Geometric Secluded Paths and Planar Satisfiability

Kevin Buchin
Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands
k.a.buchin@tue.nl

Valentin Polishchuk
Communications and Transport Systems, ITN, Linköping University, Sweden
firstname.lastname@liu.se

Leonid Sedov
Communications and Transport Systems, ITN, Linköping University, Sweden
firstname.lastname@liu.se

Roman Voronov
Institute of Mathematics and Information Technologies, Petrozavodsk State University, Russia
rvoronov@petrsu.ru

Abstract

We consider paths with low *exposure* to a 2D polygonal domain, i.e., paths which are seen as little as possible; we differentiate between *integral* exposure (when we care about how long the path sees every point of the domain) and *0/1* exposure (just counting whether a point is seen by the path or not). For the integral exposure, we give a PTAS for finding the minimum-exposure path between two given points in the domain; for the 0/1 version, we prove that in a simple polygon the shortest path has the minimum exposure, while in domains with holes the problem becomes NP-hard. We also highlight connections of the problem to minimum satisfiability and settle hardness of variants of planar min- and max-SAT.

1 Introduction and Related work

Both visibility and motion planning are textbook subjects in computational geometry – see, e.g., the respective chapters in the handbook [21] and the books [20, 34]. Visibility meets routing in a variety of geometric computing tasks. Historically, the first approach to finding shortest paths was based on searching the visibility graph of the domain; visibility is vital also in computing *minimum-link* paths, i.e., paths with fewest edges [25, 31, 32, 39]. "Visibility-driven" path planning has attracted also some recent interest [3, 37, 44]. In addition to the theoretical considerations, visibility and motion planning are closely coupled in practice: computer vision and robot navigation go hand-in-hand in many courses and real-world applications.
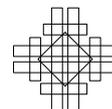
■ **Table 1** Hardness of minimum-exposure paths in polygonal environments.

	0/1 exposure	Integral exposure
With holes	Hard (Thm. 1)	PTAS
Simple	P (Thm. 2)	(Thm. 4)

■ **Table 2** Hardness of versions of planar opt2SAT (see Section 4 for definitions).

	min2SAT	max2SAT
V-cycle	Hard (Thm. 7)	
VC-cycle	Hard (Thm. 7)	Hard (Thm. 7)
Separable	P (Thm. 5)	
Monotone	P (Cor. 6)	

The question of *hiding* a path in a polygonal domain was first raised in a SoCG'88 paper [19]: it considered the *robber route problem* in which the goal is to minimize the length traveled within sight of at least one of a number of *threats* (each threat being a point); the problem reduces to finding the shortest path in the $0/1/\infty$ metric that assigns a cost of 1 to the union of the visibility polygons of the threats, and 0 to the rest of the domain (and infinite weight to the complement of the domain, where travel is forbidden). Our settings are different from [19] in two aspects: (1) we have a *continuum* of the threats (every point in the domain is a threat) and (2) in the integral version, we care for how long threats are seen from points along the path (formally: we integrate the visible area along the path); in other words, we account for the "intensity" of the visibility from the threats.

Lately, motivated by the rise of the Internet of things (IoT) and mobile computing, there has been a surge of research on anonymity, security, confidentiality and other forms of privacy preservation (in particular, in geometric environments [4]), studying paths with minimum exposure to sensors in a network [16, 17, 38, 43]. The standard model, again, assumes a finite number of point sensors, so the visibility is changing discretely, as the path goes in/out of a sensor coverage. To our knowledge, Lebeck, Mølhave and Agarwal [26, 27] were the first to introduce integration of the visibility *continuously* changing along the path (which is also one of our models). Our paper is different from Lebeck et al. in that we give algorithms with provable theoretical performance in continuous domains under the usual notion of distance-independent visibility. Lebeck et al. presented strategies with outstanding practical performance on discretized terrains, in the more realistic model of visibility deteriorating with distance.

Minimizing the integral exposure can be viewed as an extension of the *weighted region problem (WRP)* [2, 9–11, 14, 24, 33, 35] to the case of continuously changing weight, where the weight of a point is the area of its visibility polygon; in the WRP the input is a weighted polygonal subdivision of the domain (with a constant weight assigned to each cell of the subdivision) and the goal is to find the path minimizing the integral of the weight along the path. The computational complexity of the WRP is open; PTASs for the problem have running times that depend not only on the complexity of the subdivision, but also on various parameters of the input like ratio of max/min weight, largest coordinate and angles of the regions, etc. (the parameters differ between the algorithms, see [41, Ch. 31] for details). Integration of other measures of "local quality" (different from visibility) for points along a path was the subject also in the study of *high-quality* paths [1, 45] and related research [46, 47].

Recent papers [8, 18, 29, 42] explored paths adjacent to few vertices in graphs; such paths were dubbed *secluded* in [8]. Our paper may thus be viewed as studying geometric versions of the secluded path problem.

Contributions and Roadmap.   In Section 2 we prove that in a polygonal domain with holes it is NP-hard to find a path, between two given points, minimizing the area seen from the path; the reduction is from minSAT (find the truth assignment to Boolean variables so as to

satisfy the minimum number of given disjunctive clauses). In Section 2.1 we complement the hardness by showing that in a simple polygon, shortest paths are the ones that see minimum area; even more generally, we prove that in a polygon with holes, a locally shortest path sees less area than any path of the same homotopy type (because for a small number of holes the homotopy types can be efficiently enumerated, this implies that the problem is FPT parameterized by the number of holes). Section 3 gives a PTAS for minimizing the integral of the seen area along the path; we first give a generic scheme for building a piecewise-constant approximation of the visibility area for points in the domain, and then in Section 3.2 present details of an implementation which allows applying a PTAS for WRP on our "pixels" with approximately constant seen area. Finally, in Section 4 we further explore the connection between path hiding and minSAT, and determine hardness of versions of planar minSAT (and maxSAT).

Tables 1 and 2 summarize the results. We leave open designing an approximation algorithm for minimizing the seen area, as well as the complexity of the integral version of the problem.

Notation and Problems formulation. We use $|\cdot|$ to denote the measure of a set, i.e., length of a segment and area of a 2D set. Let P be a polygonal domain with n vertices and $s, t \in P$ be two given points in it. For a point $p \in P$ let $V(p) \subseteq P$ be the visibility polygon of p, i.e., the set of points seen by p. We study the following problems:

- GEOMETRIC SECLUDED PATH: Find the s-t path that sees as little area of P as possible (the area seen by a path is defined as the area seen by at least one point of the path, i.e., the so called *weak* visibility region of the path).
- INTEGRAL GEOMETRIC SECLUDED PATH: Find the s-t path π that minimizes the integral of the area of the visibility polygon over the points along the path, $\int_\pi |V(p)| \, dp$.

2　Minimizing seen area

We prove that exposure minimization is NP-hard in general, but in simple polygons the minimum-exposure path is the shortest path.

▶ **Theorem 1.** GEOMETRIC SECLUDED PATH *is NP-hard.*

Proof. We reduce from min2SAT: find truth assignment for a set of n variables, satisfying the minimum number of given two-literal disjunctive clauses. (Inside this proof n will denote the number of variables and c the number of clauses.) Figure 1, left illustrates the construction. A variable gadget is an isosceles triangle. The triangles for the variables are stacked into a *Christmas tree*, with s and t placed at the top and the root respectively. Going through the left (resp. right) vertex of a triangle represents setting the variable to True (resp. False). The clauses are all put on a horizontal line above the Christmas tree so that the segment between any literal and any clause does not intersect the tree. Each clause is connected to its literals, and all connections (including the ones forming the Christmas tree edges) are thin corridors forming the domain; a clause gadget is simply the intersection of the two corridors. The idea of the reduction is to have an s-t path go through all variable gadgets, choosing whether to go through the variable or its negation in every gadget: the fewer clause gadgets are seen, the fewer clauses are satisfied.

A few technicalities have to be taken care of:

- Two variable–clause corridors, leading to different clauses, may intersect *midway*, meaning that the intersection area may be seen twice. We have to make sure that the area of such a midway intersection is much smaller than the clause gadget area. Being smaller by a
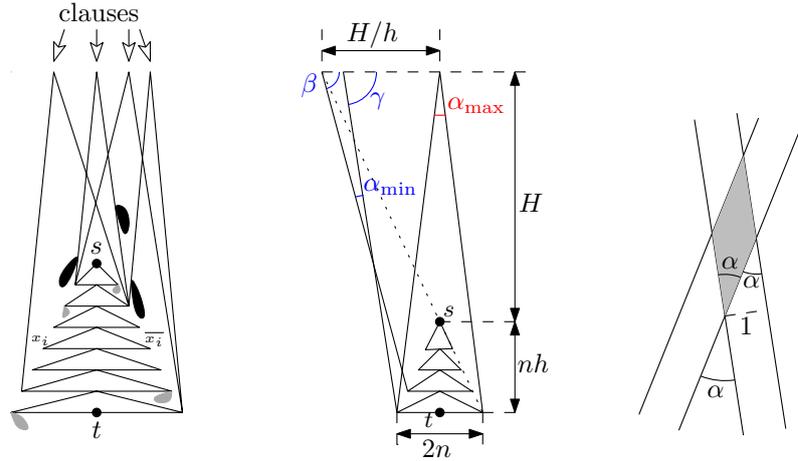
Figure 1 Left: The reduction from min2SAT. All segments are thin corridors of P. Some leakage-blocking high-area chambers are shown black and some area equalizers are shown gray (both black and gray belong to the domain). Middle: the largest angle at a clause and the smallest angle at a midway intersection. The c clauses are spread evenly on the segment of width $2H/h$; thus, the distance between clauses (the base of the triangle with angle $\alpha_{\min}$ at the apex) is $\frac{2H}{h(c-1)}$. Right: Midway intersection of unit-width corridors is area-$\frac{1}{\sin\alpha}$ rhombus with side $\frac{1}{\sin\alpha}$ and angle α.

factor $4c^3$ will suffice: even if parts of a corridor are seen due to the midway intersections with all (at most $2c - 1$) other corridors, the total seen corridor's midway area will still be smaller (by a factor $\approx 2c$) than the area of a single clause gadget. Moreover, with such small midway intersections, they may be neglected altogether when counting the areas of clause gadgets seen from literals: the total areas of all (at most $4c^2$ midway intersections) will be smaller by at least a factor of c than the area of a single clause gadget. To reduce areas of the midway intersections in comparison to the clause gadgets areas, we put the clause gadgets high above the Christmas tree – at height H, to be determined later (Fig. 1, middle). The area of intersection of two corridors (Fig. 1, right) is inversely proportional to the sine of the angle between the corridors (the corridors are all of the same width), so the smallest-area clause gadget would be the one for the clause $x_n \vee \overline{x_n}$ placed directly above the apex of the Christmas tree (since we do not control which clause goes where on the clauses line, we have to consider the worst case); let $\alpha_{\max}$ be the angle between the corridors defining the gadget. Assuming the height of every variable gadget triangle is h and their bases have lengths $2, 4, \ldots, 2n$ (refer to Fig. 1, middle),

$$\alpha_{\max} = 2 \arctan \frac{n}{H + nh}.$$

On the other hand, the *smallest angle* between two interesting corridors that do not lead to the same clause (i.e., the smallest angle that may define the area of a midway intersection) can be formed by corridors leading to last and last-but-one clause from the last-but-one and last variables x_n, x_{n-1} resp. (changing the endpoints of the corridors would only increase the angle of intersection); the angle is

$$\alpha_{\min} = \gamma - \beta = \arctan \frac{H + nh}{H/h - \frac{2H}{h(c-1)} - n} - \arctan \frac{H + (n-1)h}{H/h - (n-1)}.$$

By trigonometric formulas, the ratio $\sin\alpha_{\min}/\sin\alpha_{\max}$, after being squared a constant number of times, is a ratio of polynomials. This ratio tends to infinity as H grows; hence, at a polynomially large H, the ratio becomes larger than $4c^2$, as we need.
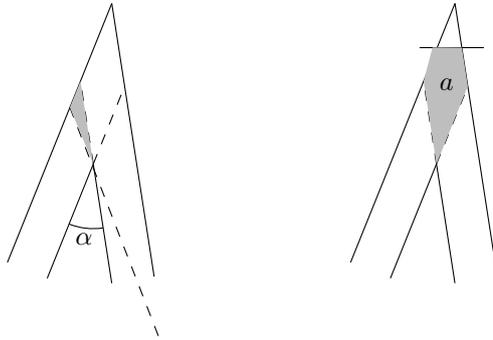
Figure 2 Left: Area seen by one literal only (gray) is negligible for small α. Right: Decreasing clause gadget area.
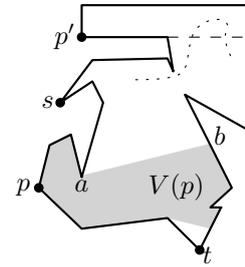
Figure 3 $V(p)$ is shaded; ab is the essential cut of p. A dotted path, crossing the cut of p' (dashed), can be short-cut along the cut.

- We make sure that the area, around a clause gadget, seen from one literal but not from the other (Fig. 2, left), is negligible in comparison with the clause gadget area (seen from both literals of the clause). This is already taken care of by the above, as the whole construction is made tall (large H).

- Leakage of paths from the Christmas tree into variable–clause corridors is prevented by attaching a large-area chamber to each corridor (between the literal and the first intersection of two corridors), so that a path going through the corridor would see the whole area of the chamber. To ensure that the area of a single chamber is larger than the area seen by any path through the Christmas tree, the whole construction is scaled up while keeping the width of the corridors fixed: since the areas available for the chambers grow quadratically with the scaling factor and the areas seen along the corridors grow linearly, a polynomial scaling will suffice to ensure that the chambers areas are large enough to prevent the path going anywhere except through the variable gadgets.

- We attach area equalizers to the literals so that no matter whether the path passes through the variable or its negation, it sees the same non-clause area (the areas may be different between the different variables; we only make sure that for any single variable the seen non-clause area does not depend on whether the variable is set to true or false by the path).

- In the construction so far, different clause gadgets may have different areas; let a denote the smallest area of a clause gadget. We make sure that all clause gadgets have area a, which can be done e.g., by appropriately cutting off the clause gadgets from the top (Fig 2, right).

Now, all s-t paths, going through the Christmas tree only, will see the same non-clause area A. The total area seen by a path is then $\approx A + ka$ where k is the number of clauses seen by the path, which is the same as the number of clauses satisfied by the truth assignment set by the path (we say that the seen area is *approximately* equal to $A + ka$ because of the non-counted areas that may be seen – midway intersections and parts seen by one literal only – which we made sure to be negligible in comparison with a). ◀

In Section 4 we discuss why we could not use *planar* min2SAT to prove hardness of GEOMETRIC SECLUDED PATH, avoiding dealing with the crossings.

2.1 Simple polygons

We show that in a simple polygon shortest paths see least area:

▶ **Theorem 2.** *If P is a simple polygon, the shortest s-t path is the solution to* GEOMETRIC SECLUDED PATH.

Proof. The visibility polygon $V(p)$ of a point $p \in P$ is bounded by edges and chords of P, with each chord connecting a vertex of the polygon to a point on its boundary. If P is a simple polygon and p does not see s ($s \notin V(p)$), then there is a unique chord separating p from s; the chord is called the *essential cut* of p [7] (Fig. 3).

If an essential cut does not separate s from t, then the shortest s-t path does not cross the cut, for otherwise, the path could be shortcut along the cut. That is, the shortest path crosses those and only those cuts that separate s from t. But any other path also has to cross all such cuts, i.e., has to see all the points seen by the shortest path. ◀

For polygons with a small number of holes one may go through all homotopy types of simple (without self-intersections) *s-t* paths: a simple argument shows that a shortcut of a path sees less than the original path, and hence the locally shortest path is the secluded path within its homotopy class.

## 3   A PTAS for minimizing integral exposure

In Section 3.1 we give a generic way to partition the domain in such a way that the visible area is approximately constant within a cell of the partition; then in Section 3.2 we present details of a slightly different partitioning, having straight-line edges, on which a PTAS for the WRP can be applied to find the path with approximately minimum integral exposure.

3.1   Reduction to WRP with curved regions

We first compute the *visibility graph* of P, i.e., the graph connecting pairs of mutually visible vertices of the domain, and extend every edge of the graph in both directions maximally within P. The extensions of the visibility edges split P into $O(n^4)$ cells such that the visibility polygon $V(p)$ is combinatorially the same for any point p within one cell of the subdivision; the subdivision is called the *visibility decomposition* of P [5]. In particular, the area $|V(p)|$ is given by the same formula for any point p in one cell σ of the decomposition. Specifically, the rays from p through the seen vertices of P split $V(p)$ into $O(n)$ triangles (Fig. 4, left). The side of any triangle, opposite to p, is a subset of an edge of P; we call this side the *base* of the triangle. Each of the other, non-base sides is formed by a ray passing through a vertex r' of P and ending at a point r on the base. (In Section 3.2 we will differentiate between *fixed-endpoint* sides for which $r = r'$ is an endpoint of the base and *rotating* rays which rotate around r' if p moves; here we treat both types of sides with a single formula, since fixed-endpoint sides may be viewed as a special case of rotating sides with $r = r'$.)

To write the formula for the area of the triangle pqr, we follow [12, Appendix A.1] and assume that the base is the x-axis and that both $p = (x, y)$ and $r' = (a, b)$ lie above the base ($y, b \geq 0$); then the abscissa of r is $x - y(x - a)/(y - b)$ (Fig. 4, right). Let q' be the vertex that defines the other side, pq, of pqr; to simplify the formulas, assume w.l.o.g. that q' lies on the y-axis: $q' = (0, d)$. The abscissa of q is then $x - yx/(y - d)$, and the area $y|rq|/2$ of the triangle pqr is

$$|pqr| = \frac{y^2}{2} \left(\frac{x - a}{y - b} - \frac{x}{y - d} \right) \tag{1}$$
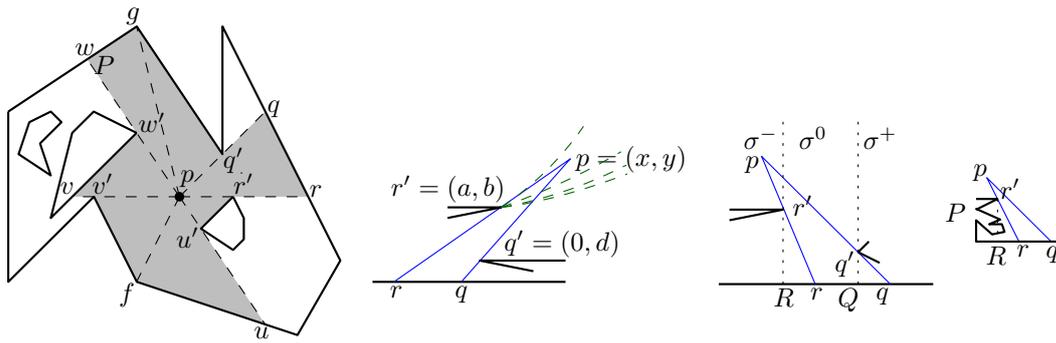
Figure 4 Left: Domain P with 3 holes and a point $p \in P$; $V(p)$ is shaded. Triangle pqr has two rotating sides, triangles puf, pvw', pwg have one fixed-endpoint and one rotating side; the other triangles have two fixed-endpoint sides. Right: $|pqr| = y|rq|/2$. Green dashed curves are level sets of $|pqr|$.

Figure 5 Left: For $p \in \sigma^-$, $|r'Rr|$ is subtracted from C while $|q'Qq|$ is added; for $p \in \sigma^0$, both areas are added; for $p \in \sigma^+$, $|r'Rr|$ is added while $|q'Qq|$ is subtracted. Right: $r'R$ is not fully inside P.

Next, to obtain a piecewise-constant $(1+\varepsilon)$-approximation of the area $|V((x,y))|$ visible from point $(x,y) \in P$, we use level sets of the area function (1). For a given area A, the equality $|pqr| = A$ is attained along the curve γ_A

$$x = \frac{2A/y^2 + a/(y-b)}{1/(y-b) - 1/(y-d)}. \tag{2}$$

Consider a cell σ of the visibility decomposition. We split σ with the curves γ_{A_i} for a set $\mathcal{A} = (A_1, \ldots, A_i, \ldots)$ of areas forming geometric progression with common ratio $1+\varepsilon$: $A_i = (1+\varepsilon)A_{i-1}$. Let S_i denote the set of points p for which the area of the triangle pqr is between A_{i-1} and A_i (that is, $S_i = \{p \in \sigma : A_{i-1} < |pqr| \leq A_i\}$ are the points between $\gamma_{A_{i-1}}$ and γ_{A_i}). We call S_i a *curved sector* because in equation (2), we have $\lim_{y \to b} x(y) = a$ for any A, i.e., all curves γ_A have $r' = (a,b)$ as a common point. (We put a GeoGebra graphics to play with the level sets to see how they look at `https://www.geogebra.org/m/cvxvhfcf`.) We assign the same *weight* A_i to all points in the curved sector; this way, for $i > 1$ the weight of any point $p \in S_i$ is within factor $1+\varepsilon$ of the area of the triangle pqr:

$$|pqr| \leq A_i \leq (1+\varepsilon)|pqr| \qquad \forall p \in S_i, \forall i > 1 \tag{3}$$

For every cell σ of the visibility decomposition, we overlay the level sets from each of the $O(n)$ triangles of $V(p)$ for $p \in \sigma$. We confine the level sets to the cell, i.e., for each curve γ_A use only the intersection $\gamma_A \cap \sigma$. We call each cell of the overlay a *region* and set the weight of the region to the sum of the weights of the curved sectors whose intersection forms the region.

To bound the number of level sets used (i.e., to determine the first area A_1 in the geometric sequence $\mathcal{A}$ and the needed length of the sequence), assume that vertices of P have integer coordinates and let L denote the largest coordinate. (This model and its variants are common for WRP; in particular, the running times of known solutions for WRP [2,9–11,24,33] depend on L.) Now, consider a triangulation T of P – any point $p \in P$ lies inside a triangle τ of T and sees all of the triangle; thus the area $|V(p)|$ is at least the area of τ. Since τ has integer coordinates, by Pick's Theorem [22] the area of the triangle is at least $1/2$:

$$|V(p)| \geq 1/2 \tag{4}$$

We are now ready to prove that it suffices to have

$$A_1 = \frac{\varepsilon}{2n} \tag{5}$$

Indeed, suppose $V(p)$ consists of K triangles of areas $\Delta_1, \ldots, \Delta_K$ and let $A_1, \ldots, A_K$ be the weights of the curved sectors that form the region to which p belongs; the weight of the region is thus $w(p) = A_1 + \cdots + A_K$. Classify the triangles as "small" and "large", with the former having area at most A_1 (and thus having p lie in the sector S_1) and the latter having area larger than A_1 (with p in a sector S_i for $i > 1$); let $l = \{k : \Delta_k > A_1\}$ be the indices of the large triangles. By (3), for every large triangle $k \in l$, $A_k \leq (1 + \varepsilon)\Delta_k$. Since $K \leq n$, we have

$$w(p) = \sum_{k \in l} A_k + \sum_{k \notin l} A_k \leq (1+\varepsilon) \sum_{k \in l} \Delta_k + nA_1 \leq (1+\varepsilon)|V(p)| + \varepsilon\frac{1}{2} \leq (1+2\varepsilon)|V(p)| \tag{6}$$

where the last inequality is due to (4).

▶ **Proposition 3.** *If WRP on N regions with curved boundaries of constant algebraic complexity can be $(1+\varepsilon)$-approximated in time $T(N, \frac{1}{\varepsilon})$, then a $(1 + \varepsilon)^2$-approximation to the minimum integral exposure path can be found in time $T(\frac{n^{10}}{\varepsilon^2} \log^2(nL), \frac{1}{\varepsilon})$.*

Proof. For an upper bound on the sector weight, note that obviously $\forall p \in P, |V(p)| \leq L^2$. Hence, the number of needed level sets is at most $\log_{1+\varepsilon}(2nL^2) = O(\frac{1}{\varepsilon} \log(nL))$. The level sets are defined for each of the $O(n^3)$ triples $r', q', \bar{q}r$ where r', q' are vertices and $\bar{q}r$ is the side of P containing qr; thus overall there are $O(\frac{n^3}{\varepsilon} \log(nL))$ level set curves. Since each curve γ_A has constant algebraic degree (cf. (2)), any two curves intersect $O(1)$ times, so the complexity of the overlay of the level sets inside the cell σ of the visibility decomposition is $O((\frac{n^3}{\varepsilon})^2 \log^2(nL))$. Since there are $O(n^4)$ cells, our construction splits P into $O(\frac{n^{10}}{\varepsilon^2} \log^2(nL))$ regions of constant weight.

By (6), region weights approximate the visibility area to within $1+\varepsilon$ (use $\varepsilon := \varepsilon/2$ to get rid of the factor 2 in front of ε); hence finding a $(1+\varepsilon)$-approximate solution to the WRP on our regions provides a $(1 + \varepsilon)^2$-approximation to the minimum integral exposure path. Formally, let π^* be the minimum integral exposure path (the optimal solution to INTEGRAL GEOMETRIC SECLUDED PATH), let $\bar{\pi}$ be the minimum-weight path through our regions (the optimal solution to WRP) and let π be the $(1+\varepsilon)$-approximate solution to WRP; then

$$\int_\pi |V(p)| \, \mathrm{d}p \; \leq \; \int_\pi w(p) \, \mathrm{d}p \; \leq \; (1+\varepsilon) \int_{\bar{\pi}} w(p) \, \mathrm{d}p \; \leq \; (1+\varepsilon) \int_{\pi^*} w(p) \, \mathrm{d}p \; \leq \; (1+\varepsilon)^2 \int_{\pi^*} |V(p)| \, \mathrm{d}p \tag{7}$$

where the first inequality is due to the left inequality of (3), the second is because π approximates $\bar{\pi}$, the third is because $\bar{\pi}$ is optimal w.r.t. w, and the last one is due to the right inequality in (3). ◀

3.2 A detailed implementation

Applicability of Proposition 3 remains questionable due to absence of an algorithm for WRP with curved regions boundaries. In this section we present another, direct approach to reduce our problem to WRP on a *polygonal* subdivision. We refine the visibility decomposition (without affecting the asymptotic complexity) and recalculate the area functions so that they have *linear* levels. This way, the regions in the overlay of the level sets are convex, so existing WRP solutions can be applied directly.
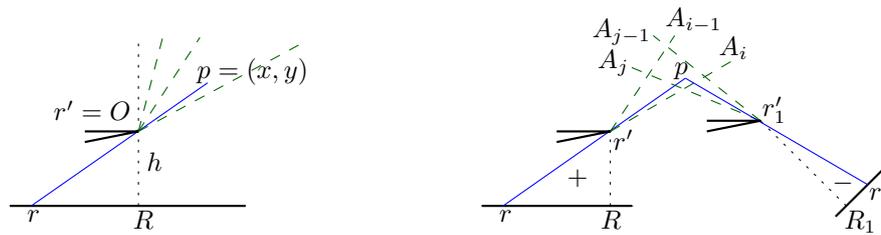
Figure 6 Left: Green dashed are level sets of the area function Δ_r (9). Right: A_i contributes positively to $w(p)$ while A_j comes with minus into $w(p)$, because for p in this region, $r' \in \oplus$ $(r'Rr \in V(p))$ while $r'_1 \in \ominus$ $(r'_1 R' r'_1 \notin V(p))$.

Specifically, we differentiate between fixed-endpoint and rotating sides of the triangles into which $V(p)$ is split: the former end at a vertex of P while the latter rotate around a vertex if p moves (see Fig. 4, left). Triangles whose both sides are fixed-endpoint are easy to handle: (while the area of each individual triangle changes as p moves,) the *total* area of all such triangles remains the same (moving p just redistributes the area between the triangles, "stealing" from some and "giving" to others). We therefore call such triangles *fixed*.

Consider now a triangle pqr whose both sides pq, pr are rotating around vertices q', r' resp. (this is the most general case: if one of the sides, say, pq' is fixed, we can just assume $q = q'$); assume that rq is horizontal (Fig. 5, left). We refine the visibility decomposition by extending the vertical segments through each of r', q' maximally up and down; let R, Q be the feet of the perpendiculars dropped from r' and q' resp. onto the supporting line of pq (any of $r'R, q'Q$ may lie only partially inside P, as in Fig. 5, right – this is not an issue). Note that $|Rr'pq'Q|$ may be added to the fixed-triangles areas – the total area of all fixed triangles plus the areas of the pentagons $Rr'pq'Q$ for all the triangles with p as the apex does not depend on p (while p remains in the same cell). Denote this total area by C. The area $|V(p)|$ is obtained from C by adding/subtracting the areas of the triangles $r'Rr$ for all vertices r' on which a side of a triangle of $V(p)$ rotates – whether $|r'Rr|$ is added or subtracted depends on whether the triangle is in $V(p)$ or not:

$$|V(p)| = C + \sum_{r' \in \oplus} \Delta_r - \sum_{r' \in \ominus} \Delta_r \tag{8}$$

where $\Delta_r = |r'Rr|$ and $\oplus$ (resp. $\ominus$) is the set of vertices whose triangles $r'Rr$ are visible (resp. invisible) from p.

Assume that r' is the origin O and that the supporting line of rR is the horizontal line $y = -h$, and let $p = (x, y)$ with $x \geq 0$ (Fig. 6, left). Then

$$\Delta_r = \frac{h^2}{2} \frac{x}{y} \tag{9}$$

and a level set $\gamma_A = \{p = (x, y) : \Delta_r = A\}$ of the function (9) is a ray (emanating from the origin) of constant x/y: since the height $r'R$ of the triangle is fixed, Δ_r is constant whenever r is fixed. As in Section 3.1, we draw the rays for a set $\mathcal{A} = (A_1, \ldots, A_i, \ldots)$ of areas forming geometric progression with common ratio $1 + \varepsilon$ and assign the weight A_i to all points in the sector $S_i = \{p \in \sigma : A_{i-1} < \Delta_r \leq A_i\}$ between $\gamma_{A_{i-1}}$ and γ_{A_i} (we again use the weight $A_1 = \varepsilon/(2n)$ for points between γ_0 and γ_{A_1}). Also as in Section 3.1, we define a *region* as a cell in the overlay of the rays emanating from the vertices r' of P. Finally, the weight $w(p)$ of any point p in a region is determined by C and the weights of the sectors forming the region: for a vertex $r' \in \oplus$ the weights of the sectors of r' are added to regions weights; for a vertex $r' \in \ominus$, the weights are subtracted (Fig. 6, right).

The fact that our subdivision into regions provides a $(1+\varepsilon)$-approximation to $|V(p)|$ can be argued similarly to Section 3.1:

▶ **Theorem 4.** *If WRP on N regions can be $(1+\varepsilon)$-approximated in time $T(N, \frac{1}{\varepsilon})$, then a $(1 + \varepsilon)^3$-approximation to the minimum integral exposure path can be found in time $T(\frac{n^4}{\varepsilon} \log(nL), \frac{1}{\varepsilon})$.*

4   On planar optimal satisfiability

In this section we return to the (non-integral) GEOMETRIC SECLUDED PATH problem (Section 2) and elaborate on its connections to planar satisfiability, identifying, in particular, polynomially solvable and hard versions of planar minSAT and maxSAT.

For a SAT instance with variables V and clauses C, the graph $G = (V \cup C, E)$ of the instance is the bipartite graph whose vertices are the variables and the clauses, and whose edges connect each variable to a clause whenever the variable or its negation appears in the clause. In a *planar* SAT, G is planar. Planar SAT has been the staple starting point for hardness reduction in computational geometry. In many cases, hardness of geometric problems was proved using *restricted* hard versions of planar SAT, such as:

V-cycle SAT: G remains planar after adding a cycle through V (G is no longer bipartite)

VC-cycle SAT: G remains planar after adding a cycle through $V \cup C$ (this version, as well as V-cycle SAT were defined already in the original paper on planar SAT [28])

Separable SAT: A further restriction of V-cycle SAT: for any variable x, the V-cycle separates clauses containing x from the clauses containing $\overline{x}$; in other words, no variable x has an x-containing clause and a $\overline{x}$-containing clause on the same side of the V-cycle (this version is from [28, Lemma 1], but has no name there; we take the name from [40])

Monotone SAT: In any clause, all variables are either non-negated or all variables are negated (this version is defined for general, not only for planar SAT).

See [15, 36, 40] for in-depth treatment of restricted planar SAT versions and their uses.

When proving hardness of GEOMETRIC SECLUDED PATH in Section 2 (Theorem 1) we spent considerable effort on dealing with crossings between variable–clause connectors. A natural question is why we did not reduce from planar minSAT. The answer is that to avoid crossings, our reduction should better start from separable minSAT (Fig. 7, left), so that for any variable x, the connections from literal x reside on one side of the Christmas tree and the connections from $\overline{x}$ – on its other side (otherwise, a connection from, say, $\overline{x}$ would cross the Christmas tree itself; Fig. 7, middle). However:

▶ **Theorem 5.** *Separable minSAT can be solved in polynomial time.*

Proof. Let A be the clauses on one side of the variable chain and $B = C \setminus A$ – the clauses on the other side. Construct the "clause conflict" graph H [30] whose vertices are the clauses and whose edges connect two clauses whenever one contains the negation of a literal in the other (Fig. 7, right). For any edge, at least one of the conflicting clauses will be satisfied in any truth assignment; thus, every edge in the graph will be incident to a satisfied clause. In particular, solving the minSAT is equivalent to finding minimum vertex cover (VC) in H. By the separability, for any variable x, all clauses with x are in A and all clauses with $\overline{x}$ are in B (or vice versa); thus, any edge of H connects a clause in A with a clause in B, i.e., H is bipartite, and the VC in it can be found in polynomial time. ◀
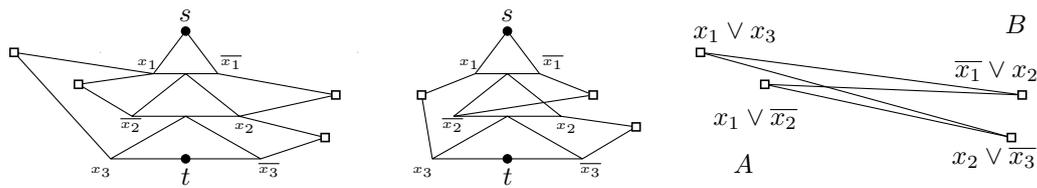
Figure 7 Left: Reduction from separable minSAT to GEOMETRIC SECLUDED PATH would have no crossings (note that some variables have their negations on one side of the Christmas tree, while others – on the other; this is fine, since the definition of separable SAT requires separability *locally* for each variable; the separability does not have to be consistent across all variables). Middle: In non-separable minSAT, clause $\overline{x_1} \lor \overline{x_2}$ could be seen not only from s-t path via $\overline{x_2}$ but also from s-t path via x_2 due to the crossing with the Christmas tree. Right: the graph H (which happens to be $K_{2,2}$) for the instance on the left.

Note that the above proof does not use the planarity. In particular, monotone minSAT can be solved similarly: the clauses with all positive variables can form the set A and the clauses with all negative variables – set B in the graph H from the proof. We thus have:

▶ **Corollary 6.** *Monotone minSAT (planar or not) can be solved in polynomial time.*

In the full version [6], we prove NP-hardness of V- and VC-cycle min2SAT, as well as hardness of all four versions of planar max2SAT (these do not have relation to secluded paths; we give the proofs just for completeness of our treatment of planar optSAT):

▶ **Theorem 7.** *The following planar versions of max2SAT are NP-hard: V-cycle, VC-cycle, monotone, separable. V- and VC-cycle min2SAT are NP-hard.*

5   Conclusion

We studied minimum-exposure paths in polygonal domains. We showed that minimizing seen area is hard in polygons with (large number of) holes, while in polygons with a small number of holes the s-t path that sees least area can be found in polynomial time. We also gave a PTAS for finding an s-t path minimizing the integral of the seen area along the path. Finally, we discussed the connection between the geometric secluded paths and optimizing planar satisfiability, and identified hard and easy cases of planar optSAT (while the planar optSAT variants, which we proved hard, were not used in reductions in this paper, we hope that they may be useful in other settings). We conclude with some remarks on each of the problems studied.

Minimizing seen area and Secluded paths in graphs

Recall that in SECLUDED PATH (the original, graph problem) the goal is to find an s-t path adjacent to fewest vertices of the graph (vertices of the path itself are also counted as adjacent to the path). The problem was proved hard in [8]. Our proof of hardness of GEOMETRIC SECLUDED PATH (Theorem 1) gives an alternative proof of hardness of SECLUDED PATH in graphs: simply remove equalizers and leakage-blocking chambers from Fig. 1 (no need to care about midway intersections and all the other geometric technicalities) and add a large number of extra vertices adjacent to each clause vertex (Fig. 8, left). While our proof is simpler than the ones in [8], it is less powerful because Chechik et al. [8] showed also hardness of approximation. In fact, the reduction in [8], shown here on Fig. 8, right, may also be seen as reduction from minSAT (in view of the connection between minSAT and VC
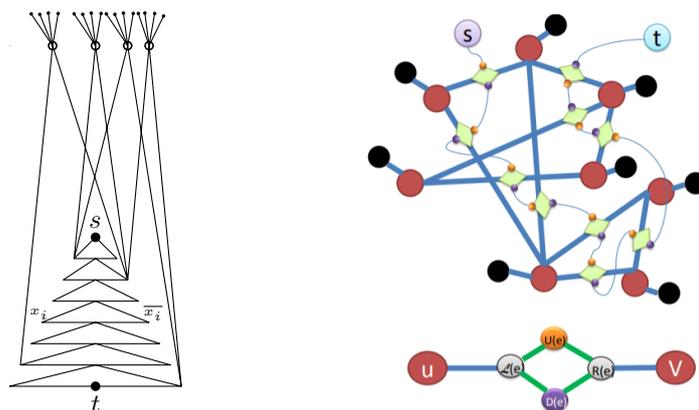
Figure 8 Left: Our reduction from min2SAT to SECLUDED PATH. To avoid high-degree vertices at the clauses (hollow), the *s-t* path will go via the Christmas tree, setting the variables; the number of seen (i.e., adjacent) clause vertices is the number of satisfied clauses. Right: The reduction from VC in a graph G [8, Fig. 3]: the new graph G' has new vertices s and t, and an *s-t* path (thin blue) crossing all edges (thick blue) is added to G, with every crossing (lightgreen rhombi) turned into a gadget (bottom) where the *s-t* path chooses which vertex of G (red) the path will see; leaking into the original vertices of G (red) is prevented in G' by attaching high-weight vertices (black).

in the clause conflict graph – see proof of Theorem 5): the choices that the *s-t* path makes in the edges of the original graph G may be seen as setting the truth values to the variables (similarly to how the path through our Christmas tree does it).

A natural question, arising in view of the effort we spent dealing with the crossings in Section 2 when proving hardness of GEOMETRIC SECLUDED PATH (Theorem 1), is why we did not reduce from SECLUDED PATH in *planar* graphs. The answer is that we are not aware of a hardness result for the problem in planar graphs. Indeed, even though Chechik et al.'s hardness proof for *general* graphs (refer to Fig. 8, right) could reduce from VC in a *planar* graph G, in order to keep the planarity also in the resulting graph G' (in which the secluded *s-t* path is sought), the added path (crossing all edges of G) must cross each edge exactly once, meaning that it is an Euler path in the planar dual of G, meaning that the dual has vertices of even degree only, meaning that G has faces with even number of edges, meaning it has only even cycles, meaning it is bipartite, meaning VC is polynomial in it. (Strictly speaking, since we need only an Euler path through the edges, not Euler cycle, G may have 2 odd faces – we believe VC is still polynomial in such graphs).

The PTAS for integral seen area minimization

Several remarks on the complexity of our solution:

- A faster algorithm for our problem could potentially be obtained by using a "1D" discretization of edges of the visibility decomposition (instead of creating a 2D "grid" of regions, as we do), as done in many algorithms for WRP (and related problems on minimizing path integral [1, 45]). Such a solution, however, would require knowing the optimal path connecting points on the boundary of the same cell of the decomposition. This, may be quite complicated, as it amounts to minimizing the integral of a function with $\Omega(n)$ terms, for which an analytical solution might not exist (though an approximation may be possible).

- An algorithm for WRP with regions whose boundaries are curves of constant algebraic degree could be interesting and would lead to a solution of our problem just using the generic scheme from Section 3.1. The biggest stumbling block for the design of such an
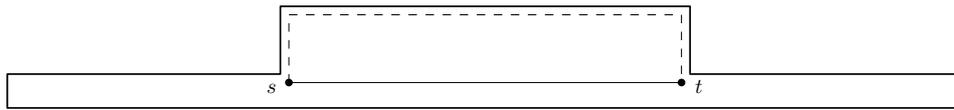
Figure 9 Shortest *s-t* path (solid) sees the niches behind s and t for its whole length; stepping to the side (dashed path) decreases the integral exposure.

algorithm may be the non-convexity of the regions, implying that a segment between two points on the boundary of a region is not guaranteed to stay inside the region. It may be possible that WRP techniques could be adapted to handle our regions from Section 3.1 by approximating their boundaries with piecewise-linear functions (since we are looking only for a $(1+\varepsilon)$-optimal path, the fineness of such piecewise-linear approximation would also be controlled by ε).

Since our problem is an extension of WRP to the case of continuously changing weight, it may be tricky to establish hardness of the problem, as the complexity of WRP has remained open for many years (see [14] for a recent proof of algebraic complexity of WRP). Differently from 0/1 exposure (Theorem 3), even in simple polygons the shortest path does not necessarily minimize the integral exposure (Fig. 9).

Optimal 2-satisfiability

Few observations on min2SAT and max2SAT:

Monotone minSAT is an example of the tractable class of submodular function minimization [23].

Planar max2SAT has a PTAS [13, Thm. 8.8].

If in a separable max2SAT with VC cycle, the cycle also separates the variables at the clauses (i.e., if at each clause the connections from the two variables come from the different sides of the cycle), then the problem can be solved in polynomial time by reduction to separable min2SAT.

References

1   Pankaj K Agarwal, Kyle Fox, and Oren Salzman. An efficient algorithm for computing high-quality paths amid polygonal obstacles. *ACM Transactions on Algorithms (TALG)*, 14(4):46, 2018.

2   Lyudmil Aleksandrov, Anil Maheshwari, and J-R Sack. Determining approximate shortest paths on weighted polyhedral surfaces. *Journal of the ACM (JACM)*, 52(1):25–53, 2005.

3   Esther M Arkin, Alon Efrat, Christian Knauer, Joseph Mitchell, Valentin Polishchuk, Günter Rote, Lena Schlipf, and Topi Talvitie. Shortest path to a segment and quickest visibility queries. *Journal of Computational Geometry*, 7(2):77–100, 2016. Special issue on SoCG'15.

4   Boris Aronov, Alon Efrat, Ming Li, Jie Gao, Joseph SB Mitchell, Valentin Polishchuk, Boyang Wang, Hanyu Quan, and Jiaxin Ding. Are friends of my friends too social? limitations of location privacy in a socially-connected world. In *Proceedings of the Eighteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 280–289. ACM, 2018.

5   Boris Aronov, Leonidas J Guibas, Marek Teichmann, and Li Zhang. Visibility queries and maintenance in simple polygons. *Discrete & Computational Geometry*, 27(4):461–483, 2002.

6   Kevin A. Buchin, Valentin Polishchuk, Leonid Sedov, and Roman Voronov. Geometric secluded paths and planar satisfiability. *CoRR*, abs/1902.06471, 2019. `arXiv:1902.06471`.

7   Svante Carlsson, Håkan Jonsson, and Bengt Nilsson. Finding the shortest watchman route in a simple polygon. *Discrete & Computational Geometry*, 22(3):377–402, 1999. `doi:10.1007/PL00009467`.

8   Shiri Chechik, Matthew P Johnson, Merav Parter, and David Peleg. Secluded connectivity problems. *Algorithmica*, 79(3):708–741, 2017.

9   Siu-Wing Cheng, Jiongxin Jin, and Antoine Vigneron. Triangulation refinement and approximate shortest paths in weighted regions. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 1626–1640. SIAM, 2014.

10   Siu-Wing Cheng, Hyeon-Suk Na, Antoine Vigneron, and Yajun Wang. Approximate shortest paths in anisotropic regions. *SIAM Journal on Computing*, 38(3):802–824, 2008.

11   Siu-Wing Cheng, Hyeon-Suk Na, Antoine Vigneron, and Yajun Wang. Querying approximate shortest paths in anisotropic regions. *SIAM Journal on Computing*, 39(5):1888–1918, 2010.

12   Otfried Cheong, Alon Efrat, and Sariel Har-Peled. Finding a guard that sees most and a shop that sells most. *Discrete & Computational Geometry*, 37(4):545–563, 2007.

13   Nadia Creignou, Sanjeev Khanna, and Madhu Sudan. *Complexity classifications of boolean constraint satisfaction problems*, volume 7. SIAM, 2001.

14   Jean-Lou De Carufel, Carsten Grimm, Anil Maheshwari, Megan Owen, and Michiel Smid. A note on the unsolvability of the weighted region shortest path problem. *Computational Geometry*, 47(7):724–727, 2014.

15   Erik Demaine. Algorithmic lower bounds: Fun with hardness proofs. MIT OCW.

16   Hristo N Djidjev. Efficient computation of minimum exposure paths in a sensor network field. In *International Conference on Distributed Computing in Sensor Systems*, pages 295–308. Springer, 2007.

17   Hao Feng, Lei Luo, Yong Wang, Miao Ye, and Rongsheng Dong. A novel minimal exposure path problem in wireless sensor networks and its solution algorithm. *International Journal of Distributed Sensor Networks*, 12(8):1550147716664245, 2016.

18   Fedor V Fomin, Petr A Golovach, Nikolay Karpov, and Alexander S Kulikov. Parameterized complexity of secluded connectivity problems. *Theory of Computing Systems*, 61(3):795–819, 2017.

19   Laxmi Gewali, Alex C. Meng, Joseph S. B. Mitchell, and Simeon C. Ntafos. Path planning in 0/1/infinity weighted regions with applications. In Herbert Edelsbrunner, editor, *Proceedings of the Fourth Annual Symposium on Computational Geometry, Urbana-Champaign, IL, USA, June 6-8, 1988*, pages 266–278. ACM, 1988.

20   Subir Ghosh. *Visibility Algorithms in the Plane*. Cambridge University Press, New York, NY, USA, 2007.

21   J.E. Goodman and J. O'Rourke, editors. *Handbook of Discrete and Computational Geometry*. Discrete Mathematics and Its Applications. Taylor & Francis, 2nd edition, 2004.

22   Branko Grünbaum and Geoffrey C Shephard. Pick's theorem. *The American Mathematical Monthly*, 100(2):150–161, 1993.

23   Dorit S Hochbaum. Complexity and approximations for submodular minimization problems on two variables per inequality constraints. *Discrete Applied Mathematics*, 250:252–261, 2018.

24   Rajasekhar Inkulu and Sanjiv Kapoor. A polynomial time algorithm for finding an approximate shortest path amid weighted regions. *Preprint*, 2015.

25   Irina Kostitsyna, Maarten Löffler, Valentin Polishchuk, and Frank Staals. On the complexity of minimum-link path problems. *JoCG*, 8(2):80–108, 2017. Special Issue on SoCG'16. URL: `http://jocg.org/index.php/jocg/article/view/328`, `doi:10.20382/jocg.v8i2a5`.

26   Niel Lebeck, Thomas Mølhave, and Pankaj K Agarwal. Computing highly occluded paths on a terrain. In *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 14–23. ACM, 2013.

27   Niel Lebeck, Thomas Mølhave, and Pankaj K Agarwal. Computing highly occluded paths using a sparse network. In *Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 3–12. ACM, 2014.

28   David Lichtenstein. Planar formulae and their uses. *SIAM journal on computing*, 11(2):329–343, 1982.

29   Max-Jonathan Luckow and Till Fluschnik. On the computational complexity of length-and neighborhood-constrained path problems. *arXiv preprint*, 2018. `arXiv:1808.02359`.

30   Madhav V Marathe and SS Ravi. On approximation algorithms for the minimum satisfiability problem. *Information Processing Letters*, 58(1):23–29, 1996.

31   J. Mitchell, G. Rote, and G. Woeginger. Minimum-link paths among obstacles. *Alg-ca'92*, 8(1):431–459, 1992.

32   Joseph Mitchell, Valentin Polishchuk, and Mikko Sysikaski. Minimum-link paths revisited. *CGTA*, 47(6):651–667, 2014. `doi:10.1016/j.comgeo.2013.12.005`.

33   Joseph SB Mitchell and Christos H Papadimitriou. The weighted region problem: finding shortest paths through a weighted planar subdivision. *Journal of the ACM (JACM)*, 38(1):18–73, 1991.

34   Joseph O'Rourke. *Art Gallery Theorems and Algorithms*. The International Series of Monographs on Computer Science. Oxford University Press, New York, NY, 1987.

35   Christos H Papadimitriou. An algorithm for shortest-path motion in three dimensions. *Information Processing Letters*, 20(5):259–263, 1985.

36   Alexander Pilz. Planar 3-sat with a clause/variable cycle. *arXiv preprint*, 2017. `arXiv:1710.07476`.

37   Valentin Polishchuk and Leonid Sedov. Gender-aware facility location in multi-gender world. In *LIPIcs-Leibniz International Proceedings in Informatics*, volume 100. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

38   Yuning Song, Liang Liu, Huadong Ma, Athanasios V Vasilakos, et al. A biology-based algorithm to minimal exposure problem of wireless sensor networks. *IEEE Trans. Network and Service Management*, 11(3):417–430, 2014.

39   Subhash Suri. A linear-time algorithm for minimum link paths inside a simple polygon. *Computer Vision, Graphics and Image Processing*, 35(1):99–110, 1986.

40   Simon Tippenhauer and Wolfgang Muzler. On planar 3-sat and its variants. *Fachbereich Mathematik und Informatik der Freien Universitat Berlin*, 2016.

41   Csaba D Toth, Joseph O'Rourke, and Jacob E Goodman. *Handbook of discrete and computational geometry*. Chapman and Hall/CRC, 2017.

42   René van Bevern, Till Fluschnik, and Oxana Yu. Tsidulko. Parameterized algorithms and data reduction for safe convoy routing. In *18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems, ATMOS 2018, August 23-24, 2018, Helsinki, Finland*, pages 10:1–10:19, 2018.

43   Giacomino Veltri, Qingfeng Huang, Gang Qu, and Miodrag Potkonjak. Minimal and maximal exposure path algorithms for wireless embedded sensor networks. In *Proceedings of the 1st international conference on Embedded networked sensor systems*, pages 40–50. ACM, 2003.

44   Haitao Wang. Quickest visibility queries in polygonal domains. In Boris Aronov and Matya Katz, editors, *33rd International Symposium on Computational Geometry, SoCG 2017, July 4-7, 2017, Brisbane, Australia*, volume 77 of *LIPIcs*, pages 61:1–61:16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.

45   Ron Wein, Jur Van Den Berg, and Dan Halperin. Planning high-quality paths and corridors amidst obstacles. *The International Journal of Robotics Research*, 27(11-12):1213–1231, 2008.

46   Ron Wein, Jur Van Den Berg, and Dan Halperin. Planning near-optimal corridors amidst obstacles. In *Algorithmic Foundation of Robotics VII*, pages 491–506. Springer, 2008.

47   Ron Wein, Jur P Van den Berg, and Dan Halperin. The visibility–voronoi complex and its applications. *Computational Geometry*, 36(1):66–87, 2007.