

Constant-Factor Approximation Algorithms for Convex Cover and Hidden Set in a Simple Polygon

Reilly Browne*, Prahlad Narasimham Kasthurirangan†, Joseph S. B. Mitchell‡, Valentin Polishchuk§

*Stony Brook University; reilly.browne@stonybrook.edu

†Stony Brook University; prahladnarasim.kasthurirangan@stonybrook.edu

‡Stony Brook University; joseph.mitchell@stonybrook.edu

§Linköping University; valentin.polishchuk@liu.se

Abstract—Given a simple polygon P , the minimum convex cover problem seeks to cover P with the fewest convex polygons that lie within P . The maximum hidden set problem seeks to place within P a maximum cardinality set of points no two of which see each other. We give constant factor approximation algorithms for both problems. Previously, the best approximation factor for the minimum convex cover was logarithmic; for the maximum hidden set problem, no approximation algorithm was known.

I. INTRODUCTION

In this paper we study two fundamental optimization problems in a geometric setting. One is a set cover problem: Cover a simple polygon P with the fewest convex polygons (we denote that number $cc(P)$) that lie within P ; this is the *convex cover problem* for a simple polygon. We also study a maximum independent set problem: Pack into a simple polygon P as many points as possible so that no two points see each other (points $p, q \in P$ are *visible*, or *see each other*, if the segment pq lies within P); this is the *hidden set problem* for a simple polygon and we let $hs(P)$ denote the maximum size of a hidden set. The hidden set problem is the maximum independent set problem in the “point visibility graph” of P , whose nodes are all points (the continuum) in P and whose edges link pairs of points that are visible to each other.

Minimum convex cover has been studied for many years, beginning with the early work of Pavlidis [27]. Most recently, Abrahamsen [1] has shown the problem to be $\exists\mathbb{R}$ -complete. In terms of approximation algorithms to compute $cc(P)$, the best result is an $O(\log n)$ -approximation algorithm, with running time $O(n^{29} \log n)$ where n is the complexity of P , found some decades ago by Eidenbenz and Widmayer [12], who also show that the problem is APX-hard. As our

main result, we give the first constant-factor approximation algorithm for computing $cc(P)$; we also drastically improve the running time.

Finding $hs(P)$ is APX-hard [11] and no prior approximation algorithm has been known for computing $hs(P)$ in general. (Alegria, Bhattacharya and Ghosh gave a $1/4$ -approximation [3] for finding a maximum hidden set of vertices.) We give the first approximation algorithm for computing $hs(P)$; our approximation factor is constant.

Other Related Work: A related covering problem that has been extensively studied is the *guarding problem* [32, Chapter 33], in which we seek to cover a polygon P with the fewest star-shaped polygons, each being a visibility polygon of some point (a “guard”) within P . This problem has long been known to be NP-hard and has recently been shown to be $\exists\mathbb{R}$ -complete [2], [29]. There have been recent advances also in computing approximately optimal guard sets [5], [10], [19]. The problem of convex covering can be contrasted with the problem of partitioning a simple polygon into a minimum number of (interior-disjoint) convex polygons; the convex partitioning problem is solvable exactly in polynomial time, as partitioning allows one to use dynamic programming to optimize, both in the case of partitioning with diagonals (chords between two vertices of P) [23] and in the case of general partitions, allowing Steiner points [7]. (There is also a very simple 4-approximation for convex partitioning of P that runs in linear time [22].)

The hidden set problem is a special case of a geometric maximum independent set problem: find a maximum independent set in the (continuous, infinite) graph whose nodes are the points within P and whose edges are determined by interpoint visibility. Given the difficulty of computing maximum independent sets in general graphs,

even approximately [18], [24], geometric instances of maximum independent set have attracted considerable attention; e.g., there has been recent progress in computing maximum independent sets among axis-parallel rectangles in the plane, including new, constant-factor approximation algorithms [17], [25].

II. PRELIMINARIES AND AN OVERVIEW

Let P be a simple polygon with n vertices, labeled v_1, v_2, \dots, v_n (ordered clockwise). The edges of P are denoted $e_i = v_i v_{i+1}$, for $1 \leq i \leq n$, with $v_{n+1} = v_1$. We consider P to be a closed region, including the boundary, denoted ∂P .

Two points $p, q \in P$ are said to be *visible* (or to *see each other*) if the line segment pq lies within P . A simple polygon P is said to be *weakly visible* from a line segment σ if every point of P sees at least one point of σ . In this paper, whenever we say that a polygon P is weakly visible we will mean that it is weakly visible from an edge W of P that is also an edge of the convex hull of P .

A *hidden set* in P is a set S of points in P such that no two points in S see each other. A *convex cover* of P is a set of convex polygons whose union equals P . Let $hs(P)$ and $cc(P)$ denote the cardinalities of a maximum hidden set and a minimum convex cover of P , respectively. Since there is at most one hidden point within any convex subset of P , we must have the following basic inequality:

$$hs(P) \leq cc(P) \quad (1)$$

Overview of Results and Methods: Our main results are as follows:

(1) We give a polynomial-time 6-approximation algorithm for computing $cc(P)$ for a simple polygon P ; see Theorem 11 of Section IV. We do this in two parts: First, we give a polynomial-time 2-approximation algorithm (Theorem 5, Section III) for computing $cc(P)$ in the case that P is weakly visible (from a convex hull edge, $W = e_n = v_n v_1$, of P). Second, we utilize a decomposition of P into weakly visible subpolygons, the *window partition* of P , and argue (Lemma 10(1)) that any convex body within P intersects at most 3 subpolygons. The 2-approximation result in weakly visible P is obtained by formulating the coverage of the edges e_i (for $1 \leq i \leq n-1$) of P as a problem of computing a minimum path cover in a directed acyclic graph (DAG) whose nodes correspond to edges of P and whose arcs correspond to pairs of “strongly visible” edges whose convex hull lies within P . Each path π in the DAG corresponds to a convex polygon, P_π , within P . We then show (Lemma 4) that for any set of k paths in a

path cover, there is a convex cover of size $2k$. One can compute a minimum path cover of a DAG in polynomial time; if it has k paths, then, by Dilworth’s theorem, there is an “antichain” of size k , which corresponds to a set of k edges of P that are “independent” in the sense that no two of them are strongly visible. Finally, we argue (Lemma 6) that there are in fact k points, one on each edge of the antichain, that form a hidden set in P , showing that $hs(P) \geq k$, so $cc(P) \geq k$, implying that our set of $2k$ convex polygons covering P is a 2-approximation (Theorem 8).

(2) We give a polynomial-time $(1/8)$ -approximation for computing $hs(P)$ (see Theorem 11). This is again obtained in two parts, first obtaining a $(1/2)$ -approximation (Theorem 9) in weakly visible polygons (utilizing the antichain result for an optimal path cover), and then using the fact (Lemma 10(2)) that a general polygon P can be decomposed (using the window partition) into 4 classes of weakly visible subpolygons, with the property that no point in a polygon of one class can see any point within another polygon of the same class.

III. WEAKLY VISIBLE POLYGONS

In this section, we assume that P is a weakly visible polygon, weakly visible from the edge $W = v_n v_1$, which is an edge of the convex hull of P . Without loss of generality, we assume that W is horizontal and that P lies above W ; see Fig. 1. Let C denote the polygonal chain, with vertices (v_1, v_2, \dots, v_n) , that comprises the boundary of P except for the one edge W . For two points $a, b \in C$ the portion of C between a and b is denoted $C(a, b)$, and for $a, b \in P$ the shortest path within P from a to b is denoted $SP(a, b)$.

A basic property of weakly visible polygons is that for checking the visibility between two points $a, b \in C$ it suffices to check that the visibility is not blocked by $C(a, b)$. Formally:

Fact 1. [Chord property] *If $ab \cap C(a, b) = \{a, b\}$, then a and b see each other.*

Proof. It is known that, for a weakly visible polygon P , $SP(a, b)$ bends only on vertices of $C(a, b)$ [15, Lemma 1]; in particular, if the shortest path does not touch any vertex of $C(a, b)$, the shortest path has no bends, i.e., is a segment. \square

A. Edges as a Poset

Consider the directed graph G , whose nodes are the edges, e_1, e_2, \dots, e_{n-1} , of C . The (directed) arcs of G are ordered pairs (e_i, e_j) such that $i < j$ and the convex hull of the edges e_i, e_j lies within P (Fig. 1). For a

(directed) path $\pi = (e_{i_1}, e_{i_2}, \dots, e_{i_m})$ in G let P_π be the convex hull of the edges e_{i_1}, \dots, e_{i_m} ; we define the *base* of the convex polygon P_π to be the segment $v_{i_1}v_{i_m+1}$. As shorthand, we refer to the first vertex of the base, v_{i_1} , as a , (or a_π) and the last vertex v_{i_m+1} as b (or, respectively, b_π). We observe that paths in G define convex polygons in P :

Lemma 2. $P_\pi \subseteq P$.

Proof. Since any two consecutive edges $e_{i_k}, e_{i_{k+1}}$ in π are connected by an arc of G , the convex hull of the edges $e_{i_k}, e_{i_{k+1}}$ belongs to P ; in particular, the subchain $C(v_{i_k}, v_{i_{k+1}})$ of the boundary of P does not intersect P_π , implying overall that the subchain $C(v_{i_k}, v_{i_{m+1}})$ does not intersect the base $v_{i_1}v_{i_m+1}$. Thus, by the chord property (Fact 1), v_{i_1} and v_{i_m+1} see each other, implying that the boundary of P does not intersect the boundary of P_π . \square

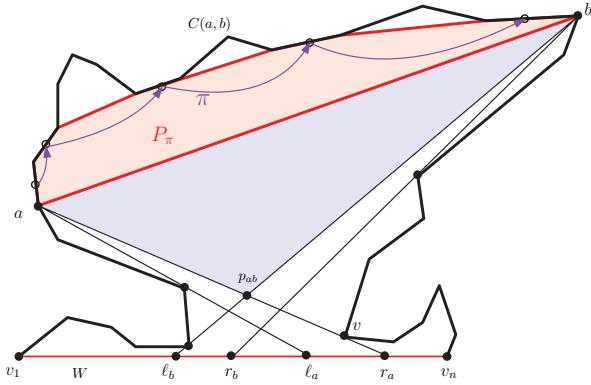


Fig. 1. A weakly visible polygon. Arcs of the graph G connect pairs of edges whose convex hull is in P ; some arcs are shown in purple, forming the path π . The convex polygon P'_π is the union of P_π (red) and the blue triangle abp_{ab} .

For every arc (e_i, e_j) of G we have $i < j$; thus, the graph G is a directed acyclic graph (DAG), whose transitive closure defines a partially ordered set (poset). In fact, Lemma 2 implies that G is its own transitive closure: if (e_i, e_j) and (e_j, e_k) are arcs of G , then (e_i, e_k) is also an arc of G . A *path cover* of a directed graph is a set of directed paths such that every node of the graph belongs to (at least) one of the paths; an *antichain* is a set of nodes such that no two nodes in the set are connected by a directed path. Since G is a DAG, both a minimum path cover, Π , of G and a maximum antichain, I , can be computed in polynomial time by the folklore reduction to a maximum bipartite matching [13]. By Dilworth's theorem, the number $|\Pi|$ of paths

equals the number $|I|$ of nodes (edges of P) in G ; let $k = |\Pi| = |I|$.

For each path $\pi \in \Pi$, we define below (Section III-B) two associated convex polygons within P . We then prove (Lemma 4) that these $2k$ convex polygons cover P . In Section III-C we show how to obtain a set H of k hidden points in P . Thus, using inequality (1), we get that $k \leq hs(P) \leq cc(P) \leq 2k$, showing that the $2k$ convex polygons and the k hidden points that we compute yield approximations, with factor 2, for the minimum convex cover and the maximum hidden set of P . This is our main technical result.

B. A Convex Cover Based on Path Cover

For any point $p \in P$, we let $\ell_p r_p \subseteq W$ be the subset of W that is seen by p , with ℓ_p (resp., r_p) the left (resp., right) endpoint. We define the polygon, $P'_\pi \supseteq P_\pi$, to be the union of P_π and the triangle, abp_{ab} , p_{ab} is the point where the segments bl_b and ar_a cross (see Fig. 1). (They must cross, since ar_a contains a chord, av , that separates b from W .) An easy observation is:

Claim 3. $Polygon P'_\pi$ is convex.

Proof. Points along the edge, $e_{i_1} = (a = v_{i_1}, v_{i_1+1})$, incident on a , must see W (since P is weakly visible from W), implying that the extension of e_{i_1} into P cannot pass above the segment av (which contains p_{ab}), as this segment separates ab from W . Thus, the vertex a is convex in polygon P'_π . The same argument applies to b . \square

While the convex polygons P_π corresponding to paths π in a path cover Π will necessarily cover the edges of C , and the polygons P'_π will cover even more of the interior of P , they need not cover all of P ; see Fig. 2 (top and middle). However, by adding as well a triangle, $bl_b r_b r_b$, for each base ab , we do obtain a covering of P (see Fig. 2, bottom), as our next lemma shows.

Lemma 4. *For any set Π of paths π that form a path covering of the edges e_i , $i = 1, 2, \dots, n-1$, of a weakly visible polygon P , the polygon P is covered by the union, over $\pi \in \Pi$, of the polygons P'_π together with the triangles $b_\pi \ell_{b_\pi} r_{b_\pi}$.*

Proof. Since we assume that Π is a path covering, all of the edges e_i , for $1 \leq i \leq n-1$, of P are covered by the convex polygons P_π and thus by the (superset) convex polygons P'_π . It remains to argue that the interior of P , as well as W , is covered by the polygons P'_π and the specified triangles.

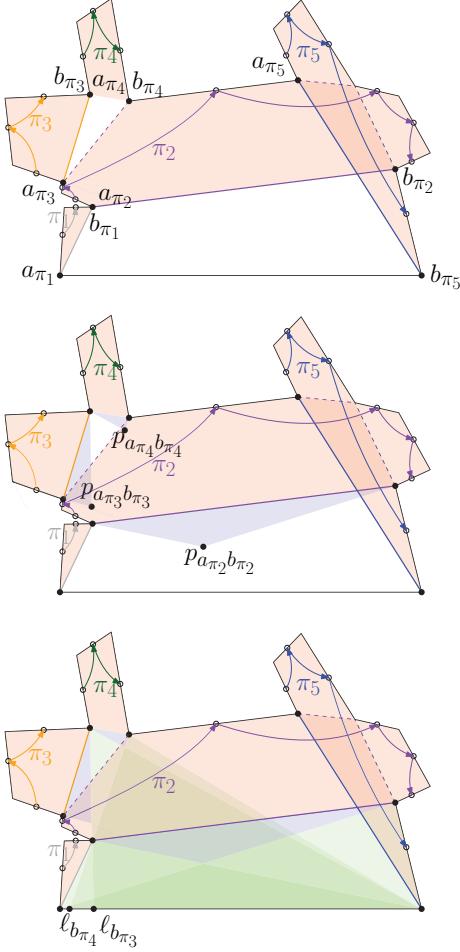


Fig. 2. Top: The convex polygons P_π for paths $\pi \in \{\pi_1, \pi_2, \dots, \pi_5\}$ in a path cover of C . Middle: The convex polygons P'_π obtained by augmenting P_π with the blue triangle abp_{ab} ; note that some (white) portions of P are uncovered. Bottom: The covering using polygons P'_π and the (green) triangles $b\ell_b r_b$.

Let $p \in P$ be any point interior to P or interior to W . Let ℓ'_p (resp., r'_p) be the point on the boundary of P first hit by a ray from p in the direction away from ℓ_p (resp., r_p), along the line through p and ℓ_p (resp., r_p). Note that it could be that $\ell_p = v_1$ or that $r_p = v_n$, or both. (If p lies on W , then $r'_p = v_1$ and $\ell'_p = v_n$.) Refer to Fig. 3.

If p lies within one of the convex polygons P_π , we are done (since $P'_\pi \supseteq P_\pi$). Otherwise, we know that the segment pr'_p (as well as the segment $p\ell'_p$) must intersect at least one base, $a_\pi b_\pi$, of a convex polygon P_π , for $\pi \in \Pi$, since its endpoint r'_p lies on the boundary of P , and the convex polygons P_π cover all of the edges e_i ($1 \leq i \leq n-1$). There are two cases:

(1) There is a base $ab = a_\pi b_\pi$ (of some P_π for $\pi \in$

Π) that intersects both pr'_p and $p\ell'_p$. In this case, we claim that p must lie within the (blue) triangle abp_{ab} . The point a and the edge W lie on the same side of the chord vr'_p ; thus, ara cannot cross the chord vr'_p , implying that the chord vr'_p (and thus the point p) must lie above the line through a and r_a . Similarly, p must lie above the line through b and ℓ_b . Thus, p lies within the triangle abp_{ab} and thus within P'_π . See the figure on the top in Fig. 3.

(2) There is not a base $a_\pi b_\pi$ that intersects both pr'_p and $p\ell'_p$. Then, there must be a base $ab = a_\pi b_\pi$ (of some P_π) that is crossed by the segment pr'_p , with the endpoint $b \in C(r'_p, l'_p)$. We distinguish two subcases:

(a) Point p sees b . In this case, p is covered by the corresponding triangle $b\ell_b r_b$ since the ray from b through p hits W .

(b) Point p does not see b . Then, the (unique) shortest path in P from p to b is not a line segment; let pv_i be the edge of this shortest path that is incident to p , with v_i a (reflex) vertex of P . Note that $v_i \in C(b, v_n)$. Therefore, the two edges incident to v_i both lie to the right of the ray from p to v_i . Then the chord, $v_i\alpha_i$, that arises from extending the edge $v_{i-1}v_i$ beyond v_i to a point $\alpha_i \in \partial P$ must have its endpoint α_i on the edge W : the path $v_i pr_p$ within P separates the interior of $v_i\alpha_i$ from the boundary $C(v_i, v_n)$, and α_i cannot lie on the boundary portion $C(v_1, v_{i-1})$, since this would imply that points interior to $v_{i-1}v_i$ cannot see W (since the chord $v_i\alpha_i$ would then separate $v_{i-1}v_i$ from W). This implies that the out-degree of $e_{i-1} = v_{i-1}v_i$ in G is zero, so in the path cover, the vertex $v_i = b_{\pi'}$ must be the base endpoint for the path π' that covers edge e_{i-1} . We conclude that point p is covered by the triangle $b_{\pi'}\ell_{b_{\pi'}}r_{b_{\pi'}}$.

Thus, every point $p \in P$ is covered. \square

The following theorem is a consequence of the lemma and the fact that optimal path covers for a DAG can be computed in polynomial time:

Theorem 5. *For a weakly visible polygon P , there is a polynomial-time algorithm to compute a set of at most $2k$ convex polygons within P that cover P , where k is the size of an optimal path cover of G .*

C. A Hidden Set from an Antichain

Let I be an antichain in G . We show how to place a set H of $|I|$ points such that no two points from H see each other.

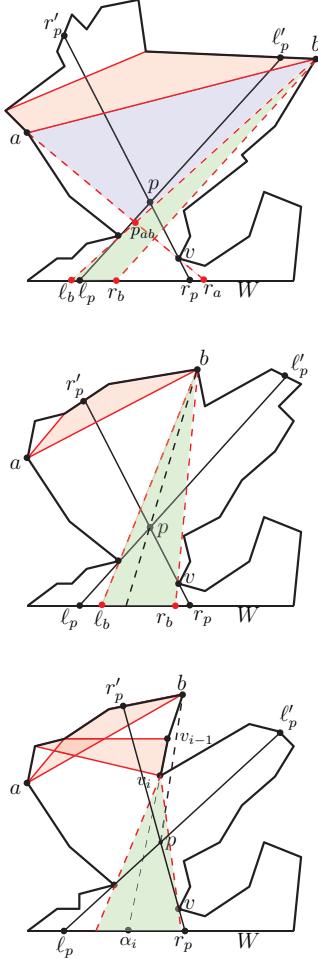


Fig. 3. Proof of the Lemma. Top: Case (1); Middle: Case (2a); and Bottom: Case (2b).

Lemma 6. *Given an antichain I in G , we can compute in polynomial time a set H of $|I|$ points, each interior to one of the edges of I , such that H is a hidden set (no two points of H are visible to each other).*

Proof. Let $I = \{e_{i_1}, e_{i_2}, \dots, e_{i_m}\}$, with $1 \leq i_1 < i_2 < \dots < i_m \leq n - 1$. We will place our hidden set at points $s_\ell \in e_{i_\ell}$ or $t_\ell \in e_{i_\ell}$, for $\ell = 1, \dots, m$ (note that ℓ enumerates the edges in I , not in C – this way we use fewer double subscripts), which are points “close” to the vertices of the edges $e_{i_\ell} \in I$ (but not equal to the vertices), defined as follows. Compute the visibility graph (VG) of P : the graph edges connect mutually visible vertices of P ; in particular, each edge of P is an edge of VG. Extend each edge of VG until the extension hits the boundary of P (and further extension would enter the exterior of P), and let $Y_\ell \subset e_{i_\ell}$, $\ell = 1 \dots m$,

denote the set of points where extended visibility graph edges hit the edge e_{i_ℓ} (Fig. 4, top left). In addition, for all $1 \leq j < k \leq m$, compute shortest paths $SP(v_{i_{j+1}}, v_{i_k})$ between the closest (along C) endpoints of the edges $e_{i_j}, e_{i_k} \in I$; if $SP(v_{i_{j+1}}, v_{i_k})$ has exactly one vertex u between $v_{i_{j+1}}$ and v_{i_k} , and if both edges e_{i_j}, e_{i_k} lie below $v_{i_{j+1}}v_{i_k}$ (Fig. 4, top right), then add to Y_j (resp. Y_k) the points where the line through u parallel to $v_{i_{j+1}}v_{i_k}$ intersects (the supporting line of) e_{i_j} (resp. e_{i_k}) if they exist. On each edge $e_{i_\ell} \in I$ we keep only the two extremal points s'_ℓ, t'_ℓ from Y_ℓ ($s'_\ell \in Y_\ell \cap e_{i_\ell}$ is closest to v_{i_ℓ} , and $t'_\ell \in Y_\ell \cap e_{i_\ell}$ is closest to $v_{i_{\ell+1}}$). Finally, s_ℓ (resp. t_ℓ) is the midpoint of $s'_\ell v_{i_\ell}$ (resp. $t'_\ell v_{i_{\ell+1}}$).

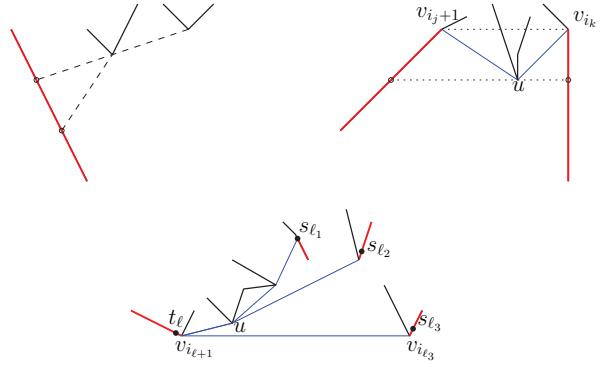


Fig. 4. Defining s_ℓ, t_ℓ on edges from I (red). Top Left: Some of the points added on edge $e \in I$. Top Right: Points added if $SP(v_{i_{j+1}}, v_{i_k}) = v_{i_{j+1}} - u - v_{i_k}$. Bottom: $C(t_\ell, u)$ is visibly obstructed from $C(u, s_{\ell'})$ by vertex u .

Fact 7. *For any $1 \leq \ell < \ell' \leq m$, there exists a vertex $u \in C$ such that no point in $C(t_\ell, u)$ (excluding u) sees any point in $C(u, s_{\ell'})$ (excluding u).*

Proof. The shortest path $SP(v_{i_{\ell+1}}, v_{i_{\ell'}})$ follows VG and makes only left turns at vertices of C [15, Lemma 2]. If the shortest path has internal vertices, then u is any such internal vertex (see $\ell' = \ell_1$ or $\ell' = \ell_2$ in Fig. 4, bottom, for an example). If $v_{i_{\ell+1}}$ and $v_{i_{\ell'}}$ see each other, then at least one of the two edges $e_{i_\ell}, e_{i_{\ell'}}$ (say, e_{i_ℓ}) lies above (the supporting line of) $v_{i_{\ell+1}}v_{i_{\ell'}}$, for otherwise, by the chord property (Fact 1) every point on e_{i_ℓ} would see every point on $e_{i_{\ell'}}$, meaning that the convex hull of the edges would lie within P , so they would be connected by an arc in G . Then $u = v_{i_{\ell+1}}$ (see point s_{ℓ_3} in Fig. 4, bottom for an example) \square

Now, for $|I| = m = 1$ the lemma is trivially true. We prove that for $m \geq 2$ we can choose a set H of m hidden points, one per edge in I , from the following $2m - 2$ points: $t_1 \cup \{s_\ell, t_\ell\}_{\ell=2}^{m-1} \cup s_m$. That is, $t_1, s_m \in H$ (we

do not place hidden points at s_1 or t_m), and each other edge e_ℓ contributes either s_ℓ or t_ℓ .

The proof is by induction on m . The base ($m = 2$) follows from Fact 7: the points t_1 and $s_m = s_2$ are visibly obstructed from each other by the vertex u (refer to Fig. 4, top right). For $m > 2$, also apply Fact 7, and let u be a vertex that (visually) separates $C(t_1, u)$ from $C(u, s_m)$. Let $I_1 \subset I$ be the edges of I that lie in $C(v_{i_1+1}, u)$, together with the edge e_{i_1} ; note that I_1 is a proper subset of I ($|I_1| < m$) because the last edge e_{i_m} is not in I_1 , and thus we can compute the hidden set H_1 (with $|H_1| = |I_1|$) for it by the inductive hypothesis. Similarly, compute the hidden set H_2 (with $|H_2| = |I_2|$) for the edges $I_2 = I \setminus I_1$ that lie in $C(u, v_{i_m})$ plus the edge e_{i_m} . By Fact 7, the union $H_1 \cup H_2$ is the required hidden set H with $|H| = |H_1| + |H_2| = |I|$. \square

See Fig. 5 for an example illustrating an antichain of edges, the convex polygons P'_π , a hidden set on ∂P , and a convex cover of P , for a weakly visible polygon P .

D. Resulting Approximations for $cc(P)$ and $hs(P)$ in Weakly Visible Polygons

Theorem 8. *For a weakly visible polygon P , a convex cover B such that $|B| \leq 2 \cdot cc(P)$ (i.e., a 2-approximation) can be found in polynomial time.*

Proof. Theorem 5 shows that we can find a convex cover, B , which has at most twice as many pieces as the size of the minimum path cover of the poset of the edges of P . Let k be this size, so at most $2k$ pieces. From Lemma 6, we know there exists a hidden set H that is at least as large as the longest antichain in the poset of the edges of P . By Dilworth's theorem [9, Theorem 1.1], we know that size of the longest antichain is equal to the size of the minimum path cover, therefore $|H| = k$. This gives us the following inequality:

$$k = |H| \leq hs(P) \leq cc(P) \leq |B| = 2k \quad (2)$$

which implies that the convex cover found in Theorem 5 is a 2-approximation. \square

Theorem 9. *For a weakly visible polygon P , a hidden set H such that $|H| \geq \frac{1}{2} \cdot hs(P)$ (i.e., a 1/2-approximation) can be found in polynomial time.*

Proof. The chain of inequalities (2) implies $k = |H| \geq \frac{1}{2} \cdot hs(P)$. \square

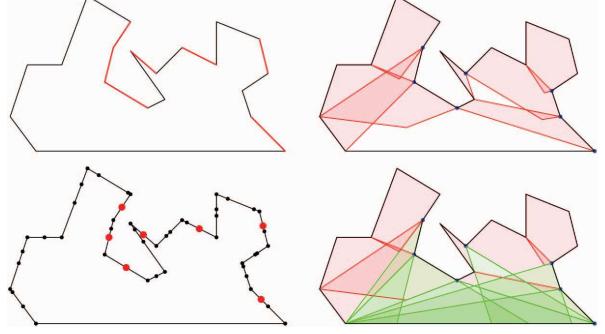


Fig. 5. An example weakly visible polygon, P , in which we show a hidden set and a convex cover. Top Left: An antichain in G . Top Right: All P'_π for a path cover in G . Bottom Left: A hidden set in P . Bottom Right: A convex cover of P .

IV. GENERAL SIMPLE POLYGONS

In this section, P is an arbitrary simple polygon, with n vertices. We utilize concepts related to “link distance” within P , so we begin with a brief review of these concepts and terminology.

A *minimum-link path* between points s, t in P is an $s-t$ path with a minimum number of edges (links); that number is the *link distance* between s and t . Algorithms for computing link distance [14], [21], [30] employ the “staged illumination” paradigm (see, e.g., the handbooks [28, Chapter 12] and [32, Chapter 31.3]): at the first stage, a light source at s illuminates the visibility polygon of s – this is the set of points with link distance 1 from s ; at the beginning of any subsequent stage, the boundary between the illuminated and the dark portions of P consists of a set of *windows*, each being a segment (a chord of P), with one endpoint being a reflex vertex of P and the other endpoint on ∂P , which bounds the weak visibility polygon of the region illuminated at the previous stage (we assume that the window is part of the cell that created it, i.e., that the illuminated region is closed while the dark region does not include its boundary with the illuminated region); see Fig. 6 (top).

The *link distance map*, denoted $LDM(s)$, with respect to the *source* point $s \in P$ is the decomposition of P into cells such that the link distance from s to any point within one cell is the same. The LDM is a by-product of the staged illumination: the edges of the map are the windows (because the windows are pairwise-disjoint, the LDM is also called a “window partition” [4], [31]). Each cell of the map is labeled with the link distance of its points to s . The cells are naturally organized into a (dual) tree T , so that the path from s arrives to a cell through the window from its parent. Further, the windows (and

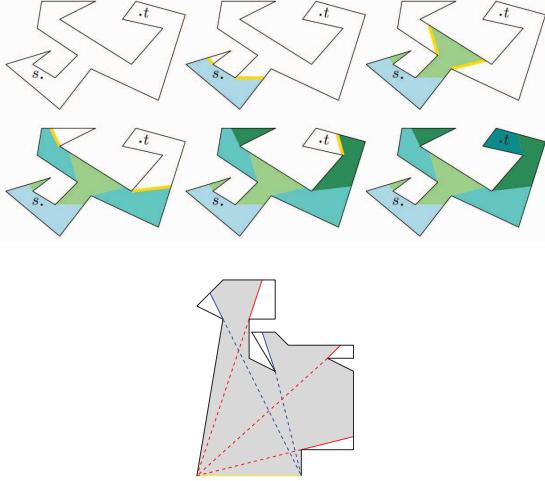


Fig. 6. Top: Staged illumination in a simple polygon (Figure 3 from [26]): The windows are yellow. Bottom: Right (red) and left (blue) windows of the weak visibility polygon (gray) of the yellow window.

hence the cells) can be classified as being *left* or *right* [3], depending on whether the child face is on the left or right of the window (Fig. 6, bottom). Let R_1 be the left cells on even levels of T ; let R_2, R_3, R_4 be the right cells on even levels, left cells on odd levels, and right cells on odd levels, respectively.

We take an arbitrary convex vertex s of P and compute $\text{LDM}(s)$, which can be done in $O(n)$ time [31]. The following properties of the LDM are important for us:

Lemma 10. (1) Any convex polygon $K \subseteq P$ intersects at most 3 cells of $\text{LDM}(s)$.
 (2) For any $i = 1, 2, 3, 4$, no point in one cell from R_i sees a point in another cell from R_i , i.e., if $f, f' \in R_i$ are two cells and $p \in f, p' \in f'$ are two points in the cells, then p does not see p' .
 (3) Every cell of $\text{LDM}(s)$ is a weakly visible polygon.

Proof. (1) Among the cells of LDM intersected by K , let f be a face with minimum associated link distance from s ; let ℓ denote this distance (the number of links in a minimum-link path from s to any point in f). Then K cannot intersect the parent face of f (since its link distance is $\ell - 1$). Since K is convex, it can intersect at most one left window and at most one right window of f , for otherwise the intersection of K and the supporting line of a window would consist from more than 1 connected component. Further, K cannot intersect any grandchild face, f' , of f , as this would imply, by the convexity of K , that there is a point in f' at link distance $\ell + 1$ from s —contradiction to the fact that the distance

from s to points in f' is $\ell + 2$. Thus, K intersects at most 3 faces: f , at most 1 left child of f , and at most 1 right child of f .

(2) This is proved in [3] with respect to the vertices but applies more generally to all points in P ; the proof is similar to the proof of (1) above: the segment between mutually visible points spans at most two levels of the tree T (as level is equal to link distance from s), and points in two right cells (or two left cells) of the same level cannot see each other. Therefore, no point in a cell of R_i can see a point in a different cell of R_i for any $i = 1, 2, 3, 4$.

(3) By definition, any cell is what is seen from a window w (the cell illuminated at stage 1 is the visibility polygon of s , which may be viewed as a degenerate, length-0 window) and is thus weakly visible from w . A cell illuminated at any stage $k > 1$, is on one side of the window and is thus a weakly visible polygon. Stage 1 starts from a convex vertex s ; hence, one can draw a line through s so that P is on one side of the line: s is thus a degenerate edge whose visibility polygon is on one side of the supporting line of the edge. \square

Lemma 10(1) implies that if we separately cover each cell of $\text{LDM}(s)$ with convex polygons, we lose only a factor of 3. Lemma 10(2) implies that we can separately find the hidden sets in each cell of the LDM (our hidden set algorithm from Lemma 6 places hidden point only in the relative interior of the edges of the chain C of a weakly visible polygon) and then choose the largest among the hidden sets in R_1, R_2, R_3, R_4 , losing only a factor of 4 (this is the same idea that was used in [3] to give a 1/4-approximation to the maximum hidden set of vertices). Combining these with Lemma 10(3) and the constant-factor approximations for convex cover and hidden set for weakly visible polygons (Section III), we obtain our main result:

Theorem 11. A 6-approximate convex cover and a 1/8-approximate hidden set in a simple polygon can be found in polynomial time.

Our algorithms run in $O(n^{2+o(1)})$ time.

- Computing the visibility graph VG of P and determining where the extensions of the visibility graph edges intersect edges of P takes $O(|\text{VG}|) = O(n^2)$ time [20].
- Minimum path cover of a DAG and a largest antichain in the poset can be found in $O(n^{2+o(1)})$ time: the problems reduce to maximum matching in a bipartite graph [13], which can be found by

computing a maximum flow, for which the fastest known algorithm [8] runs in $O(n^{2+o(1)})$ time.

- Finally, computing the $O(n)$ shortest paths for placing hidden points takes $O(n^2)$ time overall; a shortest path in a simple polygon can be computed in linear time [16], [32, Chapter 31].

Note that we do convex cover and hidden set separately in every cell of the LDM: the work is charged to the complexity of each cell, and the total complexity of the cells is $O(n)$.

Remark: We can improve the running time for computing $hs(P)$ to $O(n^2)$, with a slightly different approach: using arguments as we did for placing points s_ℓ and t_ℓ interior to edges of P , we can obtain, in time $O(n^2)$ a set of $O(n)$ points on the boundary of P that form a sufficient set for searching for a hidden set of the same size as an antichain. Then, considering these boundary points as vertices of the polygon P , we can apply the quadratic time algorithm of [15] to compute an optimal hidden subset of vertices within a weakly visible polygon.

V. CONCLUSION

We gave the first constant-factor approximation algorithms for convex cover and hidden set in simple polygons. As a by-product of our algorithms, we obtain a combinatorial result (confirming a conjecture from [6]) that $cc(P) \leq 8hs(P)$ for a simple polygon P ; for weakly visible simple polygons, we establish that $cc(P) \leq 2hs(P)$. These combinatorial bounds may be of independent interest; improving them or demonstrating their tightness is an open problem.

Perhaps the most intriguing open problem is whether our techniques can be extended to find an approximately optimal cover with star-shaped polygons, also known as the guarding problem or the Art Gallery Problem. One stumbling block is devising a lower bound stronger than the “witness number” of P (the maximum number of points having pairwise-disjoint visibility polygons): contrary to the inequality $hs(P) \leq cc(P) \leq 8hs(P)$ established in this paper, it is easy to provide examples in which the ratio of the number of guards to witnesses reaches $\Omega(n)$. Nevertheless, our results may be encouraging in the sense that $\exists\mathbb{R}$ -completeness does not preclude approximation.

In polygons with holes, maximum hidden set cannot be n^ε -approximated for some $\varepsilon > 0$, unless $P = NP$ [11]; thus, our methods do not extend to approximating convex cover in polygons with holes. The only known lower bound for the problem is APX-hardness and the best approximation ratio remains logarithmic [12].

Finally, an obvious open question is improving the approximation ratios. We believe that there are two possible fronts to achieve this, either by placing hidden points in the interior of the weakly visible polygons or by showing that only a fraction of the additional triangles are needed for the convex cover (note that, as can be seen from Fig. 2, just taking maximal extensions of our polygons P_π is not enough to cover P). It may also be interesting to improve the running time of our solutions or to give computational lower bounds.

Acknowledgements

We thank the anonymous reviewers for their helpful comments. JM is partially supported by the National Science Foundation (CCF-2007275). VP is partially supported by the Swedish Research Council and the Swedish Transport Administration.

REFERENCES

- [1] Mikkel Abrahamsen. Covering polygons is even harder. *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 375–386, 2022.
- [2] Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. The art gallery problem is NP -complete. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018*, page 65–73, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3188745.3188868.
- [3] Carlos Alegria, Pritam Bhattacharya, and Subir Kumar Ghosh. A 1/4-approximation algorithm for the maximum hidden vertex set problem in simple polygons. In *European Workshop on Computational Geometry*, 2019.
- [4] Esther M. Arkin, Joseph S. B. Mitchell, and Subhash Suri. Logarithmic-time link path queries in a simple polygon. *International Journal of Computational Geometry & Applications*, 5(04):369–395, 1995.
- [5] Édouard Bonnet and Tillmann Miltzow. An approximation algorithm for the art gallery problem. *arXiv preprint arXiv:1607.05527*, 2016.
- [6] Reilly Browne and Eric Chiu. Collapsing the hidden-set convex-cover inequality. In *Proceedings of the 38th Computational Geometry Young Researchers Forum (CG:YRF 2022)*, pages 27–32. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2022.
- [7] Bernard Chazelle and David P Dobkin. Optimal convex decompositions. In *Machine Intelligence and pattern recognition*, volume 2, pages 63–133. Elsevier, 1985.
- [8] Li Chen, Rasmus Kyng, Yang P. Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 612–623, 2022. doi:10.1109/FOCS54457.2022.00064.
- [9] R. P. Dilworth. A decomposition theorem for partially ordered sets. *Ann. of Math.*, 51(1):161–166, 1950.
- [10] Alon Efrat and Sariel Har-Peled. Guarding galleries and terrains. *Information Processing Letters*, 100(6):238–245, 2006. URL: <https://www.sciencedirect.com/science/article/pii/S0020019006001359>, doi:https://doi.org/10.1016/j.ipl.2006.05.014.

[11] Stephan Eidenbenz. Inapproximability of finding maximum hidden sets on polygons and terrains. *Computational Geometry*, 21(3):139–153, 2002. URL: <https://www.sciencedirect.com/science/article/pii/S0925772101000293>, doi:[https://doi.org/10.1016/S0925-7721\(01\)00029-3](https://doi.org/10.1016/S0925-7721(01)00029-3).

[12] Stephan J. Eidenbenz and Peter Widmayer. An approximation algorithm for minimum convex cover with logarithmic performance guarantee. *SIAM Journal on Computing*, 32(3):654–670, 2003. arXiv:<https://doi.org/10.1137/S0097539702405139>, doi:[10.1137/S0097539702405139](https://doi.org/10.1137/S0097539702405139).

[13] Delbert R Fulkerson. Note on dilworth's decomposition theorem for partially ordered sets. *Proceedings of the American Mathematical Society*, 7(4):701–702, 1956.

[14] Subir Kumar Ghosh. Computing the visibility polygon from a convex set and related problems. *Journal of Algorithms*, 12(1):75–95, 1991.

[15] Subir Kumar Ghosh, Anil Maheshwari, Sudebkumar Prasant Pal, Sanjeev Saluja, and C.E. Veni Madhavan. Characterizing and recognizing weak visibility polygons. *Computational Geometry*, 3(4):213–233, 1993. URL: <https://www.sciencedirect.com/science/article/pii/0925772193900104>, doi:[https://doi.org/10.1016/0925-7721\(93\)90010-4](https://doi.org/10.1016/0925-7721(93)90010-4).

[16] Leonidas J. Guibas, John Hershberger, Daniel Leven, Micha Sharir, and Robert Endre Tarjan. Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons. *Algorithmica*, 2:209–233, 1987. doi:[10.1007/BF01840360](https://doi.org/10.1007/BF01840360).

[17] Waldo Gálvez, Arindam Khan, Mathieu Mari, Tobias Mömke, Madhusudhan Reddy Pittu, and Andreas Wiese. *A 3-Approximation Algorithm for Maximum Independent Set of Rectangles*, pages 894–905. Society for Industrial and Applied Mathematics, 2022. URL: <https://pubs.siam.org/doi/abs/10.1137/1.9781611977073.38>, arXiv:<https://pubs.siam.org/doi/pdf/10.1137/1.9781611977073.38>, doi:[10.1137/1.9781611977073.38](https://doi.org/10.1137/1.9781611977073.38).

[18] Johan Hastad. Clique is hard to approximate within $n^{\sup 1/\text{spl epsiv}}$. In *Proceedings of 37th Conference on Foundations of Computer Science*, pages 627–636. IEEE, 1996.

[19] Simon Hengeveld and Tillmann Miltzow. A practical algorithm with performance guarantees for the art gallery problem, 2023. arXiv:[2007.06920](https://arxiv.org/abs/2007.06920).

[20] John Hershberger. An optimal visibility graph algorithm for triangulated simple polygons. *Algorithmica*, 4(1):141–155, 1989. doi:[10.1007/BF01553883](https://doi.org/10.1007/BF01553883).

[21] John Hershberger and Jack Snoeyink. Computing minimum length paths of a given homotopy class. *Computational geometry*, 4(2):63–97, 1994.

[22] Stefan Hertel and Kurt Mehlhorn. Fast triangulation of simple polygons. In *Foundations of Computation Theory: Proceedings of the 1983 International FCT-Conference Borgholm, Sweden, August 21–27, 1983 4*, pages 207–218. Springer, 1983.

[23] Mark Keil and Jack Snoeyink. On the time bound for convex decomposition of simple polygons. *International Journal of Computational Geometry & Applications*, 12(03):181–192, 2002. arXiv:<https://doi.org/10.1142/S0218195902000803>, doi:[10.1142/S0218195902000803](https://doi.org/10.1142/S0218195902000803).

[24] Subhash Khot. Improved inapproximability results for maxclique, chromatic number and approximate graph coloring. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 600–609. IEEE, 2001.

[25] Joseph S. B. Mitchell. Approximating maximum independent set for rectangles in the plane. *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 339–350, 2021.

[26] Joseph S. B. Mitchell, Valentin Polishchuk, and Mikko Sysikaski. Minimum-link paths revisited. *Computational Geometry*, 47(6):651–667, 2014. Special issue on EuroCG'11.

[27] T. Pavlidis. Analysis of set patterns. *Pattern Recognition*, 1(2):165–178, 1968. URL: <https://www.sciencedirect.com/science/article/pii/003132036890006X>, doi:[https://doi.org/10.1016/0031-3203\(68\)90006-X](https://doi.org/10.1016/0031-3203(68)90006-X).

[28] Jörg-Rüdiger Sack and Jorge Urrutia. *Handbook of computational geometry*. Elsevier, 1999.

[29] Jack Stade. The point-boundary art gallery problem is er-hard. *arXiv preprint arXiv:2210.12817*, 2023.

[30] 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.

[31] Subhash Suri. On some link distance problems in a simple polygon. *IEEE transactions on Robotics and Automation*, 6(1):108–113, 1990.

[32] Csaba D Toth, Joseph O'Rourke, and Jacob E Goodman. *Handbook of discrete and computational geometry*. CRC press, 2017.