Discrete Surveillance Tours in Polygonal Domains

Elmar Langetepe* Bengt J. Nilsson† Eli Packer‡

Abstract

The watchman route of a polygon is a closed tour that sees all points of the polygon. Computing the shortest such tour is a well-studied problem. Another reasonable optimization criterion is to require that the tour minimizes the hiding time of the points in the polygon, i.e., the maximum time during which any points is not seen by the agent following the tour at unit speed. We call such tours *surveillance routes*.

We show a linear time 3/2-approximation algorithm for the optimum surveillance tour problem in rectilinear polygons using the L_1-metric. We also present an $O(\text{polylog}\, w_{\max})$-approximation algorithm for the optimum weighted discrete surveillance route in a simple polygon with weight values in the range $[1, w_{\max}]$. Our algorithm can have superpolynomial complexity since the tour may have to see points of high weight many times.

1 Introduction

Visibility coverage of polygons with guards (mainly known as *Art Gallery* problems) have been central geometric problems for many years. Usually guards are defined as static points that see in any direction for any distance and visibility is defined by the clearance of straight lines between two features (in other words, two features see each other if the segment that connects them does not intersect (the interior of) any other feature of the input). Coverage is achieved if any point inside the polygon is visible by at least one guard.

Several art gallery variants have been proposed for different kind of settings. These include different classes of polygons, such as rectilinear and monotone polygons, and different types of guards, such as edge and segment guards; see [7, 8, 9, 11].

Allowing a guard to move inside the polygons defines a related problem but yet with very different properties. Here, a set of mobile guards walk on closed cycles (also called *tours* or *routes*) so that any point inside the polygon is seen by at least one guard during its walk along the tour. The number of guards is a parameter of the problem and the measure criteria relates to the length of the tours (e.g., minimize the longest tour). Several solutions have been proposed for the case of a single mobile guard, a *shortest watchman route* in a simple polygon. The currently fastest one combines algorithms by Tan [10] and Dror *et al.* [3], to achieve asymptotic running time $O(n^4 \log n)$.

We want to guard a given simple polygon $\mathbf{P}$, but rather than finding a shortest tour that covers the points of $\mathbf{P}$, we are interested in a tour that minimizes the maximum duration in which any of the points in $\mathbf{P}$ are not guarded. We call such a tour an *optimum surveillance route* for the polygon, abbreviated *OSR*. Kamphans and Langetepe [5] study a similar concept (*inspection paths*) but their optimization measure is the sum of the durations where features are not covered rather than the maximum duration.

We also consider a discrete version of the minimum surveillance tour problem where a given finite subset $\mathcal{S}$ of points in the polygon is to be guarded. We further generalize this version of the problem by associating weights to the points of $\mathcal{S}$.

We show a linear time 3/2-approximation algorithm for the optimum surveillance tour problem in rectilinear polygons using the L_1-metric. We also present a $O(\text{polylog}\, w_{\max})$-approximation algorithm for the optimum weighted discrete surveillance route in a simple polygon with weight values in the range $[1, w_{\max}]$.

2 Preliminaries

Let $\mathbf{V}(p)$ denote the *visibility polygon* of a point $p \in \mathbf{P}$, i.e., the set of all points q in $\mathbf{P}$ such that the segment pq fully lies inside $\mathbf{P}$. Obviously, the visibility polygon $\mathbf{V}(p)$ is a simple polygon itself. A watchman route is a closed tour within the polygon that sees all points of the polygon. Hence, a tour T is a watchman route if $\forall p \in \mathbf{P}; \mathbf{V}(p) \cap T \neq \emptyset$.

A reasonable extension of the concept of a watchman route is to require that the tour minimizes the hiding time of the points in the polygon, i.e., the maximum time during which any point in the polygon is not seen by an agent following the tour at unit speed. To formally define this, we introduce the concept of *hidden pieces* of a tour T.

Definition 1 Given a tour T and a point p, the *hidden pieces*, $\mathcal{H}_T(p)$, of T with respect to p is the set of

*Institute of Computer Science I, University of Bonn, 53117 Bonn, Germany. elmar.langetepe@cs.uni-bonn.de

†Department of Computer Science, Malmö University, SE-205 06 Malmö, Sweden. bengt.nilsson.TS@mah.se

‡Intel Corporation, Givataim, Israel. eli.packer@intel.com

maximal paths $\mathcal{H}_T(p) \stackrel{def}{=} \{T \setminus \mathbf{V}(p)\}$.

The visibility polygon $\mathbf{V}(p)$ of p subdivides the hidden pieces of T into a number of subpaths $X_1, X_2, \ldots, X_m$ that do not have any points seen from p. Hence, $\mathcal{H}_T(p) = \{X_1, X_2, \ldots, X_m\}$.

Definition 2 Given a tour T and a point p in $\mathbf{P}$, the *hiding cost* (Mitchell [4] calls it the *dark cost*), $hc_T(p)$, of T with respect to p is the length of the longest path X in $\mathcal{H}_T(p)$ if p is visible from T, i.e.,

$$hc_T(p) \stackrel{def}{=} \begin{cases} \infty, & \text{if } \mathbf{V}(p) \cap T = \emptyset, \\ \max_{X \in \mathcal{H}_T(p)}\{\|X\|\}, & \text{if } \mathbf{V}(p) \cap T \neq \emptyset, \end{cases}$$

where $\|X\|$ denotes the length of X in a given metric.

Given the definition of the hiding cost, we can define the surveillance cost or delay of a tour.

Definition 3 Given a tour T, the *surveillance cost* or *delay*, $d(T)$, of T is given by

$$d(T) \stackrel{def}{=} \max_{p \in \mathbf{P}}\{hc_T(p)\}. \tag{1}$$

We say that the tour T is a surveillance route for $\mathbf{P}$ if $d(T)$ is finite.

With this definition, it is clear that any surveillance route is also a watchman route, since all points of the polygon must be seen by the route for it to have finite surveillance cost.

Given a finite set of points $\mathcal{S}$ in $\mathbf{P}$ to be guarded, we define a discrete version of the surveillance cost or delay of a tour.

Definition 4 Given a tour T, the *discrete surveillance cost* or *discrete delay*, $d(T)$, of T with respect to a finite point set $\mathcal{S}$ to be guarded is given by

$$d_{\mathcal{S}}(T) \stackrel{def}{=} \max_{p \in \mathcal{S}}\{hc_T(p)\}. \tag{2}$$

We say that the tour T is a discrete surveillance route for $\mathcal{S}$ in $\mathbf{P}$ if $d_{\mathcal{S}}(T)$ is finite.

We make use of classical notation; see for example [2]; for the following definitions. To every reflex vertex in $\mathbf{P}$ we can associate two *extensions*, i.e., the two maximal line segments in $\mathbf{P}$ through the vertex and collinear to the two edges adjacent to the vertex; see Figure 1(a). We associate a direction to an extension e collinear to an edge e_v by giving e the same direction as e_v has when $\mathbf{P}$ is traversed in counterclockwise order. This allows us to refer to the regions to the left and right of an extension, meaning those point reached by a left turn or a right turn respectively from the directed segment e. Let $\mathbf{L}(e)$ denote the part of $\mathbf{P}$ to the left of e and $\mathbf{R}(e)$ the part to the right of e.

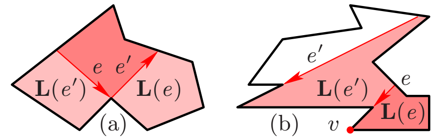

Figure 1: Illustrating definitions. (a) the two extensions issuing from a reflex vertex. (b) e dominates e', e is essential and v is an essential vertex.

We say that e is a *visibility extension* with respect to a surveillance route T, if T has some point in $\mathbf{R}(e)$.

The visibility extensions capture visibility information in the sense that a surveillance route must have points to the left of each of them.

We say that an extension e *dominates* another extension e', if $\mathbf{L}(e)$ is properly contained in $\mathbf{L}(e')$.

Definition 5 A visibility extension e is *essential*, if e is not dominated by any other visibility extension.

An essential extension e is collinear to an edge with one reflex and one convex vertex, since if both vertices are reflex, then there is another essential extension (issuing from the other reflex vertex) that dominates e, giving us a contradiction.

Definition 6 Let v be the convex vertex of the edge collinear to an essential extension e. We call the convex vertex v an *essential vertex*; see Figure 1(b).

The essential vertices play an important role for surveillance routes as we show in the next lemma.

For a polygon $\mathbf{P}$, we let OSR denote an *optimum surveillance route*, i.e., a tour T for which $d(T)$ is minimal.

Lemma 7 *If $\mathbf{P}$ is such such that $d(OSR) > 0$, then the delay of OSR is attained at some essential vertex of $\mathbf{P}$, i.e., there is an essential vertex v such that*

$$d(OSR) = hc_{OSR}(v). \tag{3}$$

Proof. Let p be a point in $\mathbf{P}$ such that $d(OSR) = hc_{OSR}(p) > 0$. Since $d(OSR) > 0$, the point p exists. Let X be a path in $\mathcal{H}_{OSR}(p)$ having the length of $d(OSR)$. The path X starts and finishes at an edge e of $\mathbf{V}(p)$ having a reflex vertex r of $\mathbf{P}$ as one endpoint. The segment e and the point p are collinear; see Figure 2. Thus, the segment e subdivides $\mathbf{P}$ into two parts, $\mathbf{P}_X$, containing the path X and $\bar{\mathbf{P}}_X$, not containing the path X. The part $\bar{\mathbf{P}}_X$ contains the point p and has e as a boundary edge, the point r is a reflex vertex of both $\mathbf{P}$ and $\bar{\mathbf{P}}_X$.

To prove that there is an *essential vertex* with hiding cost at least as high as that of p, follow the boundary of $\bar{\mathbf{P}}_X$ from r away from e until the first convex vertex u is reached and let u' be the last reflex vertex as we move along the boundary from r. We note that we could have $u' = r$. Let e' be the extension collinear to the edge

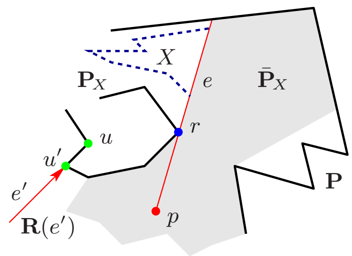

Figure 2: Illustrating the proof of Lemma 7.

$[u, u']$. Since the sequence of vertices along $\bar{\mathbf{P}}_X$ from r to u' is reflex, the extension e' is completely interior to $\bar{\mathbf{P}}_X$ and it is also a visibility extension since *OSR* has points in $\mathbf{R}(e')$, e.g., the path X in $\mathbf{P}_X$. Therefore, either e' is an essential extension or it is dominated by an essential extension. Let v denote the essential vertex for this essential extension, independently of whether the essential extension is e' or some other dominating extension. By construction, since e' is contained in $\bar{\mathbf{P}}_X$, the visibility polygon $\mathbf{V}(v)$ does not see any point in in $\mathbf{P}_X$, and hence $hc_{OSR}(v) \geq \|X\| = hc_{OSR}(p)$, proving the lemma. □

Lemma 7 shows that the optimum surveillance route in a simple polygon is the optimum discrete surveillance route of the essential vertices of the polygon, i.e., $d(OSR) = d_{\mathcal{V}}(OSR)$, where $\mathcal{V}$ is the set of essential vertices of the polygon.

3 $\mathbf{L_1}$-Surveillance Routes in Rectilinear Polygons

In [6], the authors show that the shortest watchman route and the optimum surveillance route are not necessarily the same and that the shortest watchman route is a 2-approximation to the optimum surveillance route in a simple polygon. It is still an open question whether the optimum surveillance route in a simple polygon can be computed in polynomial time, assuming P $\neq$ NP.

Even considering rectilinear polygons in the L_1-metric, a *SWR* is not necessarily the optimum surveillance tour; see Figure 3(a) and (b). The rectilinear polygon has five essential extensions, dotted lines, four of which have unit length and the fifth (the top middle one) is arbitrarily short. The length of a *SWR* is just over 8 which is also the surveillance cost. However, by revisiting the short extension we obtain a slightly longer tour with surveillance cost of just over 7.

We define a particular L_1-shortest watchman route.

Definition 8 In a rectilinear polygon, we call an L_1-shortest watchman route that has maximal interior area a *maximum shortest watchman route* and denote it by *MSWR*.

A *MSWR* has the special property that between any two consecutive essential extensions e and e', the *MSWR*

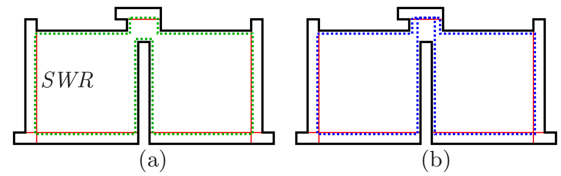

Figure 3: Counterexample showing that an L_1-optimal *SWR* is not an L_1-optimal *OSR* in rectilinear polygons.

follows a rectilinear shortest path between e and e'. A *MSWR* can be computed in linear time by a straightforward modification of the algorithm of Chin and Ntafos [1].

Theorem 9 *A MSWR is a 3/2-approximation for the L_1-optimal OSR in a rectilinear polygon.*

Proof. According to Lemma 7, there is an essential vertex v for which the hiding cost attains the delay of *OSR*. Let X be the path in $\mathcal{H}_{OSR}(v)$ with $\|X\| = d(OSR)$. We claim that $\|X\| \geq 2\|MSWR\|/3$ thus giving us that

$$d(MSWR) \leq \|MSWR\| \leq \frac{3}{2}\|X\| = \frac{3}{2}d(OSR). \quad (4)$$

To prove that $\|X\| \geq 2\|MSWR\|/3$, assume for a contradiction that $\|X\| < 2\|MSWR\|/3$. Let e be the essential extension of v and let p and q be the two endpoints of X on e. Since $[p, q] \cup X$ is a watchman route we have that $\|[p, q]\| + \|X\| \geq \|MSWR\|$ and therefore $\|[p, q]\| > \|MSWR\|/3$. Without loss of generality, we can assume that e is vertical. We construct the two maximal horizontal line segments interior to the polygon that go through the points p and q. The two segments subdivide the polygon into three pieces, $\mathbf{P}_{\mathrm{T}}$ the top piece, $\mathbf{P}_{\mathrm{M}}$ the middle piece, and $\mathbf{P}_{\mathrm{B}}$ the bottom piece; see Figure 4.

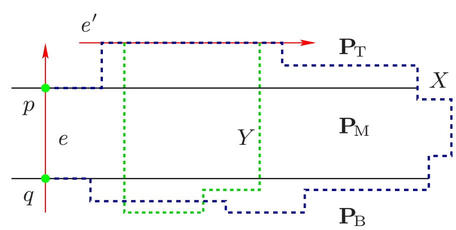

Figure 4: Illustrating the proof of Theorem 9.

Both $\mathbf{P}_{\mathrm{T}}$ and $\mathbf{P}_{\mathrm{B}}$ must contain essential extensions since otherwise, the path X is not part of the optimum surveillance route, giving us an immediate contradiction.

Therefore, let e' be an essential extension in $\mathbf{P}_{\mathrm{T}}$ with v' as the essential vertex and consider the set $\mathcal{H}_{OSR}(v')$. Some path Y in this set must visit essential extensions in $\mathbf{P}_{\mathrm{B}}$ and must therefore have length at least $2\|[p, q]\| > 2\|MSWR\|/3$; see Figure 4. Hence,

$$d(OSR) \geq hc_{OSR}(v') \geq \frac{2}{3}\|MSWR\| > \|X\| = d(OSR), \quad (5)$$

and we have a contradiction. □

Remark: Also in the case of the L_1-metric in rectilinear polygons, it is an open question whether the optimum surveillance route can be computed in polynomial time, assuming P $\neq$ NP.

4 Weighted Discrete Surveillance Routes

In this section, we consider the weighted discrete surveillance route problem in a simple polygon and define it as follows. Let $\mathbf{P}$ be a simple polygon with n edges and let $\mathcal{S}$ be a finite set of points inside $\mathbf{P}$. To each point $p \in \mathcal{S}$ is associated a weight $w(p)$. The idea is that points with higher weights have higher priority and need to be guarded more often than ones with lower weights. Given some tour T, we define the *weighted discrete delay* as

$$d^w_{\mathcal{S}}(T) \stackrel{def}{=} \max_{p \in \mathcal{S}} \{w(p) \cdot hc_T(p)\}. \tag{6}$$

We call a tour that achieves the minimum weighted delay on a finite set of points $\mathcal{S}$ in $\mathbf{P}$ with weights $w(\cdot)$ an *optimum weighted discrete surveillance route*, *OWDSR*.

For simplicity we assume that all weights are positive, that the smallest weight is equal to 1, and that the largest weight value is $w_{\max}$.

In [6], the authors show that the problem of computing an *OWDSR* is NP-hard already for the two weight values 1 and 2, that the shortest watchman route limited to see the points in $\mathcal{S}$ is a $2w_{\max}$-approximation of a *OWDSR*, and they present an $O(|\mathcal{S}|^3 n \log n)$ time constant-factor approximation algorithm for a *OWDSR* in the case of two arbitrary weight values.

4.1 A Simple Approximation Algorithm

We can immediately improve on the $2w_{\max}$-approximation in [6] as follows. Given the points in $\mathcal{S}$ and the weight values $1 = w(p_1) \leq \cdots \leq w(p_{|\mathcal{S}|}) = w_{\max}$, we scale all the weight values $w(p_i) \in [1, \sqrt{w_{\max}}[$ to 1, where $[x, y[$ denotes the right open ended interval from x to y, and all the weight values $w(p_i) \in [\sqrt{w_{\max}}, w_{\max}]$ to $\sqrt{w_{\max}}$. We next apply the c-approximation algorithm for two weight values on the scaled problem instance, giving us the following theorem.

Theorem 10 *The algorithm above computes a $c\sqrt{w_{\max}}$-approximation of an OWDSR guarding the points of $\mathcal{S}$ in $\mathbf{P}$ having arbitrary weight values in $O(|\mathcal{S}|^3 n \log n)$ time.*

4.2 An Improved Approximation Algorithm

In the following, we abuse language somewhat and say that a tour *visits* a point p, when we actually mean that the tour intersects $\mathbf{V}(p)$.

For an discrete weighted surveillance tour V, visiting the points in a finite weighted point set $\mathcal{S}$ in $\mathbf{P}$ (we assume that V is the shortest tour that visits the points in this order), we have the following inequality,

$$\forall p \in \mathcal{S} \quad hc_V(p) \leq \| V \|. \tag{7}$$

If V is such that it visits some point p' only once, then

$$\| V \| \leq 2 \cdot hc_V(p'). \tag{8}$$

Given the points of $\mathcal{S}$ in $\mathbf{P}$ with weights in the range $w(p) \in [1, w_{\max}]$, we partition the set $\mathcal{S}$ into disjoint subsets $\mathcal{S}_l$, $0 \leq l \leq M$, such that each point $p \in \mathcal{S}_l$ has $w(p) \in [w_{\max}^{l/M}, w_{\max}^{(l+1)/M}[$. We can scale the instance so that each point has weight $w_{\max}^{l/M}$. If we can find an x-approximate solution for the scaled instance, we immediately have an algorithm with approximation factor

$$x \cdot w_{\max}^{1/M} \tag{9}$$

for the original input instance.

We let $\mathcal{I}_i$, $0 \leq i \leq m$, be the nonempty sets of scaled points so that for each point $p \in \mathcal{I}_i$, the weight $w(p) = w_i = w_{\max}^{l/M}$, for some $l \geq i$. In fact, if $p \in \mathcal{I}_i$ and $p' \in \mathcal{I}_{i'}$ with $0 \leq i < i' \leq m$, then $w(p) = w_i = w_{\max}^{l/M}$ and $w(p') = w_{i'} = w_{\max}^{l'/M}$, with $l < l'$. Since the sets $\mathcal{I}_i$, $0 \leq i \leq m$, are nonempty, we have $m \leq |\mathcal{S}|$.

For each $0 \leq i \leq m$, let W_i denote a shortest tour in $\mathbf{P}$ that visits all the points in $\mathcal{I}_i$. Each such tour can be computed in $O(|\mathcal{I}_i|^3 n \log n)$ time [3, 10], and hence, all these tours can be computed in $O(|\mathcal{S}|^3 n \log n)$ time. Similarly, let T_i denote a tour in $\mathbf{P}$ with minimum delay for the scaled points in $\mathcal{I}_i$. From [6], we know that $d^w_{\mathcal{I}_i}(W_i) = d_{\mathcal{I}_i}(W_i) \leq 2d_{\mathcal{I}_i}(T_i) = 2d^w_{\mathcal{I}_i}(T_i)$ since the weights of the points in $\mathcal{I}_i$ are the same.

We furthermore define $\mathcal{I}_{i,j} = \bigcup_{i \leq \iota \leq j} \mathcal{I}_\iota$. Thus, the set $\mathcal{I}_{0,m}$ represents the scaled weight points of the original instance $\mathcal{S}$. Let $W_{i,j}$ denote a shortest tour in $\mathbf{P}$ that visits all the points in $\mathcal{I}_{i,j}$ and let $T_{i,j}$ denote a tour in $\mathbf{P}$ with minimum weighted delay for these points.

For each $0 \leq i \leq j \leq m$, we define a tour $S_{i,j}$ that visits all the points in $\mathcal{I}_{i,j}$ at least once and has short weighted delay. We have $S_{i,i} = W_{i,i} = W_i$, when $i = j$. For $i < j$, with $l \stackrel{def}{=} \lfloor (i+j)/2 \rfloor$, let $N_{i,l}$ denote the number of times points from $\mathcal{I}_{i,l}$ are visited as we follow the tour $S_{i,l}$ around once. We note that since the same point can be visited several times, $N_{i,l}$ can be substantially larger than $|\mathcal{I}_{i,l}|$. The tour $S_{i,j}$ is the tour with smallest weighted delay out of a set of tours $\{U^k_{i,j} \mid 1 \leq k \leq N_{i,l}\}$, each tour $U^k_{i,j}$ defined recursively from $S_{i,l}$ and $S_{l+1,j}$.

The tour $U^k_{i,j}$ is constructed as follows: let $r_{i,j}$ be a point on $S_{l+1,j}$ so that $\max_{p \in \mathcal{I}_{i,l}} \{SP(r_{i,j}, \mathbf{V}(p))\}$ is minimized. We denote this length by $D_{i,j}$. Let $SP(S_{i,l}, S_{l+1,j})$ be the shortest path between $S_{i,l}$ and

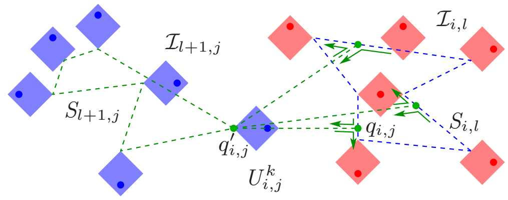

Figure 5: Schematic illustration of the construction of the tour $U^k_{i,j}$, $1 \le k \le N_{i,l}$. Red and blue regions are the visibility polygons of points in $\mathcal{I}_{i,l}$ and $\mathcal{I}_{l+1,j}$ respectively.

$S_{l+1,j}$ with endpoints $q_{i,j}$ on $S_{i,l}$ and $q'_{i,j}$ on $S_{l+1,j}$. Evidently, $\|SP(S_{i,l}, S_{l+1,j})\| \le D_{i,j}$. Let $\delta_{i,j} = \max\{D_{i,j}, \|S_{i,l}\|/2k\}$. We partition $S_{i,l}$ into at most k subpaths $Y_1, \ldots, Y_k$, each (except the last) of length $\delta_{i,j}$ and with Y_1 starting at $q_{i,j}$. $U^k_{i,j}$ is the tour obtained by first following $S_{l+1,j}$ around from $q'_{i,j}$ back to $q'_{i,j}$, then moving to $q_{i,j}$, following Y_1, moving back to $q'_{i,j}$, doing one more tour around $S_{l+1,j}$, moving to the first point of Y_2 and following Y_2, moving back to $q'_{i,j}$, make a tour around $S_{l+1,j}$, and continue alternating between following each subsequent subpath Y_κ and making a tour around $S_{l+1,j}$; see Figure 5. $U^k_{i,j}$ makes at most k rounds around $S_{l+1,j}$.

The tour among $U^1_{i,j}, \ldots, U^{N_{i,l}}_{i,j}$ with the smallest weighted delay becomes $S_{i,j}$. We show that $S_{i,j}$ has small weighted delay.

Lemma 11 *There exists a positive constant a such that*

$$d^w_{\mathcal{I}_{i,j}}(S_{i,j}) \le a^{1+\log(j-i+1)} \cdot d^w_{\mathcal{I}_{i,j}}(T_{i,j}),$$

for every $0 \le i \le j \le m$.

Proof. We make a proof by induction on $j - i$. We show the lemma to be true for $i = j$ and then proceed inductively for successively larger values of $j - i$.

From [6], we know that $d_{\mathcal{I}_i}(W_i) \le 2d_{\mathcal{I}_i}(T_i)$, for $0 \le i \le m$. Thus, all weights being equal,

$$d^w_{\mathcal{I}_i}(S_{i,i}) = d^w_{\mathcal{I}_i}(W_i) \le 2d^w_{\mathcal{I}_i}(T_i) \le a^1 \cdot d^w_{\mathcal{I}_i}(T_i), \qquad (10)$$

if $2 \le a$, proving the base case when $i = j$.

For the induction step, consider a tour $T_{i,j}$, an optimal solution for *OWDSR* in $\mathbf{P}$ that sees all the scaled points in $\mathcal{I}_{i,j}$ and has minimum weighted discrete delay.

We partition $T_{i,j}$ into subpaths as follows: let H_1 be the shortest subpath of $T_{i,j}$ that sees each point of $\mathcal{I}_{l+1,j}$ at least once, with $l \overset{def}{=} \lfloor (i+j)/2 \rfloor$ as usual, and the first visits of points in $\mathcal{I}_{i,j}$ before and after H_1 are points in $\mathcal{I}_{i,l}$. Follow $T_{i,j}$ from an endpoint of H_1 until a point of $\mathcal{I}_{l+1,j}$ is seen again. We let this subpath be L_1. Continue along $T_{i,j}$ until each point of $\mathcal{I}_{l+1,j}$ has been seen again and the next visit is to a point in $\mathcal{I}_{i,l}$, giving the subpath H_2, followed by the subpath L_2 of visits to points in $\mathcal{I}_{i,l}$, and so on. Continue subdividing $T_{i,j}$ into $2K$ subpaths, $H_1, L_1, \ldots, H_K, L_K$, for some value K, such that L_K connects back to H_1 and each H_k visits all the points of $\mathcal{I}_{l+1,j}$ and each L_k, except possibly L_K, only visits points in $\mathcal{I}_{i,l}$. The subpath L_K can visit some but not all points in $\mathcal{I}_{l+1,j}$.

For each path H_k, $1 \le k \le K$, we shortcut any detours that H_k makes to visit points in $\mathcal{I}_{i,l}$, then go back to the beginning of H_k giving us the tour H'_k. From there, we visit each (unvisited) point in $\mathcal{I}_{i,l}$ that was shortcut from H_k in the same order and continue following L_k, giving the path L'_k. Let Z be the tour $Z = \bigcup_{1 \le k \le K} H'_k \cup L'_k$. We have $\|H'_k\| \le 2\|H_k\|$ and $\|L'_k\| \le \|H_k\| + \|L_k\|$, for all $1 \le k \le K$. Hence,

$$d^w_{\mathcal{I}_{i,j}}(Z) \le 3d^w_{\mathcal{I}_{i,j}}(T_{i,j}) \text{ and } \|Z\| \le 3\|T_{i,j}\|. \qquad (11)$$

Also, for any $0 \le i \le j \le m$ and $l = \lfloor (i+j)/2 \rfloor$, we have by definition,

$$\forall k\ \forall p \in \mathcal{I}_{l+1,j} \quad hc_{T_{l+1,j}}(p) \le \|W_{l+1,j}\| \le \|H'_k\|, \qquad (12)$$

$$\forall p \in \mathcal{I}_{i,l} \quad hc_{T_{i,l}}(p) \le \|W_{i,l}\| \le \sum_{1 \le k \le K} \|L'_k\|. \qquad (13)$$

We compare the tour $U^K_{i,j}$, constructed from $S_{i,l}$ and $S_{l+1,j}$, with the tour Z constructed from an *OWDSR* $T_{i,j}$ for the point set $\mathcal{I}_{i,j}$ above. Note that we can assume that we know the value of K since we compute $U^k_{i,j}$, for all $1 \le k \le N_{i,l}$.

For a point $p \in \mathcal{I}_{l+1,j}$, the hiding cost of p is bounded by

$$\begin{aligned}
& hc_{U^K_{i,j}}(p) \le hc_{S_{l+1,j}}(p) + 2\delta_{i,j} + \|S_{i,l}\|/K \\
& \qquad \le hc_{S_{l+1,j}}(p) + 2\|S_{i,l}\|/K \\
(8) \quad & \qquad \le hc_{S_{l+1,j}}(p) + 4hc_{S_{i,l}}(p_i)/K \\
(\text{ind.}) \quad & \qquad \le a^{1+\log(j-l)} \cdot hc_{T_{l+1,j}}(p) \\
& \qquad\quad + 4a^{1+\log(l-i+1)} \cdot hc_{T_{i,l}}(p_i)/K \\
(13) \quad & \qquad \le a^{1+\log(j-l)} \cdot hc_{T_{l+1,j}}(p) \\
& \qquad\quad + 4a^{1+\log(l-i+1)} \cdot \sum_{1 \le k \le K} \|L'_k\|/K \\
& \qquad \le 4a^{1+\log(j-l)} \cdot \left(hc_{T_{l+1,j}}(p) + \max_{1 \le k \le K}\{\|L'_k\|\}\right) \\
& \qquad \le 8a^{1+\log(j-l)} \cdot \max\left\{hc_{T_{l+1,j}}(p), \max_{1 \le k \le K}\{\|L'_k\|\}\right\} \\
& \qquad \le 8a^{1+\log(j-l)} \cdot hc_Z(p) \\
(11) \quad & \qquad \le a^{1+\log(j-i+1)} \cdot hc_{T_{i,j}}(p), \qquad (14)
\end{aligned}$$

if $a \ge 24$, where $p_i \in \mathcal{I}_i$ is visited only once by $S_{i,l}$.

For a point $p \in \mathcal{I}_i$, the hiding cost of p is bounded by

$$\begin{aligned}
& hc_{U^K_{i,j}}(p) \le K \cdot \|S_{l+1,j}\| + 2K \cdot \delta_{i,j} + \|S_{i,l}\| \\
& \qquad \le K \cdot \|S_{l+1,j}\| + 2\|S_{i,l}\| \\
(8) \quad & \qquad \le 2K \cdot hc_{S_{l+1,j}}(p_{l+1}) + 4hc_{S_{i,l}}(p) \\
(\text{ind.}) \quad & \qquad \le 2Ka^{1+\log(j-l)} \cdot hc_{T_{l+1,j}}(p_{l+1}) \\
& \qquad\quad + 4a^{1+\log(l-i+1)} \cdot hc_{T_{i,l}}(p) \\
(12), (13) \quad & \qquad \le 2Ka^{1+\log(j-l)} \cdot \min_{1 \le k \le K}\{\|H'_k\|\} \\
& \qquad\quad + 4a^{1+\log(l-i+1)} \cdot \sum_{1 \le k \le K} \|L'_k\|
\end{aligned}$$

$$\begin{aligned}
&\leq 4a^{1+\log(j-l)} \cdot \Big(\sum_{1 \leq k \leq K} \|H'_k\| + \|L'_k\| \Big) \\
&\leq 4a^{1+\log(j-l)} \cdot \|Z\| \\
(8) \qquad &\leq 8a^{1+\log(j-l)} \cdot hc_Z(p) \\
(11) \qquad &\leq a^{1+\log(j-i+1)} \cdot hc_{T_{i,j}}(p), \qquad (15)
\end{aligned}$$

if $a \geq 24$, since $p \in \mathcal{I}_i$ is visited only once by $S_{i,l}$ and has maximal hiding cost among the points in $\mathcal{I}_{i.j}$, and $p_{l+1} \in \mathcal{I}_{l+1}$ is visited only once by $S_{l+1,j}$.

For a point $p \in \mathcal{I}_{i+1,l}$, the point p is in the upper half of some recursive division of the sets $\mathcal{I}_{i,l}, \mathcal{I}_{i,\lfloor(i+l)/2\rfloor}, \mathcal{I}_{i,\lfloor(i+\lfloor(i+l)/2\rfloor)/2\rfloor}, \ldots, \mathcal{I}_{i,i+1}$, for which an inequality similar to (14) and (15) applies. We omit the details.

Thus, for any point $p \in \mathcal{I}_{i,j}$, we have that there is a constant $a > 1$ such that

$$hc_{S_{i,j}}(p) \leq a^{1+\log(j-i+1)} \cdot hc_{T_{i,j}}(p)$$

and the lemma is therefore proved. □

We compute $S_{0,m}$ by establishing $U^k_{0,m}$, for every $1 \leq k \leq N_{0,\lfloor m/2 \rfloor}$, and each of these are computed recursively from $S_{0,\lfloor m/2\rfloor}$ and $S_{\lfloor m/2\rfloor+1,m}$. Given these two tours, $U^k_{0,m}$ is constructed by copying $S_{\lfloor m/2\rfloor+1,m}$ at most k times and connecting each tour to the at most k subpaths of $S_{0,\lfloor m/2\rfloor}$ using shortest paths. This takes $O(k|\mathcal{S}| + kn)$ time for each $U^k_{0,m}$. Note that $N_{0,\lfloor m/2\rfloor} \in |\mathcal{S}|^m$, thus this takes $O(|\mathcal{S}|^{m+1}n)$ time in total. At each level of the recursion we use this amount of time and we have $\log m+1$ levels, giving us $O(|\mathcal{S}|^{O(m)}n)$ time. The preprocessing step of computing all the discrete watchman routes W_i, for $0 \leq i \leq m$ takes $O(|\mathcal{S}|^3 n \log n)$ time, and hence, the total complexity is bounded by $O(|\mathcal{S}|^{O(m)} n \log n)$. We note that the complexity is superpolynomial in $|\mathcal{S}|$ and n.

By carefully considering the weight ratios in the construction, we could perhaps limit the computation to the relevant values of k, reducing the necessity to compute $U^k_{i,j}$ for all values of k up to $N_{i,l}$. We could thus potentially make the algorithm take polynomial time in the size of the output tour $S_{0,m}$.

Theorem 12 *There is an $O(|\mathcal{S}|^{O(\log w_{\max})} \cdot n \log n)$ time algorithm that computes a $O(\text{polylog}\, w_{\max})$-approximate weighted discrete surveillance tour to the original unscaled weighted point set $\mathcal{S}$ in $\mathbf{P}$ having n edges.*

Proof. Apply the algorithm described above with

$$m = \max\{3, \lceil \log w_{\max}/\log\log w_{\max})\rceil\},$$

where a is the constant in Lemma 11 and $m+1$ is the number of weight values, with all the original weights scaled to the lowest value in their respective interval $[w_{\max}^{i/m}, w_{\max}^{(i+1)/m}[$, for $0 \leq i \leq m$.

From Lemma 11, we have that the scaled instance is approximated within an approximation factor of $a^{1+\log(m+1)} \leq a^2 m^{\log a}$ and by our choice of the value m, we have $a^2 m^{\log a} \geq w_{\max}^{1/m}$ and by (9), the approximation factor for the unscaled instance is bounded by $w_{\max}^{1/m} \cdot a^2 m^{\log a} \leq a^4 m^{2\log a} \in O(\text{polylog}\, w_{\max})$.

The running time follows from the discussion above. □

5 Conclusions

We present a linear time 3/2-approximation algorithm for the optimum surveillance tour problem in rectilinear polygons in the L_1-metric. It is still an open problem whether an optimum tour can be computed in polynomial time assuming P $\neq$ NP. We believe that the same approach should also give a 3/2-approximation for general simple polygons in the L_1-metric.

We also present an $O(\text{polylog}\, w_{\max})$-approximation algorithm for the optimum weighted discrete surveillance route in a simple polygon with weight values in the range $[1, w_{\max}]$.

The deeper complexity relationships of the optimum weighted discrete surveillance tour problem in simple polygons remains to be investigated. For two weight values, the problem is NP-hard but constant factor approximable [6]. It is not evident that a polynomial time constant factor approximation algorithm exists for the general problem assuming P $\neq$ NP.

References

[1] W. Chin, S. Ntafos. Optimum Watchman Routes. *Information Processing Letters*, 28:39–44, 1988.

[2] W. Chin, S. Ntafos. Shortest Watchman Routes in Simple Polygons. *Disc. & Comp. Geometry*, 6(1):9–31, 1991.

[3] M. Dror, A. Efrat, A. Lubiw, J.S.B. Mitchell. Touring a Sequence of Polygons. In *Proc. 35th STOC*, p 473–482, 2003.

[4] J.S.B. Mitchell. Personal communication, 2017.

[5] T. Kamphans, E. Langetepe. Inspecting a Set of Strips Optimally. In *Proc. 11th WADS*, p 423–434. LNCS 5664, 2009.

[6] B.J. Nilsson, E. Packer. Weighted Discrete Surveillance Tours in Simple Polygons. In *Proc. 33rd EuroCG*, 2017.

[7] J. O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, 1987.

[8] J. O'Rourke. Visibility. In J.E. Goodman, J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 28. CRC Press, 1997, 2nd ed. 2004.

[9] T.C. Shermer. Recent Results in Art Galleries. *Proceedings of the IEEE*, p 1384–1399, 1992.

[10] X.-H. Tan. Fast Computation of Shortest Watchman Routes in Simple Polygons. *Information Processing Letters*, 77(1):27–33, 2001.

[11] J. Urrutia. Art Gallery and Illumination Problems. In J.R. Sack, J. Urrutia, editors, *Handbook on Computational Geometry*, chapter 22. Elsevier, 1999.