On the See-Through Watchman Route Problem and the Quota-TSP Problem on Infinite Lines

Joseph S. B. Mitchell
Stony Brook University, NY, USA

Linh Nguyen
Florida A&M University, Tallahassee, FL, USA

Abstract

The classic Watchman Route Problem (WRP) seeks to compute a shortest tour in a polygonal domain that sees every point of the domain. We introduce and study a novel generalization of the WRP, the See-Through Watchman Route Problem (STWRP), in which, in addition to vision-blocking "walls" of an input domain, there are obstacles to motion that are *not* opaque to vision: the watchman can see through certain obstacles or portions of the boundary of a polygonal domain P. This setting is motivated by real-world situations that may include transparent barriers (e.g., glass walls), obstacles that obstruct movement but not vision (e.g., lakes, flowerbeds, or potholes), and robotic sensors with penetration capabilities (e.g., microwave imaging). To the best of our knowledge, this version of the problem is new to the algorithms community.

Our main result is an FPTAS for the STWRP in the case that P is an opaque-walled simple polygon having within it a set of transparent obstacles. A closely related problem that arises in this setting is that of the Traveling Salesperson problem with neighborhoods (TSPN) on a set of lines in the plane, with obstacles. We give the first FPTAS for the Quota-TSPN on infinite lines with polygonal obstacles. Additionally, we show tightness of our FPTAS, in that the Quota-TSPN on infinite lines with obstacles is weakly NP-hard. In the case of the STWRP within a simple polygon P with portions of the boundary, ∂P, being transparent, we prove that the problem is NP-hard to approximate within a factor better than $O(\log n)$.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Watchman route problem, TSP with neighborhoods

Digital Object Identifier 10.4230/LIPIcs.ESA.2026.24

1 Introduction

Sensor technologies have seen rapid growth in recent years. Some robots can now be equipped with sensors that can "see through" objects, e.g., using a thermal camera or microwave imaging to penetrate various non-metallic or non-conductive materials [7, 9, 30]. This introduces many coverage path planning problems in a new, interesting setting, where obstacles may block movement but not coverage. In this paper, we consider a generalization of the classic Watchman Route Problem (WRP), which seeks a shortest route for a robot equipped with an omnidirectional camera to see every point in a polygonal domain. Traditionally, the obstacles and the boundary of the domain restrict both movement and visibility. Here, in the See-Through Watchman Route Problem (STWRP), we allow the watchman to see beyond some obstacles or portions of the boundary of the domain. Other real-world settings can also be modeled in this setting, e.g., in an open field, a lake might obstruct traversal but still permit a clear line of sight, whereas a boulder disallows both movement and vision.

The best known approximation factor in polynomial time for the WRP with opaque obstacles (holes) is poly-logarithmic [25]. Ideally, when the watchman can see beyond obstacles, we want to exploit that capability in order to obtain a shorter route. We show that in the setting in which the watchman can see through all obstacles (i.e., the STWRP within a polygonal domain P with an opaque outer boundary and transparent hole boundaries), better approximations are possible. In particular, we give a fully polynomial-time approximation scheme (FPTAS), with approximation factor $(1 + \varepsilon)$ and running time poly$(n, 1/\varepsilon)$. In contrast, when portions of ∂P are transparent, the STWRP is NP-hard to approximate within a factor better than $O(\log n)$, even if P is a simple polygon.

Connections between the WRP and variants of the Travelling Salesman Problem (TSP) are well studied (e.g., [13, 32, 35, 25]). We build on this research and consider the following variant, which arises in addressing the STWRP: the Quota-TSP on infinite lines with obstacles, where each line has an associated reward, and the goal is to find a shortest obstacle-avoiding path that collects a given quota of reward. We design an FPTAS for this problem and also prove a matching lower bound by showing that the problem is weakly NP-hard. This gives a tight characterization of its approximability in polynomial time.

2 Related Work

Visibility-based search and surveillance problems arise in a wide range of applications, including robotics, wireless sensor networks, and security; they have also been studied extensively in theoretical algorithms and computational geometry. Two classic problems in this area are the Art Gallery problem, which asks for the minimum number of static guards needed to observe a polygonal domain, and the Watchman Route Problem, introduced by Chin and Ntafos [10, 11], which seeks a shortest tour for a mobile guard to achieve full visibility. For simple polygons, although some early results were found to be flawed [8, 11], polynomial-time algorithms are now known. Specifically, for the anchored case where a starting point that the route must pass through is given, Tan, Hirata and Inagaki gave an $O(n^4)$ algorithm [35, 36], which was later improved to improved to $O(n^3 \log n)$ [13], and recently a slight improvement by a log factor has also been proposed [37]; the algorithms for the (more complicated) floating case are $O(n)$-factor slower than their anchored counterparts. In addition, several linear-time constant-factor approximation algorithms are also known [27, 33, 34], offering efficient and practical solutions.

The problem becomes much harder in the presence of obstacles. For polygons with holes, the Watchman Route Problem is NP-hard, as shown by Chin and Ntafos [10] and later revisited in [16]. Mitchell [25] further showed that it is NP-hard to approximate the minimum-length watchman route within a factor better than $O(\log n)$, where n is the combinatorial complexity of the input polygon. The best polynomial-time approximation algorithm has an $O(\log^2 n)$ factor [25], and it remains open whether an $O(\log n)$-approximation is achievable.

More recently, visibility models that allow seeing through parts of the domain boundary have also been explored. Nilsson and Schmidt [28] introduced the k-transmitter WRP, where the watchman can see through up to k walls. They proved that even in this setting, approximating the minimum-length route within a logarithmic factor is NP-hard. They also gave a polylogarithmic-factor approximation algorithm for achieving visibility of a discrete set of points in a simple polygon. Follow-up work [6] showed that the hardness persists even when restricted to some special classes of polygons. Variants of the k-transmitter model have also been considered in other coverage problems, which are of independent interest.

The WRP is closely related to the geometric Traveling Salesperson Problem with Neighborhoods (TSPN), which asks for a shortest tour visiting a collection of connected geometric regions (the neighborhoods). The classical geometric TSP on a set of points admits a polynomial-time approximation scheme (PTAS) [4, 23], but for general connected regions, the best polynominal-time approximation factor for TSPN is $O(\log n)$ [18, 21]. TSPN was originally posed for neighborhoods that are translates of the same segment or polygon, and constant-factor approximations were shown in those cases [3]. In general, the problem is APX-hard if neighborhoods are arbitrary connected regions and may overlap [12, 31].

Some special cases of TSPN admit better approximations, such as constant-factor algorithms when neighborhoods have comparable diameters [14], or a PTAS when the regions are disjoint and fat [24]. However, these results do not directly apply to the Watchman Route Problem, since the watchman must "visit" an uncountably infinite set of visibility regions – specifically, every point in the continuum introduces a visibility region that must be visited. In a simpler setting where the neighborhoods are infinite lines in the plane (related to the WRP in simple, hole-free polygons), the problem is solvable in polynomial time [13, 32, 35]. In three or more dimensions, TSPN on lines becomes NP-hard; [17] gave an $O(\log^3 n)$-approximation in polynomial time and [2] gave an $O(\log^2 n)$-approximation in quasi-polynomial time.

Obstacles make the problem even more challenging and, to our knowledge, there is little prior theoretical work on TSPN with arbitrary obstacles. The results of the papers on Watchman Route on a connected collection of infinite lines in [15, 20], although not phrased in terms of TSPN with obstacles specifically, imply a polynomial-time algorithm for a special case where movement is restricted to the lines. (The obstacles are the convex faces formed by the lines.)

3 The Watchman Route Problem with Transparent Obstacles

3.1 Preliminaries

Let $P \subset \mathbb{R}^2$ be a connected polygonal domain. The boundary of P consists of: (i) the *outer boundary*, denoted $\partial_{\text{outer}} P$, which is a simple closed polygonal curve made up of a finite number of line segments; and (ii) the *inner boundary*, which consists of h simple closed polygonal cycles, each corresponding to a hole or obstacle in the domain. We refer to such a region as a *polygonal domain*, or a *polygon with holes*. Let n denote the total number of vertices on the boundary of P. A vertex on the outer boundary is called *reflex* if its internal angle is more than 180°, and *convex* if its internal angle is less than 180°. We assume non-degeneracy: no internal angle is exactly 180°. Throughout this paper, we also assume that polygons are closed sets, meaning they include both their interior and boundary.

In the classical definition of visibility, the entire boundary of P – both outer and inner – is opaque. In this model, two points $p, q \in P$ see each other if the segment pq lies entirely within P. We extend this notion by allowing certain portions of the boundary of P (specified as part of the input) to be either *occluding*, which means they block visibility, or *transparent*, meaning they allow lines of sight to pass through. We consider two settings:

- **Transparent obstacles:** Some obstacles in P may have boundaries composed entirely of transparent segments.
- **Transparent walls:** Portions of the outer boundary of P may be transparent, in which case a watchman inside P can see through a transparent wall into the exterior, and possibly back into another part of P (via a different transparent wall).

Traversal through obstacles or through the outer boundary remains forbidden. Naturally, these two settings can be combined: P may simultaneously include transparent walls and transparent obstacles. Under this generalized visibility model, two points $p, q \in P$ *see* each other if one of the following conditions holds:

(i) The segment pq lies entirely within P; or

(ii) For every maximal subsegment of pq that lies outside of P, both of its endpoints lie on transparent (either outer or inner) boundary segments of P.

The visibility region of a point p, denoted $V(p)$, is the set of all points in P that are visible from p (see Figure 1). In general, $V(p)$ may consist of multiple connected components. For a subset $X \subseteq P$, the visibility region is defined as the union of the visibility regions of all points in it: $V(X) = \bigcup_{p \in X} V(p)$. The visibility region of a point or segment in a polygonal domain can be computed efficiently (see [29]), and standard visibility algorithms can be readily adapted to our setting (for example, to compute the visibility region of a point in a polygon P with transparent boundary segments, treat the occluding components of the boundary as a collection of "skinny" obstacles and compute the classic visibility region, then take its intersection with P).

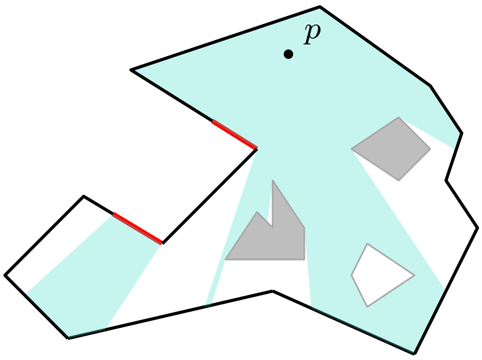

Figure 1 The visibility region of p shaded in turquoise. There are two occluding gray obstacles and one transparent obstacle. The outer boundary of P includes two transparent segments, shown in red, which allow lines of sight to pass through.

The Watchman Route Problem (WRP) asks for a tour (closed route) γ such that $V(\gamma) = P$ and $|\gamma|$ is minimized. We also use γ to denote the optimal route in the STWRP. Here, $|\cdot|$ denotes Euclidean length, although we also use $|\cdot|$ to denote the cardinality of a discrete set; the meaning will be clear from context.

In this section, we first consider the case where **all obstacles (holes) in P are transparent**. We further assume that a starting point $s \in \partial_{\text{outer}} P$ is given.

3.2 Localization and Discretization

Let $\pi_P(s,t)$ denote a *geodesic shortest path* from $s \in P$ to $t \in P$ with respect to P–that is, a path of minimum length from s to t that lies entirely within P. Since P is a polygonal domain, $\pi_P(s,t)$ is necessarily a polygonal path, there may however be exponentially many shortest paths of equal length. The *geodesic distance* between s and t is the length $|\pi_P(s,t)|$. The geodesic distance from s to a subset $X \subseteq P$ is denoted by $d(s,X)$ and defined as $d_P(s,X) := \min_{x \in X} |\pi_P(s,x)|$. ($P$ is compact, therefore so is X which means a minimum exists.)

Given a subset $X \subseteq P$, the *geodesic convex hull* of X, denoted $\mathrm{GCH}(X)$, is the inclusion-wise minimal subset of P that contains X and all geodesic shortest paths $\pi_P(s,t)$ for every pair of points $s,t \in X$. The outer boundary of $\mathrm{GCH}(X)$ is defined as the boundary of the face at infinity in the arrangement formed by these paths; it is known that the geodesic convex hull of X has the minimum-length outer boundary of all supersets of X in P. A subset of P is said to be *geodesically convex* if it is closed under geodesic shortest paths; the outer boundary of such a set is called a *geodesically convex curve*. Geodesic shortest paths and geodesic convex hulls in polygonal domains are well studied and can be computed in polynomial time; see [5, 26] for further details.

A *visibility cut* c_i with respect to a starting point s is the maximal (with respect to P) extension of an edge e incident to a reflex vertex v_i which divides P into two subpolygons so that v_i is a convex vertex in the subpolygon containing s (visibility cuts can be thought of as lines that the watchman must travel to from the other side in order to see into nooks and crannies of the domain; if the edge extension passes through s, we do not consider it a visibility cut). The subpolygon on the opposite side, i.e. the one not containing s, is the *pocket* induced by c_i. A visibility cut c_i is said to be *essential* if its pocket does not fully contain the pocket of any other visibility cut. There are $O(n)$ essential cuts as a reflex vertex induces at most one of them. See Figure 2.

Recall that in the classic WRP without obstacles (i.e., P is simple), it is well known that an optimal route visits the essential cuts in the order in which they appear around the boundary of P [11, 13]. This property enables the problem to be solved efficiently using the polygon touring algorithm [13, 37]. However, in the presence of transparent obstacles, while it remains necessary for the watchman to visit all essential cuts (proof of Lemma 1), the order in which these cuts must be visited is no longer easily computable (Figure 2).

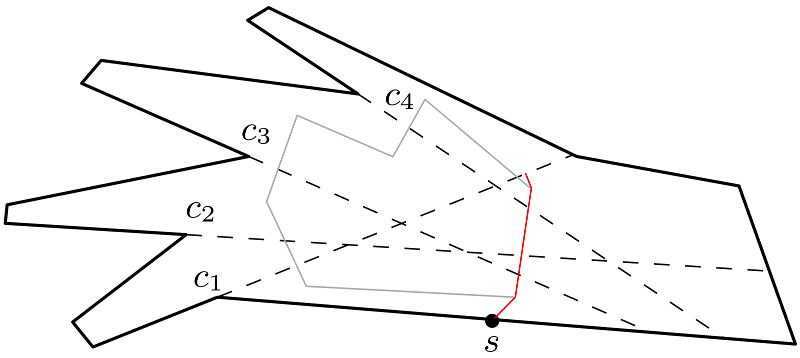

Figure 2 A polygon with one transparent obstacle and four essential cuts with respect to s. An optimal see-through watchman route (red) goes around the obstacle and visits the essential cuts in order c_3, c_2, c_4, c_1 (and then goes back to s).

We first present the structure and combinatorial bound of an optimal see-through watchman route with Lemma 1.

▶ **Lemma 1** (proof in Appendix A.1). *An optimal see-through watchman route γ is geodesically convex and has $O(n)$ vertices that lie interior to P.*

Another challenge is that, in general, there is no known polynomial-size candidate set of points within P that suffices as potential turning points for constructing an optimal watchman tour. We address this by showing how to localize the search for an optimal solution γ and discretize the domain P into a polynomial number of candidate vertices for γ.

A *geodesic disk* of radius r centered at s is defined as the set of points within P whose geodesic distance to s is at most r, i.e., $\{x \in P \mid d_P(s,x) \leq r\}$. As we increase r, the disk sees more and more of P (by definition if $X_1 \subseteq X_2$ then $V(X_1) \subseteq V(X_2)$). Denote by GD_s the *minimal illuminating geodesic disk*, which is the geodesic disk centered at s with the smallest radius such that the outer boundary of GD_s (the geodesic circle) sees all of P. To compute the radius r of GD_s, we calculate the maximum geodesic distance from s to all essential cuts, each of which may be subdivided into at most $O(n)$ connected components by the obstacles. Formally, $r = \max_c\{d_P(s,c)\}$, where c ranges over all essential cuts.

▶ **Lemma 2** (proof in Appendix A.2). $r \leq |\gamma| \leq O(nr)$

Lemma 2 implies that γ lies within an axis-aligned square B (or more precisely, within $B \cap P$) of side length $O(nr)$ (specifically, cnr for any $c > 2 + 2\pi$) centered at s. We refer to B as the *localizing square*. We now define a decomposition of P into convex cells, whose vertices serve as candidate turn points for an approximately optimal watchman route, utilizing the localizing square.

(i) Triangulate P, including s as a vertex of the triangulation.

(ii) Overlay a regular square grid onto the triangulation, partitioning the portion of the localizing square B that lies within P into pixels of side length $\delta = \frac{\varepsilon r}{12n}$.

Let $\mathcal{S}_\varepsilon$ denote the set of vertices in the resulting decomposition, $\mathcal{S}_\varepsilon$ includes intersections of triangulation diagonals with grid lines, as well as vertices of P. Therefore $|\mathcal{S}_\varepsilon| = O\left(\left(\frac{n^2}{\varepsilon}\right)^2\right)$. The decomposition yields convex cells, each of perimeter at most 4δ. The discretization lemma below shows that, by giving up an approximation factor of at most $(1+\varepsilon)$, $\mathcal{S}_\varepsilon$ can serve as the set of *candidate turn points*.

▶ **Lemma 3.** *There exists a geodesically convex tour γ' with vertices in $\mathcal{S}_\varepsilon$ such that $V(\gamma') = P$ and $|\gamma'| \leq (1+\varepsilon)|\gamma|$.*

Proof. See Figure 3. Let e be an edge of γ with endpoints located in cells σ_i and σ_j of the decomposition described above. If the endpoints of e are not themselves cell vertices, we replace e with a path (homotopic to e) composed of segments connecting vertices in $\mathcal{S}_\varepsilon$. Intuitively, imagine a rubber band stretched along e, pinned at the points where e intersects the boundaries of σ_i and σ_j. The band is then pulled outward – away from the interior of γ – until its endpoints coincide with cell vertices u' and v'. Throughout this adjustment, the band remains taut and is pinned against all obstacles in P with which it comes into contact. This transforms e into e', and ote that $|e'| \leq |uu'| + |uv| + |vv'| \leq 4\delta + |e| + 4\delta = |e| + 8\delta$. We replace e with e' and the boundary of two convex cells σ_i, σ_j.

Applying this procedure to every edge of γ with at least one endpoint in the interior of a cell in the decomposition (of which there are at most n, by Lemma 1) yields an embedded planar straight-line graph G with total edge length no greater than $|\gamma| + n \cdot 8\delta + n \cdot 4\delta \leq (1+\varepsilon)|\gamma|$. Since G surrounds γ, clearly G also sees all of P since the obstacles are transparent

Finally, we can replace G with γ', the outer boundary of the geodesic convex hull of G. It is straightforward to see that all vertices of γ' must still be a subset of $\mathcal{S}_\varepsilon$, and that $|\gamma'|$ is no greater than the total edge length of G. Thus we have demonstrated the existence of a geodesically convex, visibility covering tour whose length is no more than $(1+\varepsilon)|\gamma|$ with vertices in $\mathcal{S}_\varepsilon$. ◀

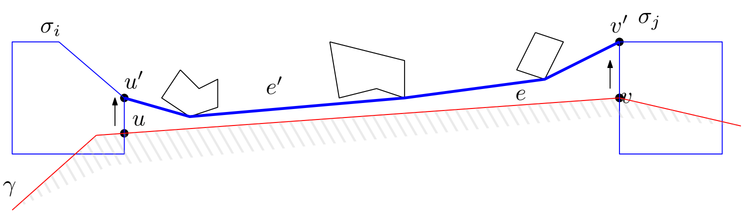

Figure 3 Replacing each edge (red) with the perimeters of the two cells containing its endpoints and a geodesic path of the same homotopy type (blue).

3.3 Structural lemmas

For $s, t \in P$, let $\pi(s, t)$ denote the *pseudo-geodesic shortest path* from s to t, defined as the shortest path with respect to only the outer boundary $\partial_{\text{outer}}P$, ignoring the holes. Equivalently, let $\overline{P}$ be the simple polygon obtained by filling in all the holes of P, then $\pi(s, t)$ is the unique shortest $s - t$ path in $\overline{P}$, since shortest paths in simple polygons are unique. Note that $\pi(s, t)$ is not necessarily a valid path in P, and is used solely for purposes of formal analysis.

For a point x lying in a hole of P, define its *pseudo-visibility region*, denoted by $V(x)$, as the set of all points in P from which a line segment can be drawn to x without crossing the unbounded face outside $\partial_{\text{outer}}P$. This definition naturally extends to subsets of $\overline{P}$ (e.g., pseudo-paths in P that possibly pass through holes), in the same manner as the standard visibility region. For convenience and to avoid notation overload, we use $V(\cdot)$ for both standard and pseudo-visibility regions, with the distinction made implicitly based on whether the argument includes points inside a hole. See Figure 4 (top) for an illustration.

Given a subset $X \subseteq P$ or $X \subseteq \overline{P}$, the *pseudo-relative convex hull*, denoted by $PCH(X)$, is defined as the inclusion-wise minimal subset of $\overline{P}$ that contains X and is closed under taking pseudo-geodesic shortest paths. (In a simple polygon, the geodesic convex hull is also called the *relative convex hull*. The pseudo-relative convex hull of X is the relative convex hull of X within $\overline{P}$ but may contain points interior to holes of P.) A subset X is said to be pseudo-relatively convex if $PCH(X) = X$, and its boundary is said to be pseudo-relatively convex (similar to how the boundary of a convex set in the plane is a convex curve). Figure 4 (bottom) illustrates the distinction between the geodesic convex hull and the pseudo-relative convex hull.

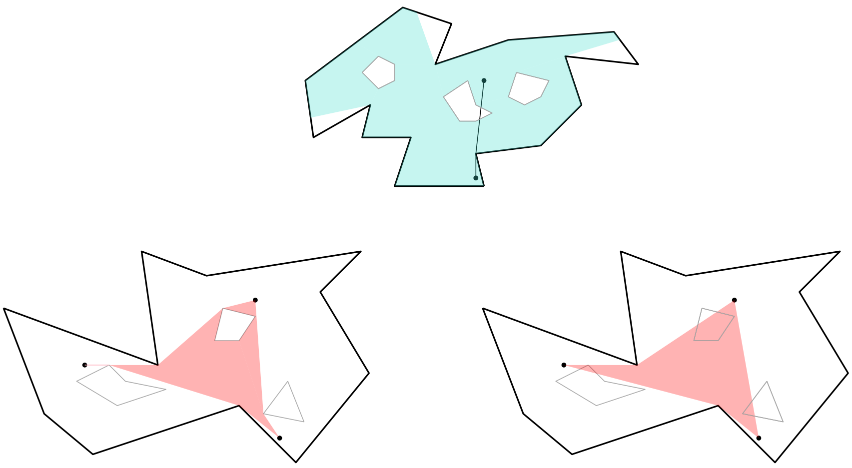

Figure 4 Top: A pseudo-path and its pseudo-visibility region. Bottom: The geodesic convex hull (left) and the pseudo-relative convex hull (right) of three points in P.

▶ **Lemma 4** (proof in Appendix A.3). *$V(\gamma) = V(PCH(\gamma))$. Generally, the pseudo-visibility region of any connected set is equal to that of its pseudo-relative convex hull.*

Lemma 4 implies that a curve γ sees all of P if and only if its relative convex hull sees all of P. Now, we claim that the problem of computing the shortest watchman route γ is equivalent to the Minimum Covering Pseudo-Relatively Convex Polygon problem: compute a pseudo-relatively convex polygon $\mathcal{P} \subset \overline{P}$ whose vertices $s, p_1, p_2, \ldots, p_m$ (in order around $\mathcal{P}$) are within P such that (i) $\mathcal{P}$ covers (intersects) all essential cuts, and (ii) $|\pi_P(s,p_1)| + |\pi_P(p_1,p_2)| + \ldots + |\pi_P(p_m,s)|$, which we refer to as the geodesic perimeter of $\mathcal{P}$, is minimized. We provide clarification and justification in the following lemma.

▶ **Lemma 5.** *An optimal solution $\mathcal{P} \subset \overline{P}$ to the* Minimum Covering Pseudo-Relatively Convex Polygon *problem yields an optimal see-through watchman tour γ, and vice versa.*

Proof. Let $\mathcal{P} \subseteq \overline{P}$ be an optimal solution to the Minimum Covering Pseudo-Relatively Convex Polygon problem, with vertices $s, p_1, p_2, \ldots, p_m$. We concatenate geodesic paths $\pi_P(s,p_1), \pi_P(p_1,p_2), \ldots, \pi_P(p_m,s)$ to form a route γ, which sees all of P. Explicitly, note that the pseudo-relatively convex polygon enclosed by γ is the pseudo-relative convex hull of its vertices. Since $\{s, p_1, \ldots, p_m\} \subset \gamma$, we have $PCH(\{s, p_1, \ldots, p_m\}) \subseteq PCH(\gamma)$, and hence $PCH(\gamma)$ sees the entire polygon.

Let $|\gamma| = |\pi_P(s,p_1)| + |\pi_P(p_1,p_2)| + \ldots + |\pi_P(p_m,s)| = L$. We claim that no shorter route than γ exists that sees all of P. Suppose, for contradiction, that there exists $\widetilde{\gamma} \subset P$ such that $V(\widetilde{\gamma}) = P$ and $|\widetilde{\gamma}| = \widetilde{L} < L$. The convex hull $\widetilde{\mathcal{P}}$ of $\widetilde{\gamma}$ (which must intersect all essential cuts $\widetilde{\gamma}$ intersects) is a feasible solution to the Minimum Covering Pseudo-Relatively Convex Polygon problem with geodesic perimeter no greater than $\widetilde{L}$, contradicting the optimality of $\mathcal{P}$.

Conversely, suppose $\gamma \subseteq P$ is a shortest see-through watchman route, and let $|\gamma| = L$. We claim that the pseudo-relative convex hull of the vertices of γ (which encloses γ and thus sees all of P) is an optimal solution to the Minimum Covering Pseudo-Relatively Convex Polygon problem. If there exists a covering pseudo-relatively convex polygon $\widetilde{\mathcal{P}} = (s, p_1, \ldots, p_m)$ of geodesic perimeter strictly less than L, then by our definition of the geodesic perimeter of $\widetilde{\mathcal{P}}$, concatenating the geodesic paths $\pi_P(s,p_1), \pi_P(p_1,p_2), \ldots, \pi_P(p_m,s)$ yields a route $\widetilde{\gamma}$ where $|\widetilde{\gamma}| < L$. This contradicts the optimality of γ. ◀

We note that in the statement of Lemma 5, $\mathcal{P}$ is not necessarily the pseudo-relative convex hull of γ. Rather, we are optimizing over pseudo-relatively convex sets with respect to the geodesic metric induced by P. See Figure 5. Our dynamic programming algorithm computes an optimal set of points $p_1, \ldots, p_m$ such that $s, p_1, \ldots, p_m$ are in *pseudo-relatively convex position*, i.e. they constitute the vertex set of a pseudo-relatively convex polygon. Then, we return $\gamma := \pi_P(s,p_1) \cup \pi_P(p_1,p_2) \cup \ldots \cup \pi_P(p_m,s)$. Lemma 5 also implies that it suffices to restrict the candidate vertices for computing $\mathcal{P}$ to the finite set $\mathcal{S}_\varepsilon$, in order to obtain a $(1+\varepsilon)$-approximate solution to the Minimum Covering Pseudo-Relatively Convex Polygon problem

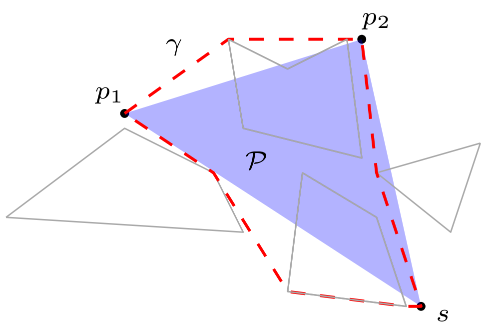

Figure 5 Joining the geodesic shortest paths between convex vertices of a pseudo-relatively polygon into a watchman route.

3.4 The dynamic programming algorithm

We can view the Minimum Covering Pseudo-Relatively Convex Polygon problem as a process of drawing a pseudo-relatively convex curve in P (ignoring the obstacles) for the watchman to "travel" (geodesically) and "collect" rewards – in this case, the essential cuts. However, the watchman may encounter the same essential cut at most twice along the tour. Naively tracking all essential cuts visited at any point would result in an exponential number of subproblems. The main challenge lies in representing each subproblem compactly while still enabling efficient identification of newly collected essential cuts at each step.

▶ **Lemma 6** (proof in Appendix A.4)**.** *For any $x, y \in P$, if the pseudo-geodesic shortest path $\pi(x, y)$ intersects a chord c of $\overline{P}$ (the simple polygon obtained by filling all holes of P), then every path from x to y, including pseudo-paths passing through holes, must intersect c.*

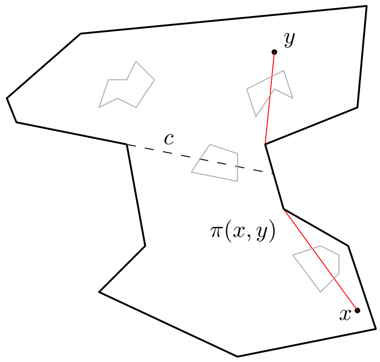

Figure 6 Illustration for Lemma 6.

Now, consider an optimal covering pseudo-relatively convex polygon $\mathcal{P}$ with vertices $(s, p_1, \ldots, p_m)$. For each i, let $\mathcal{P}_i$ be the portion of the boundary of $\mathcal{P}$ from s to p_i, and let $C(X)$, where $X \subseteq \overline{P}$, denote the set of essential cuts intersected by X. (Note that $\mathcal{P}_i$ can also be thought of as a path, and we will refer to $\mathcal{P}_i$ as such in some places.)

▶ **Lemma 7.** $C(\mathcal{P}_i) \cap C(p_i p_{i+1}) = C(\pi(s, p_i)) \cap C(p_i p_{i+1})$.

Proof. By Lemma 6, any essential cut, which is a chord of $\overline{P}$, that intersects $\pi(s, p_i)$ also intersects $\mathcal{P}_i$, therefore $C(\pi(s, p_i)) \subseteq C(\mathcal{P}_i)$, and hence $C(\pi(s, p_i)) \cap C(p_i p_{i+1}) \subseteq C(\mathcal{P}_i) \cap C(p_i p_{i+1})$.

For the opposite inclusion, let c be an essential cut that intersects both $\mathcal{P}_i$ and $p_i p_{i+1}$. We claim that c must also intersect $\pi(s, p_i)$. Indeed, since $\mathcal{P}$ is pseudo-relatively convex, c must intersect it in one contiguous segment (otherwise there are two points on c which the shortest path between them is not in $\mathcal{P}$); one end point of that segment is on $\mathcal{P}_i$, the other end point is on $p_i p_{i+1}$, thus they are on opposite sides of $\pi(s, p_i)$ (which separates $\mathcal{P}_i$ and $p_i p_{i+1}$) and must also intersect $\pi(s, p_i)$. See Figure 7. ◀

Suppose we have constructed $\mathcal{P}_i$ in the form of a sequence of points in pseudo-relatively convex position from s to p_i, and we want to determine the next point p_{i+1}. Lemma 7 implies that if p_{i+1} and the vertices of $\mathcal{P}_i$ are in pseudo-relatively convex position, then the set of essential cuts intersecting both $\mathcal{P}_i$ and the segment $p_i p_{i+1}$ is determined solely by p_i and p_{i+1}, and is independent of the precise geometry of $\mathcal{P}_i$. In other words, the set of essential cuts that are *newly* visited by $p_i p_{i+1}$ is fixed as long as pseudo-relative convexity is maintained. Our algorithm exploits this geometric optimal substructure.

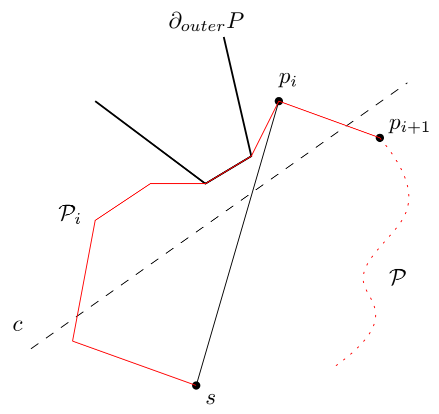

Figure 7 Proof of Lemma 7.

▶ **Corollary 8.** *For $0 \leq i \leq m$, the path $\mathcal{P}_{i+1}$ is the shortest path with vertices in pseudo-relatively convex position that intersects at least k essential cuts if and only if:*

(i) *p_{i+1} and the vertices of $\mathcal{P}_i$ are in pseudo-relatively convex position ($p_{m+1} \equiv s$), and*

(ii) *$\mathcal{P}_i$ is the shortest path that intersects at least $k - |C(p_i p_{i+1}) \setminus C(\pi(s, p_i))|$ essential cuts.*

3.5 The main theorem

▶ **Theorem 9.** *The STWRP with transparent obstacles has a fully polynomial-time approximation scheme.*

Proof. We first compute the minimal illuminating geodesic disk and then construct $\mathcal{S}_\varepsilon$, the set of candidate turn points. By Lemma 3, there exists a grid-rounded solution γ' within a factor of $(1 + \varepsilon)$ of the optimum. Consequently, there also exists a grid-rounded solution $\mathcal{P}$ to the MINIMUM COVERING PSEUDO-RELATIVELY CONVEX POLYGON problem with the same approximation factor. We now describe how to find $\mathcal{P}$ using dynamic programming.

Some preprocessing is required. We begin by establishing an ordering on the set of candidate turn points $\mathcal{S}_\varepsilon$. First, we compute $\mathcal{T}$, the tree of pseudo-geodesic shortest paths from the source s to all candidate points. Equivalently, $\mathcal{T}$ is the shortest path tree on $\mathcal{S}_\varepsilon$ within $\overline{P}$ (the simple polygon obtained by filling all holes of P) and can be computed in $O(|\mathcal{S}_\varepsilon|)$ time [19]. The path from s to a candidate point s' in $\mathcal{T}$ is the pseudo-geodesic shortest path $\pi(s, s')$.

We define a *geodesic angular order* around s as follows: for two points s_i and s_j, we say s_i precedes s_j if s_i lies to the left of $\pi(s, s_j)$ with its final segment extended until $\partial_{\text{outer}}P$. Ties are broken by increasing distance from s (farther points come later). For each reflex vertex r_i of P, we introduce an additional candidate point s_{r_i} to account for the possibility that r_i may appear twice in a pseudo-relatively convex polygonal chain. (For example, in Figure 6, tracing out $\pi(x, y)$ twice, once in each direction, yields a relatively convex polygonal tour in which some reflex vertices of P may occur twice.) Each s_{r_i} respects the geodesic angular order and is inserted after all candidate points in the subtree of $\mathcal{T}$ rooted at r_i. After sorting all candidate points accordingly, we append $\overline{s} \equiv s$ to the end of the list to ensure that the chain returns to its starting point. Let $\{s, s_1, \ldots, \overline{s}\}$ denote the sorted list of candidate turn points. A pseudo-relatively convex tour always has vertices in increasing order through this list (it suffices to assume clockwise orientation), thus this ordering is sufficient to establish a topological ordering over the subproblems. See Figure 8.

A subproblem (s_i, k) is specified by a candidate turn point and an integer $0 \leq k \leq n$, $DP(s_i, k)$ is the length of the shortest pseudo-relatively convex chain (a chain whose vertices are in pseudo-relatively convex position and ordered around the boundary of their pseudo-relatively convex hull) intersecting at least k essential cuts. By Corollary 8, the recursion is as follows:

$$DP(s_i, k) = \min_{\substack{i' < i \\ s_i \text{ sees } s_{i'}}} \{DP(s_{i'}, k - |C(s_i s_{i'}) \setminus C(\pi(s, s_i))|) + |\pi_P(s_{i'}, s_i)|\} .$$

The base case is $DP(s, 0) = 0$. The optimal grid-rounded solution to the Minimum Covering Pseudo-Relatively Convex Polygon problem is given by the subproblem $(\overline{s}, \overline{k})$, where $\overline{k} = O(n)$ is the number of essential cuts.

The recursion always returns a pseudo-relatively convex chain; otherwise, the point of tangency from s_i to the chain from s to $s_{i'}$ would yield a strictly better solution (since one could shortcut from s_i to the tangency point while intersecting the same essential cuts). This observation, together with the optimal substructure in Corollary 8, complete the proof of correctness for our algorithm.

The number of subproblems is $O\left(\left(\frac{n^2}{\varepsilon}\right)^2\right) O(n)$, to solve each subproblem we iterate through $O\left(\left(\frac{n^2}{\varepsilon}\right)^2\right)$ candidate turn points, thus the algorithm runs in time $O\left(\frac{n^9}{\varepsilon^4}\right)$, which is fully polynomial in n and $1/\varepsilon$. ◀

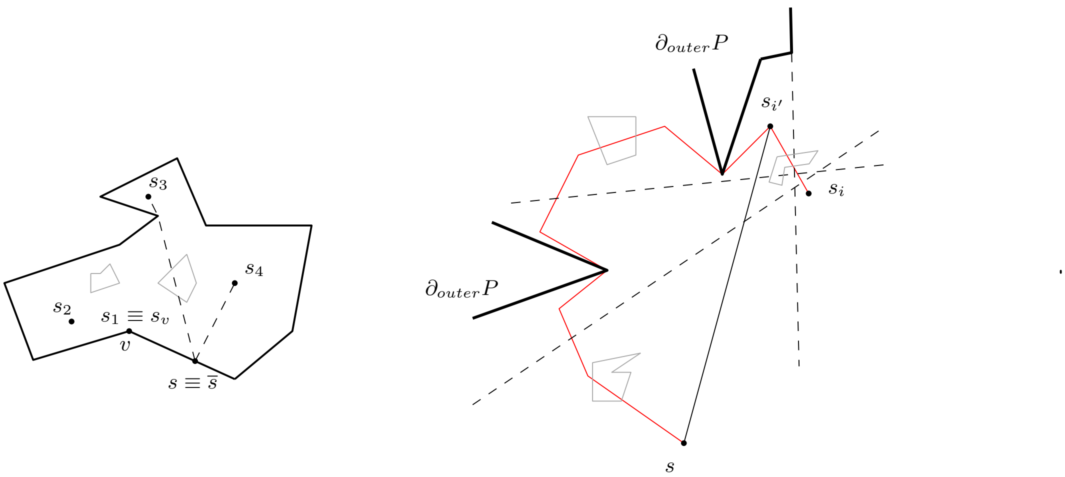

Figure 8 Left: the list of points in sorted geodesic angular order is $\{s, s_1, s_2, s_v, s_3, s_4\}$. Right: an example subproblem, $DP(s_i, 3) = DP(s_{i'}, 2) + |\pi_P(s_{i'}, s_i)|$.

4 The Quota-TSP on Infinite Lines in the Plane with Obstacles

4.1 Preliminaries

Given a set $\{l_1, l_2, \ldots, l_m\}$ of m infinite lines in the Euclidean plane $\mathbb{R}^2$, each associated with a non-negative reward $\mathcal{R}(l_i)$ ($\mathcal{R}(\cdot)$ also denotes the total reward of all the lines when the argument is a set of lines), h polygonal obstacles $O_1, O_2, \ldots, O_h$, and a quota Q, the objective is to compute a shortest obstacle-avoiding route that intersects a subset of the lines whose total reward is at least Q. Let $n = O\left(m + \sum_{i=1}^{h} |\text{vert}(O_i)|\right)$ denote the total input size, where $|\text{vert}(O_i)|$ is the number of vertices of obstacle O_i.

Many definitions and techniques carry over from the previous section on the Watchman Route Problem with transparent obstacles. We begin by recalling some relevant definitions. For any two points x, y in the free space $\mathbb{R}^2 \setminus \bigcup_{i=1}^{h} \text{int}(O_i)$ (where $\text{int}(O_i)$ is the interior of obstacle O_i), let $\pi(x, y)$ denote the geodesic shortest path between x and y that avoids obstacles. A set X is said to be *geodesically convex* if for every $x, y \in X$, we have $\pi(x, y) \subseteq X$. We use xy to denote the straight-line segment connecting x and y, which may intersect obstacles. Let γ be an optimal tour. Since it must avoid obstacles while intersecting the required lines, γ is geodesically convex and, by Lemma 1, has at most $O(n)$ vertices.

4.2 Flipping the objective function–maximizing the reward subject to a maximum length

Let γ be an optimal solution. Denote by $\mathcal{L}(\gamma)$ the set of lines intersected by γ. Note that the minimum enclosing circle of γ also intersects all these lines (and possibly more). Within this circle, there exists a locally minimum circle that intersects a set of lines whose total reward meets or exceeds the quota.

By critical placement arguments similar to those in Appendix B, there are only a polynomial number of such critical circles (i.e., locally minimum circles that intersect a given subset of lines). We can afford to compute each one and discard those that intersect lines whose total reward does not meet the quota. Let $\mathcal{C}$ be one such critical circle that intersects a sufficient-reward set of lines. Let r be its radius. We have $r \leq |\gamma| = O(nr)$, and γ lies within localizing square B of side length $O(nr)$ concentric with $\mathcal{C}$.

We now discretize the free space within B, i.e., the region $B \setminus \bigcup_{i=1}^{h} \text{int}(O_i)$. First, we triangulate this region. Then, we overlay a regular square grid over B, with pixel side length $\delta = \frac{\varepsilon \cdot \text{radius}(\mathcal{C})}{12n}$. Each convex cell in the resulting decomposition has perimeter at most 4δ. The total number of vertices in this decomposition is $O\left(\left(\frac{n^2}{\varepsilon}\right)^2\right)$; let $\mathcal{S}_\varepsilon$ denote this set of vertices. For purposes of computing a $(1 + \varepsilon)$-approximation, it suffices to consider only turn points in $\mathcal{S}_\varepsilon$ (by Lemma 3).

Our algorithm is based on a key geometric observation: a connected set intersects a line if and only if its convex hull intersects the line. We introduce an intermediary problem called Reward Collecting Convex Polygon problem, for which we propose a dynamic programming algorithm to solve a discrete version. We show how it can be used to approximately solve our Quota-TSPN variant (similar to how we use a pseudo-relatively convex polygon as a proxy for a see-through watchman route in the previous section).

In the Reward Collecting Convex Polygon problem, the goal is to compute n' points $p_1, p_2, \ldots, p_{n'}$ in convex position, so that the total reward of the lines intersecting the convex polygon $\mathcal{P} = (p_1, p_2, \ldots, p_{n'})$ is maximized, subject to a constraint that the geodesic perimeter $|\pi(p_1, p_2)| + \cdots + |\pi(p_{n'}, p_1)|$ is at most a given budget L. Since the polygon may intersect obstacles, the perimeter is measured in the geodesic metric in the free space, not the Euclidean metric.

Let γ be the optimal tour for the Quota-TSP on infinite lines, and let $L = |\gamma|$. We claim that an optimal solution to the discrete Reward Collecting Convex Polygon with budget L yields an optimal solution to the discrete Quota-TSP. Let $\mathcal{L}(\gamma)$ be the set of lines intersected by γ, and consider the convex hull $\text{CH}(\gamma)$. Since $\text{CH}(\gamma)$ intersects all the lines that γ does, $\mathcal{R}(\mathcal{L}(\text{CH}(\gamma)))$, the total reward of all lines intersected by $\text{CH}(\gamma)$, is no smaller than Q. Moreover, $\text{CH}(\gamma)$ is a convex polygon whose geodesic perimeter is no greater than $|\gamma| = L$. Thus, we have shown that there exists a convex polygon of (geodesic) perimeter at most L that collects reward at least Q, so the optimal solution $\mathcal{P} = (p_1, p_2, \ldots, p_{n'})$ to the

Reward Collecting Convex Polygon problem with budget L must also collect total reward at least Q. We concatenate the geodesic paths $\pi(p_1,p_2),\ldots,\pi(p_{n'},p_1)$ into a tour. Clearly $\mathcal{R}(\mathcal{L}(\pi(p_1,p_2)\cup\ldots\cup\pi(p_{n'},p_1)))\geq Q$ and $|\pi(p_1,p_2)|+\ldots+|\pi(p_{n'},p_1)|\leq L$. Since no tour strictly shorter than L can satisfy the quota (by the optimality of γ), it follows that $|\pi(p_1,p_2)|+\ldots+|\pi(p_{n'},p_1)|=L$ and $\pi(p_1,p_2)\cup\ldots\cup\pi(p_{n'},p_1)$ is also optimal. Similarly, one can show that an optimal Quota-TSP route γ corresponds to an optimal reward collecting convex polygon, justifying using $\mathcal{S}_\varepsilon$ as candidate turn points for solving (approximately) the Reward Collecting Convex Polygon problem (see Lemma 5).

▶ **Theorem 10.** *The Quota-TSP on infinite lines with obstacles in $\mathbb{R}^2$ has a fully polynomial-time approximation scheme.*

Proof. Let γ denote the optimal solution to the Quota-TSP on infinite lines, and γ' denote the optimal tour in the discrete version of the problem with vertices in $\mathcal{S}_\varepsilon$.

We want to find an optimal solution solution $\mathcal{P}$ to the Reward Collecting Convex Polygon problem with budget $L:=|\gamma'|$ whose vertices are in $\mathcal{S}_\varepsilon$. While we are not able to accurately determine $|\gamma'|$ easily, "guessing" L to be within a factor of $(1+\varepsilon)$ to $|\gamma'|$ is feasible. To do so, we divide the side length of the localizing square B, which is $O(n)|\gamma|$ and thus also $O(n)|\gamma'|$, into intervals of length εr, resulting in $O\left(\frac{n}{\varepsilon}\right)$ interval endpoints. The interval containing $|\gamma|$ has a right endpoint that is within a factor $(1+\varepsilon)$ to $|\gamma'|$. Our algorithm performs a binary search over these endpoints: we find the first L for which a convex polygon $\mathcal{P}$ of geodesic perimeter at most L satisfying the quota Q exists. This guarantees that $|\gamma'|\leq L\leq(1+\varepsilon)|\gamma'|\leq(1+\varepsilon)^2|\gamma|$.

For each value of L, we try all possible choices for the vertex with minimum y-coordinate in $\overline{\gamma}$ (of which there are $O\left(\left(\frac{n^2}{\varepsilon}\right)^2\right)$); call this point p_1. We then solve the discrete Reward Collecting Convex Polygon problem using dynamic programming. First, we preprocess $\mathcal{S}_\varepsilon$ by discarding all points with lower y-coordinate than p_1, then sort the remaining points in clockwise angular order around p_1, breaking ties by increasing distance (i.e., farther points come later). We also add a dummy point $\overline{p}\equiv p_1$ to close the polygon. Let $\{p_1,s_1,s_2,\ldots,\overline{p}\}$ be the sorted list of candidate points. A convex chain consists of vertices appearing in increasing order in this list (assuming clockwise orientation). Without loss of generality, we consider only the clockwise case.

Each subproblem is defined by two points $s_i,s_j\in\mathcal{S}_\varepsilon$ with $i<j$, and a budget value L'. The value $DP(s_i,s_j,L')$ denotes the maximum reward collected by a convex polygonal chain (a polygonal chain with vertices in convex position) from p_1 to s_j, ending with segment s_is_j, using total length at most L'. Let $\mathcal{L}(ab)$ denote the set of lines intersecting the segment ab. The recursion works as follows: for $i'<i$, where $s_{i'}$ lies to the right (i.e., same side as p) of the supporting line of s_is_j (oriented from s_i to s_j),

$$DP(s_i,s_j,L')=\max_{i'<i}\left\{DP(s_{i'},s_i,L'-|\pi(s_i,s_j)|)+\mathcal{R}(\mathcal{L}(s_is_j)\setminus\mathcal{L}(p_1s_i))\right\}.$$

See Figure 9. To maintain convexity in the chain associated with subproblem (s_i,s_j,L'), we recurse only on subproblems where $s_{i'}$ lies to the right of s_is_j. Recall from Corollary 8, due to convexity, the lines newly intersected by segment s_is_j are always exactly $\mathcal{L}(s_is_j)\setminus\mathcal{L}(p_1s_i)$ and do not depend on the specific chain used in previous subproblems. Thus, the chain from p_1 to s_j of total length L' collects the most reward for its length if and only if the chain from p_1 to s_i of length $L'-|\pi(s_i,s_j)|$ also collects the most reward for its length: this proves correctness and is the optimal substructure of our dynamic programming algorithm.

The base cases $DP(p_1,s,|\pi(p_1,s)|)$ for all candidate turn points s can be computed by brute force. The overall optimum is given by $\max_s\{DP(s,\overline{p},L)\}$.

However, the values $|\pi(s_i, s_j)|$ may be irrational, making exact tabulation of subproblems infeasible. To remedy this, we apply a standard "bucketing" and rounding technique: we discretize the budget into intervals, and round each segment length $|\pi(s_i, s_j)|$ up to the nearest bucket boundary. Let the width of each bucket be I. Since $\overline{\gamma}$ has at most n vertices, then the total extra length incurred due to rounding is at most nI. To ensure that this does not increase the total length by more than εL, we set $I = \frac{L}{\lceil \frac{n}{\varepsilon} \rceil}$. Thus, we need only consider the following $O\left(\lceil \frac{n}{\varepsilon} \rceil\right)$ multiples of I: $\left\{0, \frac{L}{\lceil \frac{n}{\varepsilon} \rceil}, \frac{2L}{\lceil \frac{n}{\varepsilon} \rceil}, \dots, (n + \lceil \frac{n}{\varepsilon} \rceil)\frac{L}{\lceil \frac{n}{\varepsilon} \rceil}\right\}$.

The recursion with rounded lengths is:

$$DP(s_i, s_j, L') = \max_{i'<i} \left\{DP(s_{i'}, s_i, L' - \lceil|\pi(s_i, s_j)|\rceil_I) + r(\mathcal{L}(s_i s_j) \setminus \mathcal{L}(p_1 s_i))\right\},$$

where $\lceil x \rceil_I$ denotes rounding x up to the nearest multiple of I. The optimal solution for the rounded instance is given by $DP(\overline{p}, L + nI)$, lastly we simply concatenate the geodesic paths between its vertices.

The number of subproblems is $O\left(\frac{n^4}{\varepsilon^2}\right) O\left(\frac{n^4}{\varepsilon^2}\right) O\left(\lceil \frac{n}{\varepsilon} \rceil\right) = O\left(\frac{n^9}{\varepsilon^5}\right)$. It takes $O\left(\frac{n^4}{\varepsilon^2}\right)$ to solve each subproblem, and thus the total running time is $O\left(\frac{n^{13}}{\varepsilon^7}\right)$. We obtain a tour of length at most $(1+\varepsilon)L \leq (1+\varepsilon)^2|\gamma'| \leq (1+\varepsilon)^3|\gamma| = (1+O(\varepsilon))|\gamma|$. ◀

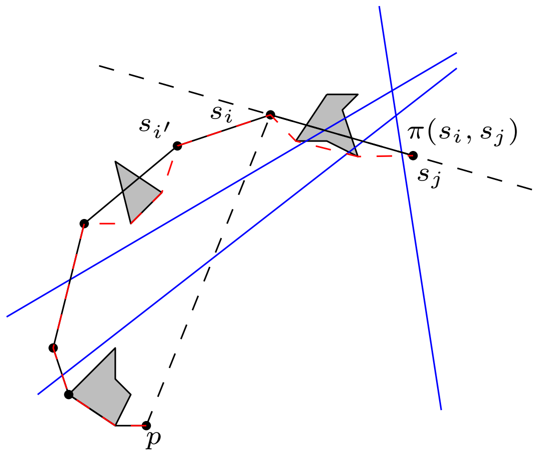

Figure 9 Subproblem (s_i, s_j, L').

5 The Watchman Route Problem with transparent walls

▶ **Theorem 11.** *The STWRP in a simple polygon with see-through boundary segments cannot be approximated to within a factor better than $O(\log n)$ in polynomial time, unless* $\mathrm{P} = \mathrm{NP}$.

Proof. Our proof is based on a reduction from the classical SET COVER problem, which has shown to be NP-hard to approximate within a factor of $c \log n$ for some $c > 0$ [1]. Given a universe $\mathcal{U} = \{x_1, x_2, \dots, x_n\}$ and a collection of subsets $\mathcal{S} = \{S_1, S_2, \dots, S_m\}$, the objective is to select the smallest number of subsets whose union covers all of $\mathcal{U}$.

We construct a simple polygon that encodes this instance as follows (see Figure 10, not drawn to scale). The polygon has a horizontal base from which $m + 1$ narrow vertical corridors rise. One of them – the leftmost – has length $L + 1$ and leads to a long horizontal

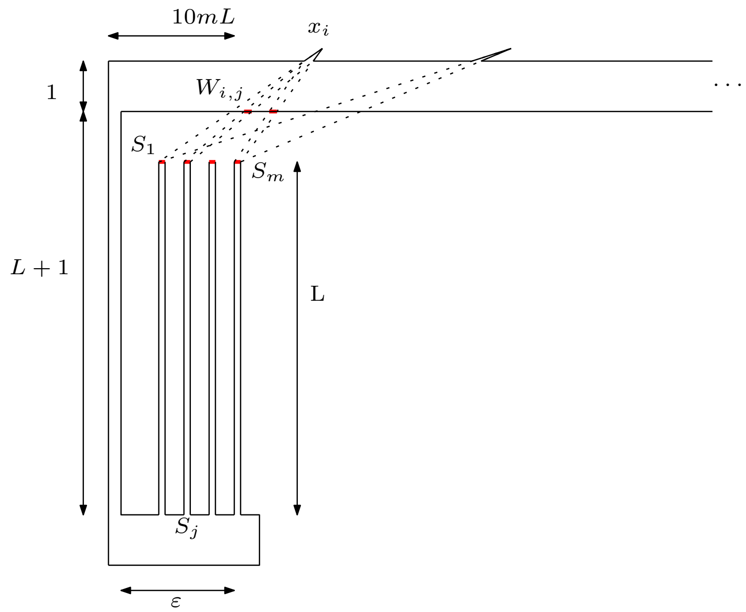

Figure 10 Reduction from SET COVER in the proof of Theorem 11.

corridor at the top. The remaining m corridors correspond to the subsets $S_1, \ldots, S_m$ and each has length L. Every vertical corridors has a transparent (see-through) wall at their top end. The corridors are spaced evenly across the base; the length of the base is $\varepsilon \ll L$. We also use S_i to denote the corridor that corresponds to S_i.

Along the top horizontal corridor, we place thin, pointed notches, each corresponding to an element $x_i \in \mathcal{U}$. Each notch has a narrow base and a sharp tip positioned at the top wall of the corridor. The placement of a notch and its two adjacent vertices, left and right, is such that the tip vertex, the left adjacent vertex and the top, left endpoint of the corridor S_1 (which is left-most of the m subset corridors) are collinear; and so are the tip, the right adjacent vertices, and the top right endpoint S_m (which is right-most).

The *feasible glass region for* x_i is the portion of the bottom wall of the top corridor through which the tip of x_i can be seen from the subset corridors, assuming the presence of some transparent boundary segments. Specifically, it is the interval between the intersection points of the bottom wall with the lines connecting the tip of x_i to the top left vertex of S_1 and to the top right vertex of S_m. By Thales' theorem, the width of this region is approximately $\varepsilon/2$. For every j, if $x_i \in S_j$ there is a small transparent window $W_{i,j}$ in the feasible glass region for x_i, that allows the to see x_i from S_j. These windows can be pairwise disjoint and arbitrarily small by making the corridors arbitrarily narrow (again via Thales' theorem), ensuring that x_i is visible only from corridors corresponding to subsets that contain it (or the top corridor, but our construction will make it suboptimal to traverse the top corridor).

The first notch x_1 is placed so that the intersection of its feasible glass region lies at distance $10mL$ from the leftmost vertical corridor. The remaining notches are positioned so that their feasible glass regions are disjoint: each is of a distance $2mL$ from the previous.

To see the upper corridor and the notches, the watchman must either traverse the entire top corridor (which is costly), or just visit the top via the left corridor and selectively ascend some of the subset corridors and see the notches through the appropriate windows. Even if the watchman ascends all m subset corridors and the top corridor, the total route length is at most $2(L + 1 + mL + 2m\varepsilon) < 10mL$, for sufficiently small ε. Therefore, any approximately-optimal watchman route will avoid traversing the top corridor altogether and instead ascend only a selected subset of the m vertical corridors. This selection corresponds exactly to a set cover. If the original Set Cover instance admits a cover of size k, then the optimal watchman route has length $O(kL)$. Hence, a polynomial-time approximation algorithm for the watchman route problem with a factor better than $O(\log n)$ would yield an approximation factor of the same order (up to a constant) for the Set Cover problem, contradicting known approximation hardness results unless P = NP. ◀

References

1. Noga Alon, Dana Moshkovitz, and Shmuel Safra. Algorithmic construction of sets for k-restrictions. *ACM Transactions on Algorithms (TALG)*, 2(2):153–177, 2006. doi:10.1145/1150334.1150336.
2. Antonios Antoniadis, Sándor Kisfaludi-Bak, Bundit Laekhanukit, and Daniel Vaz. On the Approximability of the Traveling Salesman Problem with Line Neighborhoods. In *18th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2022)*, volume 227 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:21, 2022. doi:10.4230/LIPIcs.SWAT.2022.10.
3. Esther M. Arkin and Refael Hassin. Approximation algorithms for the geometric covering salesman problem. *Discrete Applied Mathematics*, 55(3):197–218, 1994. doi:10.1016/0166-218X(94)90008-6.
4. Sanjeev Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998. doi:10.1145/290179.290180.
5. Luis Barba, Michael Hoffmann, Matias Korman, and Alexander Pilz. Convex hulls in polygonal domains. In *Proceedings of the 16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018)*, pages 8:1–8:13, 2018. doi:10.4230/LIPIcs.SWAT.2018.8.
6. Anna Brötzner, Bengt J. Nilsson, and Christiane Schmidt. The k-transmitter watchman route problem is NP-hard even in histograms and star-shaped polygons. In *40th European Workshop on Computational Geometry, Ioannina, Greece, 2024*, volume 40, pages 381–387, 2024.
7. Frantisek Burian, Petra Kocmanova, and Ludek Zalud. Robot mapping with range camera, CCD cameras and thermal imagers. In *2014 19th International Conference on Methods and Models in Automation and Robotics (MMAR)*, pages 200–205. IEEE, 2014. doi:10.1109/MMAR.2014.6957351.
8. Svante Carlsson, Håkan Jonsson, and Bengt J. Nilsson. Finding the shortest watchman route in a simple polygon. *Discrete & Computational Geometry*, 22(3):377–402, 1999. doi:10.1007/PL00009467.
9. Lidice E. Castro Jimenez and Edgar A. Martínez-García. Thermal image sensing model for robotic planning and search. *Sensors*, 16(8):1253, 2016. doi:10.3390/S16081253.
10. Wei-Pang Chin and Simeon Ntafos. Optimum watchman routes. In *Proceedings of the 2nd Annual Symposium on Computational Geometry*, pages 24–33, 1986. doi:10.1145/10515.10518.
11. Wei-Pang Chin and Simeon Ntafos. Shortest watchman routes in simple polygons. *Discrete & Computational Geometry*, 6(1):9–31, 1991. doi:10.1007/BF02574671.
12. Mark De Berg, Joachim Gudmundsson, Matthew J. Katz, Christos Levcopoulos, Mark H. Overmars, and A. Frank Van Der Stappen. TSP with neighborhoods of varying size. *Journal of Algorithms*, 57(1):22–36, 2005. doi:10.1016/J.JALGOR.2005.01.010.
13. Moshe Dror, Alon Efrat, Anna Lubiw, and Joseph S. B. Mitchell. Touring a sequence of polygons. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 473–482, 2003. doi:10.1145/780542.780612.

14 Adrian Dumitrescu and Joseph S. B. Mitchell. Approximation algorithms for TSP with neighborhoods in the plane. *Journal of Algorithms*, 48(1):135–159, 2003. doi:10.1016/S0196-6774(03)00047-6.

15 Adrian Dumitrescu, Joseph S. B. Mitchell, and Paweł Żyliński. Watchman routes for lines and line segments. *Computational Geometry*, 47(4):527–538, 2014. doi:10.1016/J.COMGEO.2013.11.008.

16 Adrian Dumitrescu and Csaba D. Tóth. Watchman tours for polygons with holes. *Computational Geometry*, 45(7):326–333, 2012. doi:10.1016/J.COMGEO.2012.02.001.

17 Adrian Dumitrescu and Csaba D. Tóth. The traveling salesman problem for lines, balls, and planes. *ACM Transactions on Algorithms (TALG)*, 12(3):1–29, 2016. doi:10.1145/2850418.

18 Joachim Gudmundsson and Christos Levcopoulos. A fast approximation algorithm for TSP with neighborhoods. *Nordic Journal of Computing*, 6(4):469–488, 1999.

19 Leonidas Guibas, John Hershberger, Daniel Leven, Micha Sharir, and Robert Tarjan. Linear time algorithms for visibility and shortest path problems inside simple polygons. In *Proceedings of the 2nd Annual Symposium on Computational Geometry*, pages 1–13, 1986.

20 Kien C. Huynh, Joseph S. B. Mitchell, Linh Nguyen, and Valentin Polishchuk. Optimizing visibility-based search in polygonal domains. In *Proceedings of the 19th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2024)*, pages 27:1–27:16, 2024. doi:10.4230/LIPIcs.SWAT.2024.27.

21 Cristian S. Mata and Joseph S. B. Mitchell. Approximation algorithms for geometric tour and network design problems. In *Proceedings of the 11th Annual Symposium on Computational Geometry*, pages 360–369, 1995.

22 Jiří Matoušek, Micha Sharir, and Emo Welzl. A subexponential bound for linear programming. In *Proceedings of the 8th Annual Symposium on Computational Geometry*, pages 1–8, 1992.

23 Joseph S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k-MST, and related problems. *SIAM Journal on computing*, 28(4):1298–1309, 1999. doi:10.1137/S0097539796309764.

24 Joseph S. B. Mitchell. A PTAS for TSP with neighborhoods among fat regions in the plane. In *Proceedings of the 8th Annual Symposium on Computational Geometry*, pages 11–18, 2007. URL: http://dl.acm.org/citation.cfm?id=1283383.1283385.

25 Joseph S. B. Mitchell. Approximating watchman routes. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 844–855. SIAM, 2013. doi:10.1137/1.9781611973105.60.

26 Joseph S. B. Mitchell. Shortest paths and networks. In *Handbook of discrete and computational geometry*, pages 811–848. Chapman and Hall/CRC, 2017.

27 Bengt J. Nilsson. Approximating a shortest watchman route. *Fundamenta Informaticae*, 45(3):235–281, 2001.

28 Bengt J. Nilsson and Christiane Schmidt. k-transmitter watchman routes. In *International Conference and Workshops on Algorithms and Computation*, pages 202–213. Springer, 2023. doi:10.1007/978-3-031-27051-2_18.

29 Joseph O'Rourke. Visibility. In *Handbook of Discrete and Computational Geometry*, pages 875–896. Chapman and Hall/CRC, 2017.

30 Nikola Petrovic, Tommy Gunnarsson, Nadine Joachimowicz, and Magnus Otterskog. Robot controlled data acquisition system for microwave imaging. In *2009 3rd European Conference on Antennas and Propagation*, pages 3356–3360. IEEE, 2009.

31 Shmuel Safra and Oded Schwartz. On the complexity of approximating TSP with neighborhoods and related problems. *computational complexity*, 14:281–307, 2006. doi:10.1007/S00037-005-0200-3.

32 Xuehou Tan. Fast computation of shortest watchman routes in simple polygons. *Information Processing Letters*, 77(1):27–33, 2001. doi:10.1016/S0020-0190(00)00146-0.

33 Xuehou Tan. Approximation algorithms for the watchman route and zookeeper's problems. *Discrete Applied Mathematics*, 136(2-3):363–376, 2004. doi:10.1016/S0166-218X(03)00451-7.

34 Xuehou Tan. A linear-time 2-approximation algorithm for the watchman route problem for simple polygons. *Theoretical Computer Science*, 384(1):92–103, 2007. doi:10.1016/J.TCS.2007.05.021.

35 Xuehou Tan, Tomio Hirata, and Yasuyoshi Inagaki. An incremental algorithm for constructing shortest watchman routes. *International Journal of Computational Geometry & Applications*, 3(04):351–365, 1993. doi:10.1142/S0218195993000233.

36 Xuehou Tan, Tomio Hirata, and Yasuyoshi Inagaki. Corrigendum to "an incremental algorithm for constructing shortest watchman routes". *International Journal of Computational Geometry & Applications*, 9(03):319–323, 1999. doi:10.1142/S0218195999000212.

37 Xuehou Tan and Bo Jiang. The touring polygons problem revisited. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, 101(5):772–777, 2018. doi:10.1587/TRANSFUN.E101.A.772.

A Proofs of some Lemmas

A.1 Lemma 1

▶ **Lemma 1.** *The optimal see-through watchman route γ is geodesically convex and has $O(n)$ vertices that lie in the interior of P.*

Proof. It is known [10, 11, 13, 35] that in a simple, hole-free polygon, an optimal watchman route is a polygonal chain with $O(n)$ vertices that visits all essential cuts (possibly weakly simple, as the watchman may traverse a segment in both directions). In a polygonal domain with entirely transparent holes, we show that the same property holds. Consider $\overline{P}$, the simple polygon whose boundary is the outer boundary of P. To see all of $\overline{P}$, a route γ needs to visit all essential cuts of $\overline{P}$ (which are also essential cuts of P). From every point in $\overline{P}$, we can draw a line of sight to γ that stays within $\overline{P}$, meaning from that same point in P we can draw the same line to γ that only crosses the transparent obstacles. Furthermore, this implies that γ corresponds to an optimal solution to an instance of the TSPN with obstacles, where the regions to be visited are the essential cuts, and thus γ must be polygonal.

Now suppose that γ is not geodesically convex. Let Γ denote the outer boundary of the geodesic convex hull GCH(γ). Then $|\Gamma| \leq |\gamma|$ [25, 26]. Consider any point $p \in P$ that lies outside of Γ. Since γ sees all of P, there exists some $x \in \gamma$ such that the segment px lies entirely within the visibility region $V(x)$. By the Jordan Curve Theorem, the segment px must intersect Γ, implying that Γ also sees p. Furthermore, since all holes are transparent, Γ sees all points in the interior of P that lie within it. Therefore, Γ sees all of P, and since $|\Gamma| \leq |\gamma|$, this contradicts the optimality of γ unless γ is itself geodesically convex.

Without loss of generality, suppose γ has at least three vertices (special cases include when γ has one vertex, i.e., s sees all of P, or when γ has two vertices, i.e., it is a weakly simple tour that traces a segment in opposite directions). Since γ is geodesically convex, its vertices (aside from s) fall into two categories: (i) vertices of P (either convex or reflex), and (ii) convex vertices that lie strictly in the interior of P. These interior convex vertices must correspond to reflection points off essential cuts [11, 13, 25].

To see this, suppose there is a convex vertex v of γ in the interior of P that is not a reflection off an essential cut, and let vv_1 and vv_2 be the two segments of γ incident to v. There are three cases to consider:

(i) If both vv_1 and vv_2 intersect essential cuts at points other than v (including the case when vv_1 or vv_2 overlaps a cut), then we can shortcut γ by connecting an intersection point on vv_1 to one on vv_2.

(ii) If either (or both) of vv_1 or vv_2 does not intersect any essential cut, then we can shortcut by connecting v_1 and v_2 directly with a segment, without missing any essential cut.

(iii) If v is contained in an essential cut, but it is not a reflection point (v_1 and v_2 are on different sides of the cut), we can also shortcut by connecting the segment v_1v_2 while still intersecting the essential cut.

In all cases, we obtain a shorter route that still intersects all essential cuts, contradicting the optimality of γ. (See Figure 11.)

Moreover, no two interior convex vertices can correspond to reflections off the same essential cut: if a vertex v of γ is reflection off of a cut c that has already been visited before, simply shortcut the two vertices adjacent to v. Note that not all essential cuts must induce a reflection: γ may follow a contiguous portion of an essential cut. However, every interior convex vertex must be a reflection off a unique essential cut. Thus, the number of convex vertices of γ that lie in the interior of P is no greater than the number of essential cuts, which is at most n. ◀

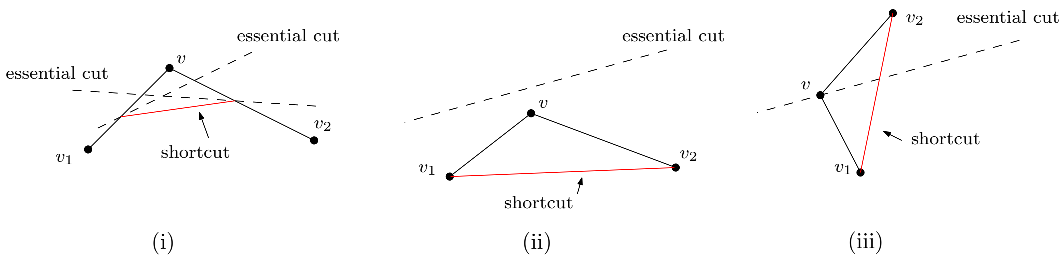

Figure 11 A convex vertex of γ interior to P must be a reflection off an essential cut, else we can shortcut γ.

A.2 Lemma 2

▶ **Lemma 2.** $r \leq |\gamma| \leq O(nr)$

Proof. For the first inequality, note that any tour of length r must have radius at most $r/2$, and thus cannot see all of P. Since γ is required to see all of P, we must have $|\gamma| \geq r$.

For the second inequality, consider the outer boundary of GD_s, denoted $\partial_{outer}GD_s$. This boundary consists of $O(n)$ line segments and $O(n)$ circular arcs [20]. By the triangle inequality, each line segment has length at most the sum of the geodesic distances from its endpoints to s, i.e., at most $2r$. Each circular arc lies on a geodesic circle of radius r, so its length is at most $2\pi r$. Therefore, $|\partial_{outer}GD_s| = O(nr) + O(nr) = O(nr)$.

Finally since $\partial_{outer}GD_s$ sees all of P, we have $|\gamma| \leq |\partial_{outer}GD_s|$. ◀

A.3 Lemma 4

▶ **Lemma 4.** $V(\gamma) = V(PCH(\gamma))$. *Generally, the pseudo-visibility region of any connected set is equal to that of its pseudo-relative convex hull.*

Proof. Since $PCH(\gamma)$ encloses γ, clearly $V(\gamma) \subseteq V(PCH(\gamma))$. On the other hand, suppose there is a point p that is seen by $PCH(\gamma)$ but not by γ. This implies that there exists a reflex vertex v of $\partial_{outer}P$ (since only a reflex vertex can obstruct visibility, given that all obstacles are transparent) that blocks vision to p from γ. Specifically, consider extending a window from v, collinear with p and away from p, maximally until the window reaches the face at infinity bounded by $\partial_{outer}P$, then γ is within the subpolygon that does not contain p (see Figure 12). On the other hand, $PCH(\gamma)$ sees p and must intersect with the window. But then we can "cut off" a part of $PCH(\gamma)$ with a line (say, parallel and close to the window), obtaining a smaller relatively convex set containing γ, which contradicts the minimality of $PCH(\gamma)$. Thus, for every point p, if p is seen by $PCH(\gamma)$ then p is seen by γ. ◀

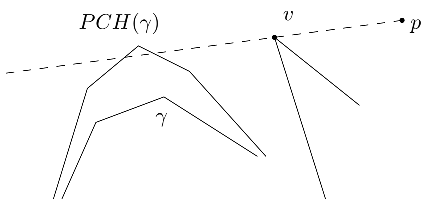

Figure 12 Proof of Lemma 4.

A.4 Lemma 6

▶ **Lemma 6.** *For any $x, y \in P$, if the pseudo-geodesic shortest path $\pi(x, y)$ intersects a chord c of $\overline{P}$ (the simple polygon obtained by filling all holes of P), including extensions of edges incident to reflex vertices and essential cuts, then every path from x to y – including pseudo-paths passing through holes – must intersect c.*

Proof. See Figure 6. If c intersects x or y directly, then clearly every path from x to y must intersect c. Otherwise, suppose c intersects $\pi(x, y)$ at an interior point. Since $\pi(x, y)$ is the shortest path in $\overline{P}$, this implies that c separates x and y into different subpolygons of $\overline{P}$. Therefore, by the Jordan Curve Theorem, every path from x to y must intersect c. ◀

B The floating STWRP

Different locations of s yield different essential-cut sets, but only $O(n^2)$ combinations arise: extending all edges incident to reflex vertices decomposes P into convex cells (Figure 13), and each face (2-, 1-, or 0-dimensional) of this arrangement corresponds to one combination of essential cuts.

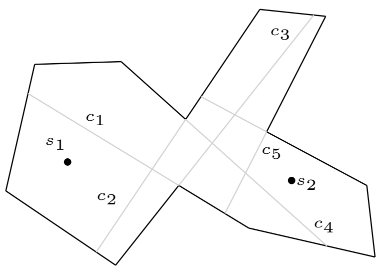

Figure 13 If the watchman starts from s_1 or any point in the same cell, c_1 and c_2 are the essential cuts an optimal route needs to visit. If the watchman starts from s_2 or any point in the same cell, an optimal route needs to visit c_3, c_4.

For the floating case, our algorithm computes a polynomial number of localizing squares, discretizes each into candidate turn points, and tries each as a potential start; preprocessing and DP then proceed as in Theorem 9 with the corresponding essential cuts. We now show how to find the localizing squares.

Let mec(γ) be the minimum enclosing circle of an optimal route γ. Since γ sees all of P, so does $\mathcal{R}$, the connected component of mec(γ) $\cap P$ enclosing γ. Hence there is a *locally minimum illuminating circle* $\mathcal{C} \subseteq$ mec(γ): a circle with a connected component $\mathcal{C} \cap P \subseteq \mathcal{R}$ that sees all of P, whose radius is locally minimal among such circles (mec(γ) itself qualifies but need not be locally minimum). A circle has three degrees of freedom (translation and scale), so a standard critical-placement argument lets us shrink mec(γ), while preserving the illuminating property, until it is *pinned* by three (not necessarily distinct) edge extensions, by tangency or at an endpoint. (If instead $\mathcal{C}$ shrinks to a point, then $\overline{P}$ – the simple polygon bounded by $\partial_{outer}P$ – is star-shaped with kernel meeting P; we compute this kernel in $O(n)$ time [22] and pick a point of it in P.) This yields $\mathcal{C}$ with three critical contacts c_1, c_2, c_3, lying in the same component of $\mathcal{C} \cap P$.

- **Case 1:** c_1, c_2, and c_3 are pairwise distinct. There are only a constant number of combinatorial types of triple contacts (e.g., c_1 and c_2 are contacted by tangency, c_3 is contacted at an endpoint, etc.). For each triple of essential cuts, we can afford to compute all locally minimum illuminating circles, one for each combinatorial type of contacts.
- **Case 2:** Two of the contacts are the same, say $c_1 = c_2 \neq c_3$. Then $\mathcal{C}$ is initially pinned by only two distinct edge extensions, c_1 and c_3. We analyze two subcases based on the nature of these contacts:
  - *Subcase A: At least one contact is at an endpoint.* Suppose, for example, that c_1 is contacted at an endpoint, say p, and c_3 is contacted by tangency. Then the center of $\mathcal{C}$ lies along the line passing through p perpendicular to the line supporting c_3. Any attempt to shift the center to reduce the radius would cause $\mathcal{C}$ to lose contact with c_1 or c_3, contradicting local minimality – unless a new critical contact arises, in which case we are back in Case 1. The argument is identical when both contacts are endpoints.
  - *Subcase B: Both contacts are by tangency.*
    - If c_1 and c_3 are *not parallel*, then their supporting lines intersect at a point (possibly outside P). The center of $\mathcal{C}$ lies along the angle bisector of the two lines. We are assuming $\mathcal{C}$ is pinned only by c_1 and c_3, thus sliding its center does not affect intersection with any other essential cut. If we continuously slide the center of $\mathcal{C}$ toward the intersection point p of the lines supporting c_1 and c_3, its radius decreases. If, during this process, $\mathcal{C}$ loses contact with another essential cut, then a new critical contact has occurred and we are back in Case 1. Otherwise, we can continue sliding indefinitely. However, the only thing that prevents us from sliding $\mathcal{C}$ all the way to when $\mathcal{C}$ degenerates to a point is that it must maintain intersection with both c_1 and c_3 within P. Therefore, $\mathcal{C}$ must have its center and radius chosen so that it *just* intersects both c_2 and c_3 within P. To compute the exact $\mathcal{C}$ in this case: let x and y be the intersections of c_1 and c_3 with the outer boundary of P, on the side closer to p. Draw perpendicular lines to c_1 and c_3 from x and y, respectively. Let these lines intersect the angle bisector; the further of the two intersection points from p is the center of $\mathcal{C}$, and its radius is the distance from the center to the corresponding point (either x or y).
    - If c_1 and c_3 are *parallel*, then $\mathcal{C}$ lies centered somewhere along the bisector line between them. Consider sliding the center of $\mathcal{C}$ along that bisector line. If at any point, $\mathcal{C}$ becomes tangent with another essential cut or come into contact with an endpoint of one, then we are back in Case 1. Else, it must be that all essential cuts other than c_1 and c_3 are contained within the strip between c_1 and c_3. Thus any path from c_1 to c_3 will intersect all essential cuts, and an optimal route can be obtained by computing the geodesic shortest path between c_1 and c_3.

- **Case 3:** $c_1 = c_2 = c_3$. Then $\mathcal{C}$ degenerates to a single point, implying that all of P is visible from that point. This means $\overline{P}$ is star-shaped, and P contains part of the kernel of $\overline{P}$ (set of all points from which $\overline{P}$ can be seen). The optimal (0-length) route γ is a point in said kernel.

Altogether, only a polynomial number of locally minimum illuminating circles exist, and each is computable as above. As in Lemma 2, the optimal solution lies within a localizing square B, concentric with $\mathcal{C}$, of side length $O(n) \cdot \text{radius}(\mathcal{C})$.

C NP-Hardness of the Quota-TSP on infinite lines with obstacles

▶ **Theorem 7.** *The* Quota-TSP on infinite lines *with obstacles is weakly NP-hard.*

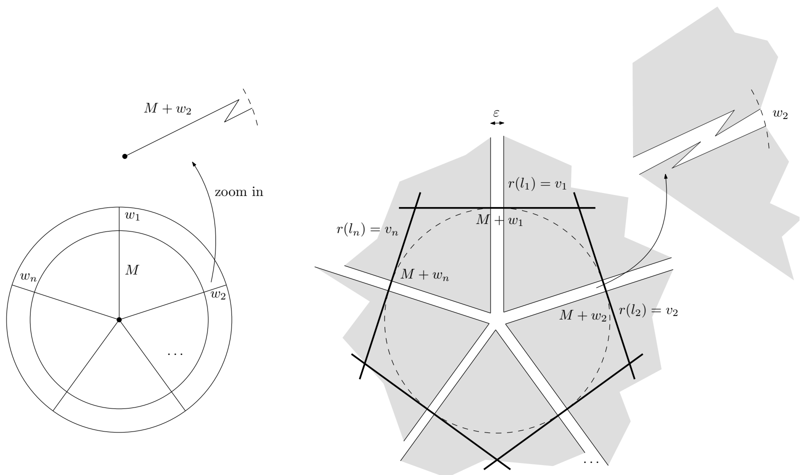

Figure 14 Reduction from Min-Knapsack to Quota-TSP on infinite lines with obstacles.

Proof. We reduce from Min-Knapsack: given n items with weights $\{w_1, \dots, w_n\}$, values $\{v_1, \dots, v_n\}$, and a target value V, select a subset of total value at least V with minimum total weight (this is weakly NP-hard, as its decision version coincides with that of classic Knapsack).

Refer to Figure 14. Draw a circle of radius $M = 1000 \cdot \max\{w_1, \dots, w_n\}$ and a concentric circle of radius $M + \min\{w_1, \dots, w_n\}$; assume without loss of generality this minimum is w_1. From the center, draw n evenly spaced narrow radial corridors reaching the outer circle, each corresponding to an item and terminating at an infinite line tangent to that circle. Corridor i has length exactly $M + w_i$.

If item i's weight exceeds w_1, its corridor includes a zigzag section, confined to a narrow cone of angle much smaller than $\frac{2\pi}{1000n}$ about the radial direction, to bring its length up to exactly $M + w_i$; the zigzag is constructed to avoid self-intersection (shown exaggerated in the figure).

We insert obstacles between adjacent corridors, leaving a clearance of width $\varepsilon < \frac{\max\{|w_i - w_j|\}}{1000n^2}$; these obstacles are narrow, zigzag in congruence with the corridors, and extend well past the outer circle so that leaving it is always suboptimal.

At the end of each corridor i we place an infinite line, tangent to the outer circle, providing reward v_i; the construction ensures line i's reward is collectible only via corridor i.

This construction can clearly be built in polynomial time and has n infinite lines with obstacles of total complexity $O(n)$: since $\frac{M+w_i}{M+w_1} < 2$, each corridor needs only one zigzag section, so each obstacle has constant complexity. Any optimal route for quota V must select exactly the corridors corresponding to an optimal item subset in the Min-Knapsack instance, so solving this Quota-TSP instance solves Min-Knapsack. ◀