `OPTSol`, with complexity $O(n^2 + \frac{1}{d}n)$, such that if s_0 is a fixed point on the convex hull of \mathcal{P} , $CH(\mathcal{P})$, finds an optimal solution S to the MinStation Problem under the restriction that $s_0 \in S$; here, L is the perimeter of \mathcal{P} . This yields either an optimal or quasi-optimal solution to the unconstrained MinStation Problem (without requiring $s_0 \in S$). The problem of finding an optimal unconstrained solution is equivalent to that of finding the location of a single station in an optimal solution. We leave as an open problem that of designing a polynomial time algorithm for the general unconstrained case.
- For the *MinDistance Problem*, we show how to approximate an optimal solution up to an additive constant. The approach is based on an algorithm, `AppSol`, which solves a discretized version of the MinStation Problem.
- We implemented `AppSol` and ran experiments on a polygon corresponding to an island (using GIS data); in many cases, the algorithm returns a certifiably optimal solution on this real data.

1.1. Related work

Drones have become the natural choice for the deployment of air surveillance systems [2,21,24]. We mention three areas where problems close to ours can be found: facility location, wireless networks, and computational geometry.

In facility location problems we are interested in finding the best places to locate a set of resources (e.g. airports, pharmacies, gas stations, markets, etc.) to better serve a community, as well as creating optimal routes to visit them. The facilities can be isolated points in space [19] or 2-dimensional structures such as straight lines, line-segments, polygonal

curves, or circles [10]. The location-routing problem is a research area within locational analysis, with the distinguishing property of paying special attention to vehicle routing aspects [27]. Recently, applications in the aerial robotics community, such as finding the best places to locate drone base stations and creating flying routes for the drones, have arisen in areas such as border patrolling [31]. Cities on the borders of countries are modeled as demand points, and airports are considered as base stations or hubs. In [39] the authors studied a base location and path planning problem in maritime target reconnaissance problems. Their problem is formulated as an integer linear program where the total score obtained from visiting points of interest by flight routes of drones is maximized; a novel ant colony optimization metaheuristic approach is proposed. In a more recent paper [22], both capacity constraints on base stations and endurance limitations on drones are taken into account and two heuristic algorithms are designed to solve the problem. A similar problem is considered in [8]. They investigate the 3D location problem of multiple drone base stations as well as the allocation of their dynamic capacities to the users. The service provided by the stations is also dynamic in terms of the data rate level provided to the users. In addition to border patrolling, the problem of deploying a number of charging stations to cover a demand region has been considered in several areas. For example, [15] proposes a coverage model to figure out the optimal positions of a given number of charging stations from a discretized candidate set, with the objective of maximizing the coverage of customers. More recently, a connectivity requirement on the stations is required in [16] in order to guarantee the delivery to customers located far from a depot. They argue that any two neighbor charging stations should be within a certain range such that a drone with a fully charged battery can reach one from the other. In other words, it is required that the deployed charging stations should be connected to the depot. Then, a drone that departs from the depot can arrive at any charging station via a subset of other charging stations, and it can then service customers near this charging station.

Another field of research close to our work can be found in wireless sensor networks. In [18], the authors study the k -barrier problem: how to deploy a set of sensors in a belt region surrounding a castle in such a way that any intruder is detected by at least k sensors. In [4] the following problem is studied: Given that an intruder has been detected by a set of sensors, how can they be moved in an optimal way to prevent further intrusions? An interesting survey of problems similar to ours can be found in [1], where they study the problem of protecting several types of holes that can occur in a wireless sensor network, where a hole is a region not covered by the sensing disks of a set of sensors. In the same paper, other problems related to ours are considered, including routing in static and mobile sensor networks. See also [5,11,37].

Finally, in computational geometry, there is a whole area of research devoted to problems, collectively known as Art Gallery problems, which are related to ours. The oldest problem in this area requires finding a minimum set of points S within an art gallery, usually modeled by a polygon \mathcal{P} , such that every point in the polygon is visible from at least one point in S . Many variations of this first problem have since been studied and the interested reader is referred to the book by O'Rourke [29] or the surveys by Shermer [32] and Urrutia [36]. Some variations of the art gallery problem, more closely related to our problem, consider the use of mobile guards. The *Watchman Route problem*, introduced in [7], is that of finding a path of minimum length that a guard can follow in order to guarantee that every point within \mathcal{P} is visible from some point in the path. There are two main variations of this problem, one in which the starting point of the route is given [35], as well as the unrestricted case [33]. Even more closely related to our problem is that of finding the minimum watchman route in the exterior of \mathcal{P} [28]. Some variations of the watchman route problem require the path to visit a set of k sites represented as polygons in \mathcal{P} . In the *Safari Route problem* [34] we are allowed to enter the sites, while in the *Zoo Keeper Route problem* [38,3] the guard has to visit their boundary but is not allowed to go inside (as when feeding an animal without entering its cage). Both problems are NP-hard in the general case, but can be solved in polynomial time if the sites are adjacent to the boundary of \mathcal{P} . The *Aquarium Keeper's problem*, studied in [9], deals with the problem of computing the shortest closed path inside \mathcal{P} which visits each of its edges at least once.

2. Terminology and problem formulation

In what follows a polygon \mathcal{P} is represented by a sequence $\langle p_0, \dots, p_{n-1} \rangle$ of its vertices given in clockwise order around its boundary. Thus, the edges of \mathcal{P} are the line segments $\overline{p_i p_{i+1}}$, with addition taken mod n . We assume that our polygons are always *simple*, i.e. that no two non-consecutive edges intersect. We use $Int(\mathcal{P})$ and $Ext(\mathcal{P})$ to denote, respectively, the interior and exterior of the region enclosed by \mathcal{P} , and use \mathcal{P} itself to refer to the boundary of this region (often referred to by $\partial\mathcal{P}$ in the literature). Accordingly, the length of \mathcal{P} is the sum of the lengths of its edges. A (polygonal) path is a sequence of points $\langle q_0, \dots, q_k \rangle$ together with the set of edges $\overline{q_i q_{i+1}}$, $i = 0, \dots, k-1$; the length of a path is the sum of the lengths of its edges.

Given two distinct points $a, b \in \mathcal{P}$ the interval $[a, b]$ is the set of points of \mathcal{P} traversed while moving from a to b in the clockwise direction along the boundary of \mathcal{P} . The distance $\delta_{\mathcal{P}}(a, b)$ between a and b in \mathcal{P} is the length of the interval $[a, b]$. Observe that since $a \neq b$, $[a, b] \neq [b, a]$, $[a, b] \cup [b, a] = \mathcal{P}$, and that $\delta_{\mathcal{P}}(a, b) + \delta_{\mathcal{P}}(b, a)$ is the length of \mathcal{P} .

The following definitions of what we will call d -paths and d -hulls arise from the restriction that the flight range of a drone is a fixed number d .

An open line segment contained in $Ext(\mathcal{P})$ joining two points $a, b \in \mathcal{P}$ will be called a *bridge*; if its length is at most d it is called a d -bridge of \mathcal{P} . Note that a drone cannot fly along a bridge of \mathcal{P} with length greater than d ; the base stations are restricted to be on \mathcal{P} , thus if a drone chooses to fly over a bridge with length greater than d it would run out of fuel and fall to the sea.




Fig. 2. The interval $[a, b]$ and a d -path $\pi_{x,y}$ joining x and y are shown in black dashed lines. The segment $\overline{p_i p_j}$ is a d -bridge contained in $\pi_{x,y}$.

A polygonal path joining two points $a, b \in \mathcal{P}$ is called a d -path if all of its edges are d -bridges of \mathcal{P} , or segments of edges of \mathcal{P} . We say that a polygon \mathcal{C} is a d -hull of \mathcal{P} if it encloses \mathcal{P} , and all of its edges are contained in edges of \mathcal{P} or are d -bridges of \mathcal{P} . Observe that a polygon has many (in fact an infinite number) of d -hulls, indeed \mathcal{P} is a d -hull of itself.

The drone distance $\delta(a, b)$ from a to b is the length of a shortest clockwise d -path joining a to b . As an example, Fig. 2 shows the shortest d -path from x to y . For simplicity, we will refer to the drone distance as the distance from a to b . Observe that $\delta(a, b)$ is in general different from $\delta(b, a)$. Further observe that the drone distance and the geodesic distance from a to b (understood as the length of the shortest clockwise path from a to b disjoint from $Int(\mathcal{P})$) coincide whenever one of them is at most d . Finally, note that if a drone with flight range d can fly between two points $a, b \in \mathcal{P}$ without recharging, then there is a d -path joining them of length at most d .

Our island guarding problem can now be restated as follows:

Problem 1 (MinStation). Given a polygon \mathcal{P} find a set of base stations $S = \{s_0, \dots, s_{k-1}\}$ with minimum cardinality such that for every $0 \leq i \leq k - 1$ there is a d -path π_i of length at most d joining s_i to s_{i+1} , and such that $\mathcal{C} = \pi_0 \cup \dots \cup \pi_{k-1}$ is a d -hull of \mathcal{P} ; $s_i \in \mathcal{P}$, $i = 0, \dots, k - 1$, addition taken mod k .

By a solution to the MinStation problem we refer simply to a set S of base stations together with the collection of d -paths π_i whose union is a polygon that encloses \mathcal{P} . Recall that a solution is optimal if it contains the least possible number of base stations, and quasi-optimal if it contains one more base station than an optimal solution.

We also study the next problem, a kind of dual problem to the MinStation problem. Suppose that we have a budget that allows us to build k base stations, and want to find the locations along \mathcal{P} where to build them such that the flight range of the drones used to patrol \mathcal{P} is minimized, formally:

Problem 2 (MinDistance). Given a simple polygon \mathcal{P} and an integer $k \geq 2$, find the smallest d and a set $S = \{s_0, s_2, \dots, s_{k-1}\}$ of k stations on \mathcal{P} such that, for $i = 0, \dots, k - 1$, there is a d -path π_i of length at most d joining s_i to s_{i+1} , addition taken mod k , and $\mathcal{C} = \pi_0 \cup \dots \cup \pi_{k-1}$ is a d -hull of \mathcal{P} .

Computing a d -hull that minimizes the number of base stations needed to solve the MinStation problem is more subtle than it may at first look. There are polygons \mathcal{P} for which, given d , the smallest number of base stations needed to solve the MinStation problem, lie on the shortest d -hull enclosing \mathcal{P} . An example is shown in Fig. 1. However, there are examples for which the stations of an optimal solution do not lie on the shortest d -hull enclosing \mathcal{P} . An example is given in Fig. 3. It is easy to see that placing a base station at any point other than the black points shown there, increases the number of base stations needed to solve the MinStation problem. In fact, it is not hard to construct polygons such that the number of stations required for the shortest d -hull is almost twice the number of stations given by the MinStation problem. This is the case, for example, for a star shaped polygon such that the distance between adjacent vertices on the boundary of the convex hull is $\frac{d}{2} + \epsilon$, for some arbitrarily small ϵ , as shown in Fig. 4.

We remark that in the optimal solutions of the MinStation and the MinDistance problems, the base stations lie on \mathcal{P} but not necessarily on vertices of \mathcal{P} or \mathcal{C} .

3. Preliminary results

Given a fixed point $s \in \mathcal{P}$ we define a total order $O_s(\mathcal{P}, \preceq)$ on the points in \mathcal{P} as follows:

1. for any point $a \in \mathcal{P}$, $s \preceq a$
2. for any $a, b \in \mathcal{P}$, both different from s , $a \preceq b$ if $[s, a] \subseteq [s, b]$. (Note that possibly $a = b$.)




Fig. 3. Example. The optimal d -hull requires 6 base stations. Replacing it with one with smaller perimeter increases the number of base stations to 7.




Fig. 4. Example. The shortest d -hull (dotted) requires almost twice as many stations as the optimal d -hull that solves the MinStation problem (dashed). This example can be extended to polygons with arbitrarily many vertices.

For convenience we will add an extra element s' to our order such that for any $a \in \mathcal{P}$, $a \preceq s'$; that is, s and s' are, respectively, the minimum and the maximum elements of $O_s(\mathcal{P}, \preceq)$. We can think of s' as a copy of s , and refer to $O_s(\mathcal{P}, \preceq)$ simply as \preceq .

Consider a point $s \in \mathcal{P}$ and the order \preceq it defines on the points on \mathcal{P} . We define a distance δ_d on the points on \mathcal{P} as follows:

1. $\delta_d(a, a) = 0$.
2. If $a \preceq b \in \mathcal{P}$, $\delta_d(a, b) = 1$ if there is a d -path of length at most d from a to b .
3. $\delta_d(a, b) = k$ if k is the smallest integer such that there is a sequence of points $p_0 = a, \dots, p_k = b$ such that $\delta_d(p_i, p_{i+1}) = 1$, $i = 0, \dots, k - 1$.

The following technical Lemma will be crucial in the proposed approach to solve the MinStation problem.

Lemma 1 (The Sandwich Lemma). Let $w, x, y, z \in \mathcal{P}$ such that $w \preceq x \preceq y \preceq z$ on \mathcal{P} , such that $\delta_d(w, y) \leq 1$, and $\delta_d(x, z) \leq 1$. Then $\delta_d(w, z) \leq 2$, $\delta_d(w, x) \leq 2$ and $\delta_d(y, z) \leq 2$.

Proof. Suppose that $\delta_d(w, z) > 1$, for otherwise we are finished. Since $w \preceq x \preceq y \preceq z$ the shortest d -paths $\pi_{w,y}$ and $\pi_{x,z}$ joining w to y , and x to z intersect, see Fig. 5 for an example. Let p be a point in the intersection of $\pi_{w,y}$ and $\pi_{x,z}$. If the distance $\delta(x, p)$ along $\pi_{x,z}$ between p and x is smaller than the distance $\delta(p, y)$ between p and y along $\pi_{w,y}$, then $\delta(w, p) + \delta(p, x) \leq 1$ and therefore $\delta_d(w, z) \leq 2$. The case when $\delta(p, x) \geq \delta(p, y)$ follows the same way. The inequalities $\delta_d(w, x) \leq 2$ and $\delta_d(y, z) \leq 2$ are proved in a similar way. \square

The Sandwich Lemma suggests that in an optimal solution to the MinStation problem, a drone flies around a non-crossing curve \mathcal{C} that encloses \mathcal{P} . We formalize this observation in the lemma that follows:




Fig. 5. Illustration of Lemma 1.

Lemma 2. Suppose that $s_0 \in \mathcal{P} \cap CH(\mathcal{P})$ and let $S = \{s_0, s_1, \dots, s_{k-1}\}$ be a solution to the MinStation problem that goes around P in the clockwise direction and which has the least number of stations among all solutions starting from s_0 . Then $s_0 \preceq s_1 \preceq s_2 \preceq \dots \preceq s_{k-1}$.

Proof. Assume that all of the π_i paths joining s_i to s_{i+1} are of minimum length. Since $\mathcal{C} = \pi_0 \cup \dots \cup \pi_{k-1}$ encloses \mathcal{P} , any point p in $CH(\mathcal{P})$ lies on \mathcal{C} . It is now easy to see that \mathcal{C} covers p exactly once. It follows now that s_0 is not in the interior of π_0 , and that π_0 is a simple curve that always advances in the clockwise direction along \mathcal{C} .

Now, suppose that $s_i \preceq s_{i-1}$ for some $i > 1$, and let j be the maximum value such that $s_j \preceq s_i$; that is, $s_j \preceq s_i \preceq s_{j+1}$. Using Lemma 1 it follows that $s_j \preceq s_i \preceq s_{j+1} \preceq s_{i-1}$. Thus, by Lemma 1, $\delta_d(s_j, s_i) \leq 2$, and since S is an optimal solution, it follows that $s_{j+1} = s_{i-1}$.

Let r be the minimum value such that $s_{i-1} \preceq s_r$. Then, we have that $s_i \preceq s_{r-1} \preceq s_{i-1} \preceq s_r$. It follows that $s_i = s_{r-1}$.

Now, since $s_{j+1} = s_{i-1}$ and $s_i = s_{r-1}$ we have that $s_j \preceq s_i \preceq s_{i-1} \preceq s_r$, where $\delta_d(s_j, s_{i-1}) = \delta_d(s_i, s_r) = 1$. Therefore, by Lemma 1, $\delta_d(s_j, s_r) \leq 2$. This is a contradiction, and thus $s_i \preceq s_{i+1}$ for all i . Hence, $\{s_0, s_1, \dots, s_r\}$ continues to make forward progress and the result follows. \square

A similar argument shows that for any optimal solution $S = \{s_0, s_1, \dots, s_{k-1}\}$ of the MinStation Problem (no longer subject to the condition that s_0 is fixed) \mathcal{C} is a simple closed curve.

4. The algorithm

We consider the following algorithm, which constructs a solution to the MinStation problem starting at a point $v \in \mathcal{P} \cap CH(\mathcal{P})$.

Algorithm 1: OPTSOL.

Input: Polygon \mathcal{P} , $s_0 \in \mathcal{P} \cap CH(\mathcal{P})$, $d > 0$

Output: The stations in an optimal or quasi-optimal d -hull for \mathcal{P}

1 Let $s_0 = s'_0 = y_{-1} = v$ and $y_0 = \max\{y : \delta_d(s_0, y) = 1\}$

2 Set $S_0 = \{s_0\}$ and $i = 0$

3 **repeat**

4 $i = i + 1$

5 $y_i = \max\{y : \delta_d(s_{i-1}, y) = 2\}$

6 $s_i = \text{any } s \in \{w : \delta_d(s_{i-1}, w) = 1 \text{ and } \delta_d(w, y_i) = 1\}$

7 Set $S_i = S_{i-1} \cup \{s_i\}$

until $y_i = s'_0$

return the last generated set $S = S_i$

We claim that if we further require that $v \in S$, then the set S returned is, indeed, an optimal solution to MinStation. On the other hand, we observe that this algorithm always gives a solution that is globally optimal or quasi-optimal (no longer subject to the restriction that $v \in S$).

Theorem 3. Given a starting point $s_0 \in \mathcal{P} \cap CH(\mathcal{P})$, if k is the least value such that $s_k = s'_0$, then the set of points $S = \{s_0, \dots, s_{k-1}\}$ returned by the OPTSOL algorithm is an optimal solution to the MinStation problem with the additional requirement that a base station be located at s_0 .

Proof. Suppose that S has more than one element, for otherwise our result is obvious. Suppose that $Z = \{z_0, \dots, z_{n-1}\}$ is an optimal solution for the MinStation problem such that there is a base station located at $s_0 = z_0$. We prove now that $n = k$




Fig. 6. The 2-hop problem. (a) s_{i+1} is a projection of the vertex p_j on the edge e . (b) s_{i+1} is a d -projection of the edge f on the edge e .

By Lemma 2, we may assume that $v = z_0 \leq z_1 \leq \dots \leq z_{n-1}$, that for all i , $\delta_d(z_i, z_{i+1}) = 1$, and that $\delta_d(z_{n-1}, s'_0) = 1$. Consider now the set $S = \{s_0 = v, \dots, s_{k-1}\}$ returned by the OPTSOL algorithm. Recall that $y_{k-1} = s'_0$. While the relationship between S and Z is unclear, the relationship between Z and $Y_{k-1} = \{y_{-1}, y_0, \dots, y_{k-1}\}$ is more straightforward; indeed we prove by induction that, for all i , $z_{i+1} \leq y_i$.

This clearly holds for $i = 0$, as $z_1 \leq y_0$ by definition. Now, suppose $z_i \leq y_{i-1}$. Let j be minimal so that $z_i \leq y_{j-1}$. Then $j \leq i$, and as $v = y_{-1} < z_i$ we have that $j \geq 1$. Now, by the minimality of j , we know that $y_{j-2} \leq z_i \leq y_{j-1}$ and by definition of y_{j-2} , we have that $s_{j-1} \leq y_{j-2}$. Combining, $s_{j-1} \leq z_i \leq y_{j-1}$. Then as $\delta_d(s_{j-1}, y_{j-1}) = 1$ by construction, Lemma 1 implies that $z_{i+1} \leq y_j \leq y_i$. This completes the inductive step, and hence the proof.

Therefore, the minimal k such that $\delta_d(s_k, s'_0) = 1$ and the minimal n such that $\delta_d(y_n, s'_0) = 1$ are the same, and thus $S = \{s_0 = v, \dots, s_{k-1}\}$ is an optimal solution to the MinStation problem, with the additional requirement that there is a base station at $s_0 = v$. \square

Theorem 4. Let $s_0 = v \in \mathcal{P} \cap CH(\mathcal{P})$. The set $S = \{s_0, \dots, s_{k-1}\}$ returned by the OPTSOL algorithm is an optimal solution or a quasi-optimal solution to MinStation problem.

Proof. Suppose that $Z = \{z_0, \dots, z_{n-1}\}$ is an optimal solution to the MinStation problem, and that $S = \{s_0, \dots, s_{k-1}\}$ is the solution returned by the MinStation algorithm. We prove now that $k = n$ or $k = n + 1$.

Since s_0 is on the convex hull of \mathcal{P} there is a shortest d -path between some z_i and z_{i+1} that contains s_0 . Hence, adding s_0 to $Z = \{z_0, \dots, z_{n-1}\}$ yields an optimal or quasi-optimal solution to the MinStation problem including s_0 . \square

Remark 1. The OPTSOL algorithm can be easily adapted to solve the MinStation problem for a polygonal line \mathcal{Q} with endpoints q_0 and q_k , contained in \mathcal{P} . This may be useful to patrol a section of the coastline instead of a complete island. If both endpoints of \mathcal{Q} are contained in an interval $[a, b]$ such that \overline{ab} is a d -bridge of \mathcal{P} , then we need at most two base stations depending on the length of \overline{ab} . Otherwise, if q_0 is not contained in such an interval $[a, b]$, we run OPTSOL clockwise starting from q_0 and stop when we reach a point s such that $q_k \in [q_0, s]$. In the remaining case we run OPTSOL counterclockwise starting from q_k . As the solution returned by OPTSOL for \mathcal{P} is optimal or quasi-optimal, the solution obtained for \mathcal{Q} is also optimal or quasi-optimal.

Remark 2. Although there are polygons such that the optimal solution contains no points in $CH(\mathcal{P})$ (a simple modification of Fig. 1 yields one such example), from a practical point of view, it is convenient to assume that at least one station v lies in $\mathcal{P} \cap CH(\mathcal{P})$, as otherwise any solution that includes v could have an arbitrarily large number of stations (imagine that it is located in a large pocket where bridges cannot be established).

4.1. Time complexity

We prove now that we can implement the OPTSOL algorithm to run in $O(n^2)$ time.

Given a point s_i we want to find a point $y_{i+1} = \max\{y : \delta_d(s_i, y) = 2\}$ with respect to \leq and a point $s_{i+1} \in \{w : \delta_d(s_i, w) = 1 \wedge \delta_d(w, y_{i+1}) = 1\}$. We refer to the problem of finding s_{i+1} and y_{i+1} as the 2-hop problem, see Fig. 6.

We will show that by applying a quadratic time pre-processing on \mathcal{P} the 2-hop problem can be solved in linear time for each s_i .

A point x of an edge e is a projection of a vertex p_i on e if $x \leq p_i$ and the line segment joining them is a d -bridge of \mathcal{P} perpendicular to e . See Fig. 6a.

In a similar way, we say that a point x of an edge e of \mathcal{P} is called a d -projection of an edge f on e if there is a point $y \in f$ such that the line segment joining them is a bridge of \mathcal{P} of length d perpendicular to e . See Fig. 6b.

Lemma 5. Given s_i, s_{i+1} is either a vertex of \mathcal{P} , the projection of a vertex on an edge, the d -projection of an edge or a point with $\delta(s_i, s_{i+1}) = d$.

Proof. Suppose that s_{i+1} is not a vertex of \mathcal{P} and $\delta(s_i, s_{i+1}) < d$. Let e be the edge of \mathcal{P} containing s_{i+1} , see Fig. 6. Note that $\delta(s_{i+1}, y_{i+1}) = d$ by the choice of y_{i+1} . If s_{i+1} is neither the projection of a vertex on e nor the d -projection of an edge on e , then it can be moved slightly along edge e and advance y_{i+1} . This contradicts the definition of y_{i+1} . \square

There might be $O(n)$ points at distance d from a previously placed station s_i . However, we only need to consider the maximum with respect to \preceq among them as a candidate for placing s_{i+1} , as we prove next.

Lemma 6. Let w, x, y, z be points in \mathcal{P} such that $w \preceq x \preceq y \preceq z$. Suppose that $\delta(w, x) = \delta(w, y) = d$, $\delta(w, z) > d$, and $\delta(x, z) = \ell$. Then, $\delta(y, z) \leq \ell$.

Proof. Let r be an intersection point of the shortest d -path $\pi_{w,y}$ from w to y and the shortest d -path $\pi_{x,z}$ from x to z . Note that r always exists by the choice of the four points on \mathcal{P} . Let $\delta(w, r)$ and $\delta(r, y)$ be the distance along $\pi_{w,y}$ between w and r , and between r and y , respectively. Let $\delta(x, r)$ and $\delta(r, z)$ be the distance along $\pi_{x,z}$ between x and r , and between r and z , respectively. Since $\delta(w, z) > d$, we have $\delta(r, y) < \delta(r, z)$. Now suppose that $\delta(r, y) > \delta(x, r)$. Then we have that $\delta(w, r) + \delta(r, x) < d$, which is a contradiction to our assumption that $\delta(w, x) = d$. Thus, $\delta(r, y) \leq \delta(x, r)$ and $\delta(y, z) \leq \ell$. \square

We claim that, even though there might be $O(n^2)$ projections of vertices and d -projections of edges, $O(n)$ candidate points are sufficient to compute s_{i+1} .

Let e and f be edges of \mathcal{P} . We say that $e \prec f$ if for any point x in the interior of e and any point y in the interior of f , $x \preceq y$.

Lemma 7. For each edge e of \mathcal{P} we need to store at most three points:

1. The minimum d -projection (with respect to \preceq) of an edge e' on e such that $e \prec e'$.
2. The endpoint not in e of the bridge generating the maximum d -projection (with respect to \preceq) on e of an edge e' such that $e' \prec e$. In this case the stored point lies on e' .
3. The minimum projection (with respect to \preceq) of a vertex on e .

Proof. *Case 1.* Let x and x' be d -projections of two distinct edges f and f' , respectively, on e such that $e \prec f$ and $e \prec f'$. Let \overline{xy} be the d -bridge perpendicular to e having x as an endpoint, i.e., $y \in f$ and \overline{xy} has length d . Define $\overline{x'y'}$ analogously. Because of the length of \overline{xy} (respectively, $\overline{x'y'}$), if we place a station at x (respectively, x') then we also need to place a station at y (respectively, y'). Suppose w.l.o.g. that $x \preceq x'$, see Fig. 7a. Since all the bridges defining d -projections of edges on e are parallel, this implies that $f' \prec f$ and $y' \preceq y$. Moreover, as the interval $[x, y]$ contains the interval $[x', y']$, placing a station at x guarantees that both intervals of \mathcal{P} are guarded. Hence, we maximize y_{i+1} with respect to \preceq by choosing the minimum d -projection of an edge on e as s_{i+1} .

Case 2. This case is analogous to the first one, see Fig. 7b.

Case 3. Let x and x' be the projections of two distinct vertices p_i and p_j , respectively, on an edge e . Let $\overline{xp_i}$ and $\overline{x'p_j}$ be their corresponding d -bridges. Suppose w.l.o.g. that $x \preceq x'$. This implies that $p_j \preceq p_i$ and that placing a station at x guarantees that both intervals $[x, p_i]$ and $[x', p_j]$ are guarded, see Fig. 8a. It remains to be proven that by placing a station at x we can advance further on \mathcal{P} with respect to \preceq . Let $w \in \mathcal{P}$ be a point such that $p_i \preceq w$, and let $p_{x,w}$ and $p_{x',w}$ be the shortest d -paths joining x to w and x' to w . Let r be the intersection point of $\overline{xp_i}$ and $p_{x',w}$. Notice that the points x, x' and r form a right triangle that is right-angled at x . Therefore, the length of $p_{x,w}$ is smaller than the length of $p_{x',w}$, which implies that we can maximize y_{i+1} by choosing the minimum projection of a vertex on e as s_{i+1} . \square

In order to compute the candidate points on \mathcal{P} , we first find, for each edge $e \in \mathcal{P}$, the subset containing each point $x \in \mathcal{P}$ for which there is a segment perpendicular to e joining x and e , and completely contained in $\text{Ext}(\mathcal{P})$. In such case we say that x is *orthogonally visible* from e .

We define a *lid* as an edge of the convex hull of \mathcal{P} that is not an edge of \mathcal{P} . Each lid $h = \overline{ab}$ defines a polygon \mathcal{P}_h , which is the union of h and the interval of \mathcal{P} determined by a and b which has no points in the convex hull of \mathcal{P} besides a and b . Note that any projection of a vertex or d -projection of an edge is defined by a segment whose endpoints are contained in the same \mathcal{P}_h , for otherwise the segment would intersect $\text{Int}(\mathcal{P})$. Therefore, we only need to compute the set of points orthogonally visible from each edge e contained in a \mathcal{P}_h ; moreover, we only need to look at the polygon \mathcal{P}_h containing e to find these points.

For the next lemma, we assume that we have computed the polygons defined by all the lids of \mathcal{P} , as well as the triangulation of each such polygon. This can be done in $O(n)$ time overall, see [25] and [6].




Fig. 7. (a) Case 1: we only need to store the point x on edge e . (b) Case 2: we only need to store the point x on edge f .

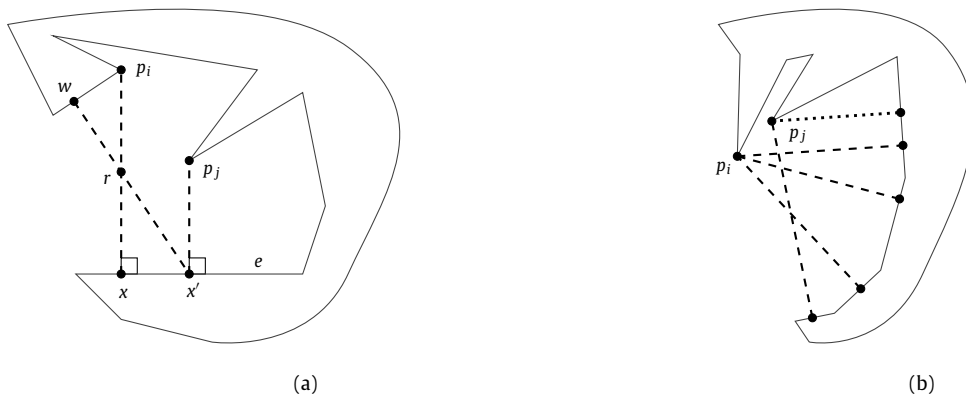


Fig. 8. (a) Case 3: The distance from x to w is smaller than the distance from x' to w . (b) We need to store all vertex projections except the one that is the endpoint of the dotted segment.

Lemma 8. We can find the set containing all the segments of \mathcal{P} orthogonally visible from any edge of \mathcal{P} in $O(n)$ time. Moreover, each such set has $O(n)$ size.

Proof. Let $h = \overline{ab}$ be a lid of \mathcal{P} and let $e = \overline{uv}$ be an edge of \mathcal{P}_h . We proceed as follows: Compute the set $VP(\mathcal{P}_h, e)$ of points of \mathcal{P}_h visible from a point in e . $VP(\mathcal{P}_h, e)$ can be computed in $O(n)$ time [13].

Suppose w.l.o.g. that $u < v$. Let R be the region contained between the lines perpendicular to e through u and v , and to the left of the line directed from u to v . It is easy to see that any point of \mathcal{P} orthogonally visible from e must lie in $\mathcal{R}_e = VP(\mathcal{P}_h, e) \cap R$, which can be computed in $O(n)$ time by intersecting $VP(\mathcal{P}_h, e)$ with both lines. We suppose w.l.o.g. that e is horizontal and that the interior of \mathcal{R}_e lies above e .

We say that a vertex $p \in \mathcal{R}_e$ is a *turn vertex* if the maximal vertical segment \overline{xy} through p and completely contained in \mathcal{R}_e separates \mathcal{R}_e into three subpolygons, see Fig. 9a. If two of these subpolygons lie to the right (left) of \overline{xy} , we say that p is a *right (left) turn vertex*. Let x be the top endpoint of \overline{xy} . The segment \overline{px} separates \mathcal{R}_e into two subpolygons, one of them containing e . Let $\mathcal{R}_e(p)$ denote the subpolygon generated by \overline{px} not containing e . It is easy to see that any point in \mathcal{R}_e not being orthogonally visible from e lies in the subpolygon $\mathcal{R}_e(p)$ for some turn vertex p , and that any point in $\mathcal{R}_e(p) \setminus \overline{px}$ is not orthogonally visible from e .

Note that the internal angles at both vertices of $e = \overline{uv}$ are convex in \mathcal{R}_e . Ghosh et al. [12] proved that for any vertex p in \mathcal{R}_e , the shortest path from u to p , denoted as $\rho_{u,p}$, makes a left turn at every vertex of the path, and $\rho_{v,p}$ makes a right turn at every vertex of the path. This also holds true for the points in the interior of any edge of \mathcal{R}_e .

Let p be a turn vertex of \mathcal{R}_e and let x be the top endpoint of the maximal vertical segment through p completely contained in \mathcal{R}_e . We claim that the vertical line through p , ℓ_p , does not intersect any point of $\mathcal{R}_e(p) \setminus \overline{px}$. Suppose otherwise that there is a point x' in $\mathcal{R}_e(p) \setminus \overline{px}$ contained in ℓ_p . Then, there exists a vertex q in $\mathcal{R}_e(p) \setminus \overline{px}$ such that $\rho_{v,x'}$ makes a left turn at q or $\rho_{u,x'}$ makes a right turn at q , which is a contradiction [12], see Fig. 9b. It follows that $\mathcal{R}_e(p) \cap \ell_p = \overline{px}$. This fact yields the following algorithm for removing $\mathcal{R}_e(p)$ from \mathcal{R}_e for each turn vertex p .

We deal with the right turn vertices by traversing the edges of \mathcal{R}_e clockwise from v to u . We set a variable *edgesVisible* to *true*. Let $f = \overline{qr}$, $q < r$, be the current edge in the traversal.




Fig. 9. (a) p is a left turn vertex, r is a right turn vertex, and q is not a turn vertex. (b) Neither x' nor any point in the shaded region is in R_e : the shortest path from v to x' makes a left turn at q .

- If *edgesVisible* is *true* we check if r is a right turn vertex. In the affirmative case, we set *edgesVisible* to *false* and store the vertical line through r , ℓ_r and the edge f .
- If *edgesVisible* is *false*, then we had previously stored the last visible edge $g = \overline{op}$, where p is a right turn vertex, and the vertical segment through p , ℓ_p . We check if $x = f \cap \ell_p$ is not empty. In such a case, we replace the interval $[p, x]$ of \mathcal{R}_e with the vertical segment \overline{px} , set *edgesVisible* to *true*, and discard g and ℓ_p .

We can remove the sub-polygons defined by the left turn vertices analogously by traversing \mathcal{R}_e counter-clockwise from u to v . As each edge of \mathcal{R}_e is visited at most twice, the removal of the sub-polygons defined by all the turn vertices takes $O(n)$ time. Let \mathcal{R}'_e be the polygon obtained by these traversals.

To obtain the subset of \mathcal{P} orthogonally visible from e , we only need to discard e , the segment contained in the lid of \mathcal{P}_h , and the vertical segments added in the previous process (at most one per turn vertex) from \mathcal{R}'_e .

Since \mathcal{P}_h has no holes, each edge of \mathcal{P}_h provides at most one segment to \mathcal{R}'_e . Therefore, the set of segments of \mathcal{P} orthogonally visible from any edge e has $O(n)$ size. \square

Lemma 9. For any edge e of \mathcal{P} we can find the projections described in Lemma 7 in $O(n)$ time.

Proof. Suppose that $e = \overline{uv}$, $u \leq v$. By Lemma 8, we can find the set W of all the segments of edges and vertices of \mathcal{P} orthogonally visible from e in $O(n)$ time; moreover, W has $O(n)$ size. Let W_B be the subset of elements of W smaller than u and let W_A be the subset of the elements of W greater than v with respect to \leq .

We find the d -projections corresponding to the first two cases of Lemma 7 as follows. Let ℓ be the line parallel to e , to the left of the line directed from u to v and at distance d from e . We first compute the intersection of ℓ with both W_B and W_A , which by the size of W can be obtained in $O(n)$ time. To obtain the point described in the first case of the proof of Lemma 7 we take the maximum point q with respect to \leq in $\ell \cap W_A$ and store the intersection point of e with the line through q perpendicular to e . To obtain the described in the second case of the proof of Lemma 7 we store the minimum point in $\ell \cap W_B$ with respect to \leq , if any.

We find the projection of the maximum vertex on e described in the third case of Lemma 7 as follows. For each vertex of \mathcal{P} in W_B we compute its distance with respect to e . We then store the maximum with respect to \leq of the vertices at distance less or equal than d from e . \square

Now we need to solve the following subproblem: given a point $x \in \mathcal{P}$, find the maximum w , $x \leq w$, such that $\delta(x, w) = d$. Guibas et al. [13] proved that, given the triangulation of a polygon \mathcal{R} and a point $p \in \mathcal{R}$, the euclidean shortest paths from p to all the vertices of \mathcal{R} can be found in linear time (see also [20]). The union of all the shortest paths from the source point p to the vertices of \mathcal{R} is a planar tree called the *shortest-path tree* of \mathcal{R} with respect to p .

Let \mathcal{R} be the polygon obtained by enclosing \mathcal{P} in a sufficiently large rectangle and connecting one of the sides of the rectangle to the starting point of the sequence, x_0 , with a thin corridor. The polygon \mathcal{R} can be obtained in $O(n)$ time, see [28]. Note that \mathcal{R} has $m \leq n + 8$ vertices and \mathcal{P} is contained in the exterior of \mathcal{R} . We assign to the points in \mathcal{R} that are also points in \mathcal{P} the same order as in \mathcal{P} .

Henceforth we assume that \mathcal{R} has been computed along with its triangulation, which as proven by Chazelle [6] can be found in $O(n)$ time.

Lemma 10. Given any point $x \in \mathcal{P}$, the point $w \in \mathcal{P}$ with $\delta(x, w) = d$ such that w is maximum with respect to \leq can be found in $O(n)$ time.

Proof. Let x be a point in \mathcal{P} and let x' be its corresponding point in \mathcal{R} . We compute the shortest path $\rho(x', y)$ from x' to every vertex $y \in \mathcal{R}$ such that y is also a vertex of \mathcal{P} and $x' < y$. Let T be the shortest-path tree obtained by the union of

these shortest paths. Let M be the set of vertices of T such that, for any $w \in M$, $\delta(x', w) \leq d$, and w shares an edge of \mathcal{R} with a vertex y such that $\delta(x', y) > d$. The set M can be found in $O(n)$ time by traversing T from its root x' .

Observe that any point of \mathcal{R} at distance d from x' is one of the following:

- An element of M .
- A point in an edge $e = \overline{uv}$, $u < v$, of \mathcal{R} such that $e \in E(T)$. In this case, $u \in M$ and $\delta(x', v) > d$.
- A point in an edge $e = \overline{uv}$, $u < v$, of \mathcal{R} such that $e \notin E(T)$. Notice that, in this case, $\delta(x', v) > d$. Moreover, there is exactly one $z \in M$ such that $(z, u) \in E(T)$.

Hence, in order to find all the points at distance exactly d from x' it is sufficient to check the edges having a neighbor of an element of M in T as an endpoint. Since each vertex is adjacent to at most one element of M , this can be done in $O(n)$ time. At the final step we need to find the maximum among all the points at distance d from x' , which can also be done in $O(n)$ time. Our result follows. \square

Theorem 11. *Let \mathcal{P} be a polygon with n vertices and let $s_0 \in \mathcal{P}$ be a point on the convex hull of \mathcal{P} . Then OPTSOL returns an optimal solution S to the MinStation Problem such that $s_0 \in S$ in $O(n^2 + \frac{L}{4}n)$ time, where L is the length of \mathcal{P} .*

Proof. By Lemma 5, given s_i , the point s_{i+1} is either a point on \mathcal{P} at distance exactly d from s_i , a vertex of \mathcal{P} , the projection of a vertex onto an edge, or the d -projection of an edge onto another edge.

By Lemma 6, we only need to consider the maximum point with respect to \leq at distance d from x_i , which can be found in $O(n)$ time as stated in Lemma 10.

There might be $O(n^2)$ projections of vertices and d -projections of edges. However, Lemma 7 states that in the set of candidates we need to store at most three projections for each edge of \mathcal{P} . Moreover, these projections can be found in $O(n)$ time for each edge.

The set of candidate points to compute all the elements of the set S has $O(n)$ size. For each candidate x , we compute the maximum point at distance d from x and associate this point to x , which by Lemma 10 takes $O(n)$ time per candidate.

It is easy to see that we only need to consider the candidates contained in the interval of \mathcal{P} from s_i to the maximum point with respect to \leq at distance d from s_i . From all these candidates, we choose as s_{i+1} the candidate which maximizes y_{i+1} , which can be done in $O(n)$ time. Since we might need to place $O(\frac{L}{d})$ stations, this step takes time $O(\frac{L}{d}n)$. Therefore, the set S can be found in $O(n^2 + \frac{L}{4}n)$ time. \square

5. Discretization

In this section, we present a discretization algorithm that is easy to implement for the MinStation problem, and then show how it can be utilized to obtain a solution to the MinDistance problem which is close to optimal. This algorithm avoids computing projections, drone distances (geodesic paths) and orthogonal visibility, which makes it very practical. The idea is to construct a graph and apply a slight modification of Dijkstra algorithm.

Fix $0 < \epsilon \leq d$ and let $X = \{s_0 = x_0 \leq \dots \leq x_{r-1}\} \subseteq \mathcal{P}$ be a set of points so that s_0 lies on $\mathcal{P} \cap CH(\mathcal{P})$ and the distance between x_i and x_{i+1} along \mathcal{P} is at most ϵ , addition taken mod r . For technical reasons that will become apparent later, we also ask that the vertices of \mathcal{P} are contained in X . Consider the graph $G_d(X)$ such that $V(G_d(X)) = X$ in which two elements $x_i, x_j \in X$ are adjacent if the length of the geodesic path π_{x_i, x_j} in $\mathcal{P} \cup Ext(\mathcal{P})$ connecting them is at most d (as we will show soon, computing $G_d(X)$ does not require the shortest-path trees mentioned in Lemma 10). We then solve the problem of finding a shortest cycle in $G_d(X)$ from x_0 to itself going around \mathcal{P} . The set of vertices of that cycle, including x_0 , is a valid solution to our problem, but not necessarily an optimal one.

Note that the problem of finding a shortest cycle from x_0 to itself can be reduced to that of finding a shortest path from x_0 to a copy $x'_0 = x_r$ of x_0 . To this end, we insert x'_0 in $V(G_d(X))$ in such a way that, if the length of the interval $[x_i, x_0]$ is at most d , then x_i is adjacent to x'_0 instead of x_0 .

Now we show in detail how the algorithm works, including how to compute $G_d(X)$.

It is possible to check whether a directed edge (x_i, x_j) belongs to E_2 in $O(n)$ time. This leads to a total time complexity of $O((\frac{L}{\epsilon})^3 + (\frac{L}{\epsilon})^2n)$ for APPSOL, where L denotes the total length of \mathcal{P} .

This algorithm, while simpler to implement than OPTSOL, does not directly yield an approximation to the MinStation problem (this is discussed in more detail in the next section). We now show how we can improve on this by applying this algorithm more than once: two applications of the MinStation APPSOL algorithm can be used to certify the sharpness of a single application of this result, and a logarithmic number of applications can be used to give an additive approximation for MinDistance.

Denote by $\alpha(\mathcal{P}, s_0, d, \epsilon)$ be the number of base stations returned by the APPSOL algorithm for given \mathcal{P} , $s_0 \in \mathcal{P} \cap CH(\mathcal{P})$, flight range d , and $\epsilon > 0$. Let k be the minimum number of base stations among all solutions that have s_0 as one of their base stations. The key is the following result:

Theorem 12. $\alpha(\mathcal{P}, s_0, d + \epsilon, \epsilon) \leq k \leq \alpha(\mathcal{P}, s_0, d, \epsilon)$. In particular, if $\alpha(\mathcal{P}, s_0, d + \epsilon, \epsilon) = \alpha(\mathcal{P}, s_0, d, \epsilon)$, the solution is best possible among those containing s_0 .

Algorithm 2: APPSOL.

Input: Polygon \mathcal{P} , $s_0 \in \mathcal{P} \cap CH(\mathcal{P})$, $d > 0$, $\epsilon > 0$
Output: List of stations in a d -hull of \mathcal{P}

- 1 **if** $s_0 = x_0$ is not a vertex of \mathcal{P} **then**
- 2 \lfloor Split the edge containing s_0 in two in such a way that s_0 becomes a vertex
- 3 Let V be the set of vertices of \mathcal{P} and set $X = V$
- 4 Add a copy x'_0 of x_0 to X
- 5 **for each edge** of \mathcal{P} of length $\ell > \epsilon$ **do**
- 6 \lfloor Add $\lceil \frac{\ell}{\epsilon} \rceil$ points to X dividing the edge into segments of length $\leq \epsilon$
- 7 Let $X = \{x_0, x_1, \dots, x_{m-1}, x_m = x'_0\}$ be the set of points in clockwise order around \mathcal{P}
- 8 Construct a weighted directed graph $H_d(X) = (X, E)$ with $E = E_1 \cup E_2$ defined as follows:
 - (a) $(x_i, x_j) \in E_1$ if $j = i + 1$ and x_i, x_{i+1} are on the same edge of \mathcal{P}
 - (b) $(x_i, x_j) \in E_2$ if $i < j$, and the open segment from x_i to x_j has length $\leq d$ and is contained in $Ext(\mathcal{P})$
 - (c) The weight of each edge $(x_i, x_j) \in E$ is the Euclidean distance between x_i and x_j
- 9 **for each** $x_i \in X$ **do**
- 10 \lfloor Use Dijkstra's algorithm to compute X_i , the set of vertices of X that can be reached from x_i by a directed path of total weight $\leq d$
- 11 Construct a graph with vertex set X where x_i is adjacent to x_j iff $x_j \in X_i$ or $x_i \in X_j$. Since the vertices of \mathcal{P} belong to X , one can easily check that this graph is actually $G_d(X)$.
- 12 Using BFS (or Dijkstra's algorithm with weights 1), compute a shortest path from x_0 to $x'_0 = x_m$ of minimum length in $G_d(X)$

return the set of vertices of $G_d(X)$

Proof. Clearly, $\alpha(\mathcal{P}, s_0, d, \epsilon) \geq k$. It suffices to show that $\alpha(\mathcal{P}, s_0, d + \epsilon, \epsilon) \leq k$. Consider an optimal set of k base stations S^* . Let S be a set of base stations obtained by selecting the nearest point in X for each point in S^* , then the geodesic distance between consecutive base stations in S is at most $d + \epsilon$. Therefore $\alpha(\mathcal{P}, s_0, d + \epsilon, \epsilon) \leq k$. \square

Since s_0 lies on the boundary of the convex hull of \mathcal{P} , every solution to MinStation must contain a station on a point $z \in \mathcal{P}$ such that $\delta(x_0, z) \leq d$. This can easily be seen to imply that Theorem 13 can be adapted to work for general solutions (and not only those that contain x_0) by modifying the algorithm so that it searches for the shortest path in $G_d(X)$ from x_i to itself for all x_i with $\delta(x_0, x_i) \leq d + \frac{\epsilon}{2}$, and then returns the shortest one among all of these. This slight variant of APPSOL will be called APPSOL2.

This has immediate implications for MinDistance; if the least number of stations in a solution in $G_{d+\epsilon}(X)$ is at most k , then the optimal solution to the MinDistance Problem (find the smallest flight range such that k stations are sufficient) lies between d and $d + \epsilon$. Thus by using binary search on d , the optimal flight range can be approximated up to an additive constant.

Theorem 13. Given a positive integer k and an $\epsilon > 0$, it is possible to find a solution to the MinDistance problem using k base stations such that the flight capacity of the drones is at most ϵ larger than the optimal one. This is achieved by running $O(\log |X|) = O(\log(\frac{L}{\epsilon} + n))$ iterations of APPSOL2 to perform a binary search on the set of all distinct drone (geodesic) distances between pairs of points of X .

Corollary 14. Given a positive integer k and an $\epsilon > 0$,

- An additive ϵ -approximation for the MinDistance problem with one fixed base station can be computed in $O((n^2 + Ln) \log(\frac{L}{\epsilon} + n))$ time.
- A quasi-optimal additive ϵ -approximation for the MinDistance problem (i.e. with k or $k + 1$ base stations) can be computed in $O((n^2 + Ln) \log(\frac{L}{\epsilon} + n))$ time.

6. Experiments

We implemented algorithm APPSOL. The program is written in Java and is available at [14]. We run experiments on data from Salamis Island using data provided by Harvard WorldMap.⁹ There are 596 vertices in the polygon representing the island. The vertices are given by latitude and longitude and we converted them to (x, y) coordinates in meters. According to our data, the perimeter of the island is 113639.9 meters.

Our goal was to find the optimal number of base stations using approximation algorithm APPSOL and the sufficient condition provided by Theorem 12. For different values of d and a fixed base station ($s_0 = 0$) on the island, we apply the following approach. We start with epsilon equal to the drone capacity d . We divide ϵ by 1.2 each time it does not satisfy the sufficient condition of Theorem 12. Interestingly, the optimal number of base stations was found in all experiments. The results are shown in Table 1. The program was executed on a Linux server with 32 core CPUs and 64 GB RAM. Observe that k is monotone with respect to d but ϵ is not. Two solutions for $d = 2000$ and $d = 2400$ are shown in Fig. 10.

⁹ <http://worldmap.harvard.edu>.

Table 1

Columns: distance d in meters, k is the optimal number of base stations, ϵ of the last iteration is in meters, T_1 is the time (in seconds) of the last iteration, and T_2 is the total time of AppSol.

d	k	ϵ	T_1	T_2
1000	85	10.48	55.502	378.434
1200	69	12.58	40.774	360.072
1250	67	2.54	1034.178	7696.363
1300	63	1.06	6912.250	49544.916
1400	54	8.49	90.388	792.151
1500	51	15.72	36.228	317.915
1700	44	63.85	3.861	38.557
1750	43	12.74	53.491	508.694
1800	42	2.12	2443.456	17636.010
1900	38	16.60	44.015	363.716
2000	36	52.17	7.316	73.641
2100	34	1.43	6772.667	49193.534
2400	26	90.15	2.676	39.461
2500	25	93.90	2.950	36.413
3000	20	2.45	3688.721	26533.493
3200	19	120.20	2.753	44.280




Fig. 10. Salamis Island in the Saronic Gulf. (a) 36 base stations for $d = 2000$. (b) 26 base stations for $d = 2400$. The number of base stations in (a) and (b) is optimal among those containing s_0 by Theorem 12.

7. Conclusions

In this paper we consider the problem of finding the minimum number of refueling stations along the boundary of an island, modeled as a polygon \mathcal{P} with perimeter L , in such a way that a drone with flight range d can follow a polygonal path enclosing P . We describe an $O(n^2 + \frac{L}{d})$ -time algorithm that attains an optimal solution under the restriction that a base station is a point in the intersection of the boundary of \mathcal{P} and its convex hull. Moreover, if we remove this restriction, our algorithm returns a solution with at most one additional base station with respect to a globally optimal solution.

The setting of the problem allowed us to suppose that the drones fly at constant height, and therefore the assumption that any drone is always able to fly between base stations at distance at most d is not unreasonable. However, some applications may require to consider the elevation differences or the presence of obstacles between base stations, which means that re-computing the maximum flight distance each time a base station is placed might be necessary. If this value can be obtained in linear time per base station, then our algorithm could be adapted for these settings while keeping the original time complexity. To accomplish this we only have change the value of d in the steps described in Lemma 9 and Lemma 10, as these steps are done for each base station.

It remains as an open problem to determine if the MinStation problem without the restriction that one base station has to lie on the convex hull can be optimally solved in polynomial time. This is relevant, since there exist examples in which an optimal solution contains no base station on the convex hull of the island.

We also presented an algorithm to obtain an additive approximation to the problem of minimizing the fuel capacity required for the drones to patrol an island when we are allowed to place at most k base stations around its boundary. The main tool in this solution is a discretization of the original MinStation problem. This discretized approach also yields an easier to implement algorithm to approximate the MinStation problem, albeit without any theoretical guarantees on the

quality of the solution. It is also an open problem to determine if an exact solution to the MinDistance problem can be obtained in polynomial time.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] N. Ahmed, S.S. Kanhere, S. Jha, The holes problem in wireless sensor networks: a survey, *ACM SIGMOBILE Mob. Comput. Commun. Rev.* 9 (2) (2005) 4–18.
- [2] B. Alzahrani, O.S. Oubbati, A. Barnawi, M. Atiquzzaman, D. Alghazzawi, UAV assistance paradigm: state-of-the-art in applications and challenges, *J. Netw. Comput. Appl.* 166 (2020) 102706.
- [3] S. Bespamyatnikh, An $o(\log n)$ algorithm for the zoo-keeper's problem, *Comput. Geom.* 24 (2) (2003) 63–74.
- [4] B. Bhattacharya, M. Burmester, Y. Hu, E. Kranakis, Q. Shi, A. Wiese, Optimal movement of mobile sensors for barrier coverage of a planar region, *Theor. Comput. Sci.* 410 (52) (2009) 5515–5528.
- [5] P. Bose, P. Morin, I. Stojmenović, J. Urrutia, Routing with guaranteed delivery in ad hoc wireless networks, *Wirel. Netw.* 7 (6) (2001) 609–616.
- [6] B. Chazelle, Triangulating a simple polygon in linear time, *Discrete Comput. Geom.* 6 (3) (1991) 485–524.
- [7] W.-P. Chin, S. Ntafos, Optimum watchman routes, in: *Proceedings of the Second Annual Symposium on Computational Geometry*, 1986, pp. 24–33.
- [8] C.T. Cicek, H. Gultekin, B. Tavli, The location-allocation problem of drone base stations, *Comput. Oper. Res.* 111 (2019) 155–176.
- [9] J. Czyzowicz, P. Egyed, H. Everett, D. Rappaport, T. Shermer, D. Souvaine, G. Toussaint, J. Urrutia, The aquarium Keeper's problem, in: *Proceedings of the Second Annual ACM-SIAM Symposium on Discrete Algorithms*, 1991, pp. 459–464.
- [10] J.M. Díaz-Báñez, J.A. Mesa, A. Schöbel, Continuous location of dimensional structures, *Eur. J. Oper. Res.* 152 (1) (2004) 22–44.
- [11] S. Ganeriwal, A. Kansal, M.B. Srivastava, Self aware actuation for fault repair in sensor networks, in: *IEEE International Conference on Robotics and Automation*, 2004, in: *Proceedings. ICRA'04. 2004*, vol. 5, IEEE, 2004, pp. 5244–5249.
- [12] S.K. Ghosh, A. Maheshwari, S.P. Pal, S. Saluja, C.V. Madhavan, Characterizing and recognizing weak visibility polygons, *Comput. Geom.* 3 (4) (1993) 213–233.
- [13] L. Guibas, J. Hershberger, D. Leven, M. Sharir, R.E. Tarjan, Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons, *Algorithmica* 2 (1) (1987) 209–233.
- [14] M. Haghpanah, Border protection project. [GitHub repository](https://github.com/mhaghpanah/RocioBorderProtection), [GitHub/mhaghpanah/RocioBorderProtection](https://github.com/mhaghpanah/RocioBorderProtection), 2022.
- [15] I. Hong, M. Kuby, A.T. Murray, A range-restricted recharging station coverage model for drone delivery service planning, *Transp. Res., Part C, Emerg. Technol.* 90 (2018) 198–212.
- [16] H. Huang, A.V. Savkin, A method of optimized deployment of charging stations for drone delivery, *IEEE Trans. Transp. Electr.* 6 (2) (2020) 510–518.
- [17] Y. Karaca, M. Cicek, O. Tatli, A. Sahin, S. Pasli, M.F. Beser, S. Turedi, The potential use of unmanned aircraft systems (drones) in mountain search and rescue operations, *Am. J. Emerg. Med.* 36 (4) (2018) 583–588.
- [18] S. Kumar, T.H. Lai, A. Arora, Barrier coverage with wireless sensors, *Wirel. Netw.* 6 (13) (2007) 817–834.
- [19] G. Laporte, S. Nickel, F. Saldanha-da Gama, Introduction to location science, in: *Location Science*, Springer, 2019, pp. 1–21.
- [20] D.-T. Lee, F.P. Preparata, Euclidean shortest paths in the presence of rectilinear barriers, *Networks* 14 (3) (1984) 393–410.
- [21] M. Li, L. Zhen, S. Wang, W. Lv, X. Qu, Unmanned aerial vehicle scheduling problem for traffic monitoring, *Comput. Ind. Eng.* 122 (2018) 15–23.
- [22] Y. Liu, Z. Liu, J. Shi, G. Wu, C. Chen, Optimization of base location and patrol routes for unmanned aerial vehicles in border intelligence, surveillance, and reconnaissance, *J. Adv. Transp.* (2019) 2019.
- [23] T. Long, M. Ozger, O. Cetinkaya, O.B. Akan, Energy neutral Internet of drones, *IEEE Commun. Mag.* 56 (1) (2018) 22–28.
- [24] S.G. Manyam, S. Rasmussen, D.W. Casbeer, K. Kalyanam, S. Manickam, Multi-UAV routing for persistent intelligence surveillance & reconnaissance missions, in: *2017 International Conference on Unmanned Aircraft Systems (ICUAS)*, IEEE, 2017, pp. 573–580.
- [25] A.A. Melkman, On-line construction of the convex hull of a simple polyline, *Inf. Process. Lett.* 25 (1) (1987) 11–12.
- [26] A. Merwaday, A. Tuncer, A. Kumbhar, I. Guvenc, Improved throughput coverage in natural disasters: unmanned aerial base stations for public-safety communications, *IEEE Veh. Technol. Mag.* 11 (4) (2016) 53–60.
- [27] G. Nagy, S. Salhi Location-routing, Issues, models and methods, *Eur. J. Oper. Res.* 177 (2) (2007) 649–672.
- [28] S. Ntafos, L. Gewali, External watchman routes, *Vis. Comput.* 10 (8) (1994) 474–483.
- [29] J. O'Rourke, *Art Gallery Theorems and Algorithms*, vol. 57, Oxford New York, NY, USA, 1987.
- [30] R.G. Ribeiro, L.P. Cota, T.A. Euzébio, J.A. Ramírez, F.G. Guimarães, Unmanned-aerial-vehicle routing problem with mobile charging stations for assisting search and rescue missions in postdisaster scenarios, *IEEE Trans. Syst. Man Cybern. Syst.* (2021).
- [31] İ. Sarıççek, Y. Akkuş, Unmanned aerial vehicle hub-location and routing for monitoring geographic borders, *Appl. Math. Model.* 39 (14) (2015) 3939–3953.
- [32] T.C. Shermer, Recent results in art galleries (geometry), *Proc. IEEE* 80 (9) (1992) 1384–1399.
- [33] X. Tan, Fast computation of shortest watchman routes in simple polygons, *Inf. Process. Lett.* 77 (1) (2001) 27–33.
- [34] X. Tan, T. Hirata, Finding shortest safari routes in simple polygons, *Inf. Process. Lett.* 87 (4) (2003) 179–186.
- [35] X. Tan, T. Hirata, Y. Inagaki, Corrigendum to an incremental algorithm for constructing shortest watchman routes, *Int. J. Comput. Geom. Appl.* 9 (03) (1999) 319–323.
- [36] J. Urrutia, Art gallery and illumination problems, in: *Handbook of Computational Geometry*, Elsevier, 2000, pp. 973–1027.
- [37] G. Wang, G. Cao, T.F. La Porta, Movement-assisted sensor deployment, *IEEE Trans. Mob. Comput.* 5 (6) (2006) 640–652.
- [38] C. Wei-Pang, S. Ntafos, The zookeeper route problem, *Inf. Sci.* 63 (3) (1992) 245–259.
- [39] E. Yakıcı, Solving location and routing problem for UAVs, *Comput. Ind. Eng.* 102 (2016) 294–301.