

Polygon Simplification by Minimizing Convex Corners

Yeganeh Bahoo^a, Stephane Durocher^a, J. Mark Keil^b, Debajyoti Mondal^b,
Saeed Mehrabi^c, Sahar Mehrpour^d

^a*Department of Computer Science, University of Manitoba, Winnipeg, Canada.*

^b*Department of Computer Science, University of Saskatchewan, Saskatoon, Canada.*

^c*School of Computer Science, Carleton University, Ottawa, Canada.*

^d*School of Computing, University of Utah, Utah, USA.*

Abstract

Let P be a polygon with $r > 0$ reflex vertices and possibly with holes and islands (i.e., components of the polygon inside holes). A subsuming polygon of P is a polygon P' such that $P \subseteq P'$, each connected component R of P is a subset of a distinct connected component R' of P' , and the reflex corners of R coincide with those of R' . A subsuming chain of P' is a minimal path on the boundary of P' whose two end edges coincide with two edges of P . Aichholzer et al. proved that every polygon P has a subsuming polygon with $O(r)$ vertices, and posed an open problem to determine the computational complexity of computing subsuming polygons with the minimum number of convex vertices.

In this paper, we prove that the problem of computing an optimal subsuming polygon is NP-hard (but the complexity remains open for simple polygons; i.e., polygons without holes). Our NP-hardness result holds even when the subsuming chains are restricted to have constant length and lie on the arrangement of lines determined by the edges of the input polygon. We show that this restriction makes the problem polynomial-time solvable for simple polygons.

Keywords:

2010 MSC: 68Q25, 68U05

1. Introduction

Polygon simplification is well studied in computational geometry, with numerous applications in cartographic visualization, computer graphics and data compression [2, 3]. Techniques for simplifying polygons and polylines have appeared in the literature in various forms. Common goals of these simplification

[★]A preliminary version of this paper appeared in proceedings of the 22nd International Computing and Combinatorics Conference (COCOON 2016) [1].

Email addresses: bahoo@cs.umanitoba.ca (Yeganeh Bahoo),
durocher@cs.umanitoba.ca (Stephane Durocher), keil@cs.usask.ca (J. Mark Keil),
dmondal@cs.usask.ca (Debajyoti Mondal), saeed.mehrabi@carleton.ca (Saeed Mehrabi),
mehrpour@cs.utah.edu (Sahar Mehrpour)




Figure 1: (a) A polygon P , where the polygon is filled and the holes are empty regions. (b) A subsuming polygon P' , where P' is the union of the filled regions. A subsuming chain is shown in bold; notice that the end edges of this chain overlap with two edges of P . (c) A min-convex subsuming polygon P'_{min} , where $\mathcal{A}_e(P'_{min}) = \mathcal{A}_e(P)$. (d) A polygon P such that for any min-convex subsuming polygon P'_{min} , $\mathcal{A}_e(P) \neq \mathcal{A}_e(P'_{min})$.

algorithms include to preserve the shape of the polygon, to reduce the number of vertices, to reduce the space requirements, and to remove noise (extraneous bends) from the polygon boundary (e.g., [4, 5, 6]). In this paper we consider a specific version of polygon simplification introduced by Aichholzer et al. [7], which keeps reflex corners intact, but minimizes the number of convex corners. Aichholzer et al. showed that such a simplification can help achieve faster solutions for many geometric problems such as answering shortest path queries (as shortest paths stay the same), computing Voronoi diagrams, and so on.

A simple polygon is a connected region without holes. Let P be a polygon with r reflex vertices and possibly with holes and islands. An *island* is a simple polygon that lies entirely inside a hole. A *reflex corner* of P consists of three consecutive vertices u, v, w on the boundary of P such that the angle $\angle uvw$ inside P is more than 180° . We refer the vertex v as a *reflex* vertex of P . The vertices of P that are not reflex are called *convex* vertices. By a *component* of P , we refer to a maximally connected region of P . A polygon P' *subsumes* P if $P \subseteq P'$, each component R' of P' contains a distinct component R of P (i.e., $R \subseteq R'$), and the reflex corners of R coincide with the reflex corners of R' . A *k-convex subsuming polygon* P' contains at most k convex vertices. A *min-convex subsuming polygon* is a subsuming polygon that minimizes the number of convex vertices. Figure 1(a) illustrates a polygon P , and Figures 1(b) and (c) illustrate a subsuming polygon and a min-convex subsuming polygon of P , respectively. A *subsuming chain* of P' is a minimal path (of length at least 2) on the boundary of P' whose end edges overlap with a pair of edges of P ; see Figure 1(b) for an example.

Aichholzer et al. [7] showed that for every polygon P with n vertices, $r > 0$ of which are reflex, one can compute in linear time a subsuming polygon P' with at most $O(r)$ vertices. Note that although a subsuming polygon with $O(r)$ vertices always exists, no polynomial-time algorithm is known for computing a min-convex subsuming polygon. Finding an optimal subsuming polygon seems challenging since it does not always lie on the arrangement of lines $\mathcal{A}_e(P)$ (resp., $\mathcal{A}_v(P)$) determined by the edges (resp., pairs of vertices) of the input polygon.

Figure 1(c) illustrates an optimal polygon P'_{min} for the polygon P of Figure 1(a), where $\mathcal{A}_e(P'_{min}) = \mathcal{A}_e(P)$. On the other hand, Figure 1(d) shows that a min-convex subsuming polygon may not always lie on $\mathcal{A}_e(P)$ or $\mathcal{A}_v(P)$. Note that the input polygon of Figure 1(d) is a *simple polygon*, i.e., it does not contain any hole. Hence determining min-convex subsuming polygons seems challenging even for simple polygons. In fact, Aichholzer et al. [7] posed an open question that asks to determine the complexity of computing min-convex subsuming polygons, where the input is restricted to simple polygons.

1.1. Our Contribution

In this paper we show that the problem of computing a min-convex subsuming polygon is NP-hard for polygons possibly with holes (Section 2). We noted earlier that discretizing the solution space is a potential challenge, i.e., that the optimal polygon may not always lie on the line arrangement determined by the input polygon (Figure 1(d)). Interestingly, our NP-hardness result does not seem to utilize this challenge, instead, the hardness holds even when we restrict the subsuming chains to have constant length and to lie on $\mathcal{A}_e(P)$.

A question that naturally appears in this context is whether such restrictions on subsuming chains can make the problem easier for nontrivial classes of polygons. For example, consider an x -monotone polygon³ P , e.g., see Figure 2(a). Then it is not difficult to see that there exists a min-convex polygon such that each subsuming chain has constant length and lies on $\mathcal{A}_e(P)$. The argument is simple except for the subsuming chains that covers the two ends of P , e.g., see Figure 2(b). A simple proof is included in Section 3.




Figure 2: (a) An x -monotone polygon P . The dashed lines illustrate the monotonicity. (b) A min-convex subsuming polygon of P . The dashed lines illustrate the separation between successive chains.

We then show that the question can be answered affirmatively for arbitrary simple polygons, i.e., for any simple polygon P , one can compute in polynomial time, a min-convex subsuming polygon P_{min} under the restriction that the subsuming chains are of constant length and lie on $\mathcal{A}_e(P)$.

³ P is x -monotone if every vertical line intersects P at most twice.

1.2. Organization

The rest of the paper is organized as follows. Section 2 presents the NP-hardness result for polygons with holes. Section 3 presents our observations on monotone polygons. Section 4 describes the techniques for computing subsuming polygons for simple polygons. Finally, Section 5 concludes the paper discussing directions to future research.

2. NP-hardness of Min-Convex Subsuming Polygon

In this section we prove that it is NP-hard to find a subsuming polygon with minimum number of convex vertices. We denote the problem by MIN-CONVEX-SUBSUMING-POLYGON. We reduce the NP-complete problem monotone planar 3-SAT [8], which is a variation of the 3-SAT problem as follows: Every clause in a monotone planar 3-SAT consists of either three negated variables (*negative clause*) or three non-negated variables (*positive clause*). Furthermore, the bipartite graph constructed from the variable-clause incidences, admits a planar drawing such that all the vertices corresponding to the variables lie along a horizontal straight line l , and all the vertices corresponding to the positive (respectively, negative) clauses lie above (respectively, below) l . The problem remains NP-hard even when each variable appears in at most four clauses [9].

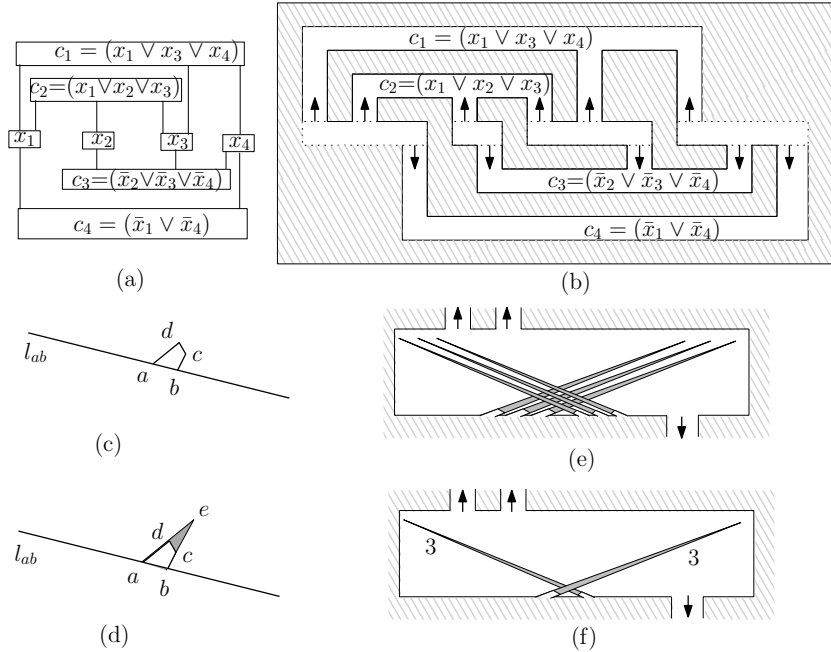


Figure 3: (a) An instance I of monotone planar 3-SAT. (b) The orthogonal polygon P_o corresponding to I . (c)–(f) Illustration for the variable gadget.

The idea of the reduction is as follows. Given an instance of a monotone planar 3-SAT I with variable set X and clause set C , we create a corresponding instance \mathcal{P}_I of MIN-CONVEX-SUBSUMING-POLYGON. Let λ be the number of convex vertices in \mathcal{P}_I . The reduction ensures that if there exists a satisfying truth assignment of I , then \mathcal{P}_I can be subsumed by a polygon with at most $\lambda - |X||C|^2 - 3|C|$ convex vertices, and vice versa.

Given an instance I of monotone planar 3-SAT, we first construct an orthogonal polygon P_o with holes. We denote each clause and variable using a distinct axis-aligned rectangle, which we refer to as the *c-rectangle* and *v-rectangle*, respectively. Each edge connecting a clause and a variable is represented as a thin vertical strip, which we call an *edge tunnel*. Figures 3(a) and (b) illustrate an instance of monotone planar 3-SAT and the corresponding orthogonal polygon, respectively. While adding the edge tunnels, we ensure for each v-rectangle that the tunnels coming from the top lie to the left of all the tunnels coming from the bottom. Figure 3(b) marks the top and bottom edge tunnels by upward and downward rays, respectively. The v-rectangles, c-rectangles and the edge tunnels may form one or more holes, as it is shown by diagonal line pattern in Figure 3. We now transform P_o to an instance \mathcal{P}_I of MIN-CONVEX-SUBSUMING-POLYGON.

We first introduce a few notations. Let $abcd$ be a convex quadrangle and let l_{ab} be an infinite line that passes through a and b . Assume also that l_{bc} and l_{ad} intersect at some point e , and c, d, e all lie above l_{ab} , as shown in Figures 3(c)–(d). Then, we call the quadrangle $abcd$ a *tip* on l_{ab} , and the triangle cde a *cap* of $abcd$. The tip $abcd$ is called *top-right* tip if the slope of ad is positive; otherwise, it is called a *top-left* tip. If the points c, d and e lie below l_{ab} , then we call $abcd$ a *down-right* tip (resp., *down-left* tip) if the slope of ad is negative (resp., positive).

2.1. Variable Gadget

We construct variable gadgets from the v-rectangles. We add some top-right (and the same number of top-left) tips at the bottom side of the v-rectangle, as shown in Figure 3(e). There are three top-right and top-left tips in the figure. For convenience we show only one top-left and one top-right tip in the schematic representation, as shown in Figure 3(f). However, we assign weight to these tips to denote how many tips there should be in the exact construction. We will ensure a few more properties: (I) The caps do not intersect the boundary of the v-rectangle, (II) no two top-left caps (or, top-right caps) intersect, and (III) every top-left (resp., top-right) cap intersects all the top-right (resp., top-left) caps.

Observe that each top-left tip contributes to two convex vertices such that covering them with a cap reduces the number of convex vertices by 1. The peak of the cap reaches very close to the top-left corner of the v-rectangle, which will later interfere with the clause gadget. Specifically, this cap will intersect any downward cap of the clause gadget coming through the top edge tunnels. Similarly, each top-right tip contributes to two convex vertices, and

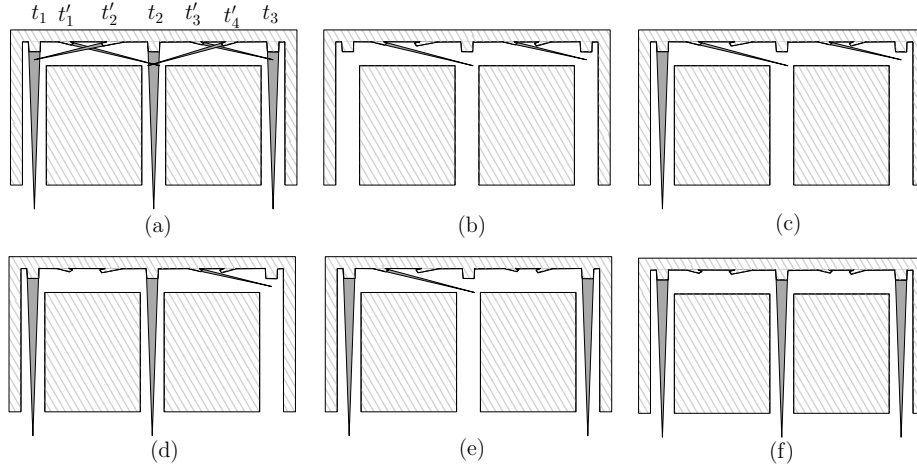


Figure 4: Illustration for the clause gadget.

the corresponding cap intersects any upward cap coming through the bottom edge tunnels.

Note that the optimal subsuming polygon P cannot contain the caps from both the top-left and top-right tips. We assign the tips with a weight of $|C|^2$. In the hardness proof this will ensure that either the caps of top-right tips or the caps of top-left tips must exist in P , which will correspond to the true and false configurations, respectively.

2.2. Clause Gadget

Recall that, by definition, each clause consists of three variables and so it is incident to three edge tunnels. Figure 4(a) illustrates the transformation for a c -rectangle. Here we describe the gadget for the positive clauses, and the construction for negative clauses is symmetric. We add three downward tips incident to the top side of the c -rectangle, along its three edge tunnels. Each of these downward tip contributes to two convex vertices such that covering the tip with a cap reduces the number of convex vertices by 1. Besides, the corresponding caps reach almost to the bottom side of the v -rectangles, i.e., they would intersect the top-left caps of the v -rectangles. Let these tips be t_1, t_2, t_3 from left to right, and let $\gamma_1, \gamma_2, \gamma_3$ be the corresponding caps.

We then add a down-left and a down-right tip at the top side of the c -rectangle between t_i and t_{i+1} , where $1 \leq i \leq 2$, as shown in Figure 4(a). Let the tips be t'_1, \dots, t'_4 from left to right, and let the corresponding caps be $\gamma'_1, \dots, \gamma'_4$. Note that the caps corresponding to t'_j and t'_{j+1} , where $j \in \{1, 3\}$, intersect each other. Therefore, at most two of these four caps can exist at the same time in the solution polygon. Observe also that the caps corresponding to t_1, t_2, t_3 intersect the caps corresponding to $\{t'_2\}, \{t'_1, t'_4\}, \{t'_3\}$, respectively. Consequently, any optimal solution polygon containing none of $\{\gamma_1, \gamma_2, \gamma_3\}$ has

at least 12 convex vertices along the top boundary of the c-rectangle, as shown in Figure 4(b).

We now show that any optimal solution polygon P containing at least $\alpha > 0$ caps from $\Gamma = \{\gamma_1, \gamma_2, \gamma_3\}$ have exactly 11 convex vertices along the top boundary of the c-rectangle. We consider the following three cases:

Case 1 ($\alpha = 1$): If γ_1 (resp., γ_3) is in P , then P must contain $\{\gamma'_1, \gamma'_3\}$ (resp., $\{\gamma'_2, \gamma'_4\}$). Figure 4(c) illustrates the case when P contains γ_1 . If γ_2 is in P , then P must contain $\{\gamma'_2, \gamma'_3\}$. In all the above scenarios the number of convex vertices along the top boundary of the c-rectangle is 11.

Case 2 ($\alpha = 2$): If P contains $\{\gamma_1, \gamma_3\}$, then either γ'_1 or γ'_4 must be in P . Otherwise, P contains either $\{\gamma_1, \gamma_2\}$ or $\{\gamma_2, \gamma_3\}$. If that P contains $\{\gamma_1, \gamma_2\}$, as in Figure 4(d), then γ'_3 must lie in P . In the remaining case, γ'_2 must lie in P . Therefore, also in this case the number of convex vertices along the top boundary of the c-rectangle is 11.

Case 3 ($\alpha = 3$): In this scenario P cannot contain any of $\gamma'_1, \dots, \gamma'_4$. Therefore, as shown in Figure 4(e), the number of convex vertices along the top boundary of the c-rectangle is 11.

As a consequence we obtain the following lemma.

Lemma 1. *If a clause is satisfied, then any optimal subsuming polygon reduces exactly three convex vertex from the corresponding c-rectangle.*

2.3. Reduction

Although we have already described the variable and clause gadgets, the optimal subsuming polygon still may come up with some unexpected optimization that interferes with the convex corner count in our hardness proof. Figure 5(left) illustrates one such example. Therefore, we replace each convex corner that does not correspond to the tips by a small polyline with alternating convex and reflex corners, as shown Figure 5(right). By construction, it is now straightforward to observe the following fact.

Remark 1. *Let x be a reflex vertex of \mathcal{P}_I , and let x' be the first reflex vertex after x while walking clockwise on the boundary of \mathcal{P}_I starting at x . Then, the number of convex vertices that can appear between x and x' is at most two. Furthermore, if there are two convex vertices, then either they correspond to a tip, or an axis-parallel dent of the form \square or \sqcup (up to rotation).*

We now prove the NP-hardness of computing an optimal subsuming polygon.

Theorem 2. *Finding an optimal subsuming polygon is NP-hard.*

PROOF. Let $I = (X, C)$ be an instance of monotone planar 3-SAT and let \mathcal{P}_I be the corresponding instance of MIN-CONVEX-SUBSUMING-POLYGON. Let λ be the number of convex vertices in \mathcal{P}_I . We now show that I admits a satisfying truth assignment if and only if \mathcal{P}_I can be subsumed by a polygon having at most $\lambda - |X||C|^2 - 3|C|$ convex vertices.




Figure 5: Refinement of \mathcal{P}_I .

First assume that I admits a satisfying truth assignment. For each variable x , we choose either the top-right caps or the top-left caps depending on whether x is assigned true or false. Consequently, we save at least $|X||C|^2$ convex vertices. Consider any clause $c \in C$. Since c is satisfied, one or more of its variables are assigned true. Therefore, for each positive (resp., negative) clause, we can have one or more downward (resp., upward) caps that enter into the v-rectangles. By Lemma 1, we can save at least three convex vertices from each c-rectangle. Therefore, we can find a subsuming polygon with at most $\lambda - |X||C|^2 - 3|C|$ convex vertices.

Assume now that some polygon P with at most $\lambda - |X||C|^2 - 3|C|$ convex vertices can subsume \mathcal{P}_I . We now find a satisfying truth assignment for I . By Remark 1 we can restrict our attention only to c- and v-rectangles. Note that the maximum number of convex vertices that can be reduced from the c-rectangles is at most $3|C|$. Therefore, P must reduce at least $|C|^2$ convex vertices from each v-rectangle. Recall that in each v-rectangle, either the top-right or the top-left caps can be chosen in the solution, but not both. Therefore, the v-rectangles cannot help reducing more than $|X||C|^2$ convex vertices. If P contains the top-right caps of the v-rectangle, then we set the corresponding variable to true, otherwise, we set it to false. Since P has at most $\lambda - |X||C|^2 - 3|C|$ convex vertices, and each c-rectangle can help to reduce at most 3 convex vertices (Lemma 1), P must have at least one cap from $\gamma_1, \gamma_2, \gamma_3$ at each c-rectangle. Therefore, each clause must be satisfied. Recall that the downward (resp., upward) caps coming from edge tunnels are designed carefully to have conflict with the top-left (resp., top-right) caps of v-variables. Since top-left and top-right caps of v-variables are conflicting, the truth assignment of each variable is consistent in all the clauses that contains it. \square

It is straightforward from the construction of \mathcal{P}_I that no optimal subsuming polygon P that subsumes \mathcal{P}_I can have a subsuming chain of length larger than 2, and there always exists an optimal solution that lies on $\mathcal{A}_e(\mathcal{P}_I)$. Hence, Theorem 2 holds even under the restriction that the subsuming chains must be of constant length and lie on $\mathcal{A}_e(\mathcal{P}_I)$. In Sections 3 and 4, we will show that these restrictions make the problem polynomial-time solvable for simple polygons.

3. Monotone Polygons

In this section, we give a straightforward algorithm to compute a min-convex subsuming polygon of x -monotone polygons. In fact, we prove a stronger claim that every x -monotone polygon P admits a min-convex subsuming polygon such that the subsuming chains are of constant length and lie on $\mathcal{A}_e(P)$.

Let $\mathbf{upper}(P)$ and $\mathbf{lower}(P)$ denote the upper and lower chains of P , respectively. Moreover, let u_1 (resp., u_m) be the leftmost (resp., rightmost) vertex of P ; notice that the vertices u_1 and u_m are both convex (as P is x -monotone) and are shared between $\mathbf{upper}(P)$ and $\mathbf{lower}(P)$. Let u_2, \dots, u_{m-1} (resp., $l_2, \dots, l_{m'-1}$) be the set of reflex vertices of P that lie on $\mathbf{upper}(P)$ (resp., on $\mathbf{lower}(P)$); we let $l_1 = u_1$ and $l_{m'} = u_m$. For a reflex vertex u_i , where $2 \leq i < m$, let $\ell^-(u_i)$ (resp., $\ell^+(u_i)$) denote the line determined by the edge of P whose right endpoint (resp., left endpoint) is u_i . Similarly, define the lines $\ell^-(l_i)$ and $\ell^+(l_i)$ for all $2 \leq i < m'$. For u_1 and u_m , only $\ell^+(u_1)$ and $\ell^-(u_m)$ are defined. We next describe the algorithm.

First, consider $\mathbf{upper}(P)$. Initially, let the simplified polygon P' be P . For each reflex vertex u_i on $\mathbf{upper}(P)$, where $2 \leq i < m - 1$, consider the vertical slab defined by the two vertical lines through u_i and u_{i+1} . If there is no convex vertex of $\mathbf{upper}(P)$ in this slab, then the edge $u_i u_{i+1}$ must be an edge of any feasible solution. So, such an edge stays in P' . Otherwise, $\ell^+(u_i)$ and $\ell^-(u_{i+1})$ intersect each other at some point p outside and above P or on P . Then, we remove the chain of convex vertices between u_i and u_{i+1} and add the two edges $u_i p$ and $u_{i+1} p$ to P' ; hence, reducing the number of convex vertices of $\mathbf{upper}(P)$ between u_i and u_{i+1} to one. Next, we consider $\mathbf{lower}(P)$ and apply the same process to every two each vertex l_i on $\mathbf{lower}(P)$, where $2 \leq i < m' - 1$.

It remains to show how to deal with the convex vertices that appear before the leftmost reflex vertex or after the rightmost reflex vertex on each chain. We show that these convex vertices can be reduced to at most two convex vertices, depending on the slopes of the edges incident to such *reflex* vertices. In the following, as the second part of the algorithm, we discuss the details for convex vertices that appear before the leftmost reflex vertices; the convex vertices on the other end of the polygon can be handled similarly.

Consider u_1 (i.e., the leftmost vertex of P) and let π be the chain on the boundary of P that connects u_2 to l_2 in counter-clockwise order (i.e., it contains u_1). To reduce the number of convex vertices on π , it is sufficient to check if $\ell^-(u_2)$ and $\ell^-(l_2)$ intersect at some point p whose x -coordinate is less than that of both u_2 and l_2 . For instance, if the slope of $\ell^-(u_2)$ is positive and the slope of $\ell^-(l_2)$ is negative, then $\ell^-(u_2)$ and $\ell^-(l_2)$ intersect at such point p ; see Figure 6(a). In this case, we can replace π with two edges $u_2 p$ and $l_2 p$; hence, reducing the number of convex vertices on π to one. Therefore, we can simplify π as follows: if $\ell^-(u_2)$ and $\ell^-(l_2)$ intersect at such point p , then we replace π with two edges $u_2 p$ and $l_2 p$ (hence, reducing the number of convex vertices on π to one). Otherwise, if no such point p exists (e.g., when the slope of $\ell^-(u_2)$ is negative, but the slope of $\ell^-(l_2)$ is positive), then both $\ell^-(u_2)$ and $\ell^-(l_2)$ must intersect the line passing through at least one of the edges incident to u_1 . See




Figure 6: How to reduce the convex vertices around the leftmost vertex u_1 of P .

Figure 6(b-c). So, we can replace π with three edges and reducing the convex vertices on π to two. We perform a similar process on the path π corresponding to the rightmost convex vertex of P and its “closest” reflex vertices from each chain on the other end of P . Let P^* be the resulting simplified polygon.

To see the monotonicity of P^* , we note that in each slab considered in the first part of the algorithm, at most two new edges are introduced that lie inside the slab. Therefore, the edges from different slabs are disjoint. Moreover, the edges introduced in the second part of the algorithm (i.e., when dealing with the leftmost and rightmost convex vertices u_1 and u_m) do not violate the monotonicity of P^* .

For every two consecutive reflex vertices, P^* has exactly one convex vertex (resp., has no convex vertex) between them if P had at least one (resp., had none) between them. Since there must be at least one convex vertex between every two reflex vertices that had at least one convex vertex between them, any simplified polygon must have at least as many convex vertices as P^* between every two consecutive reflex vertices. Moreover, one can easily verify that P^* has the minimum number of convex vertices generated in the second part of the algorithm. Therefore, P^* is optimal and we hence have the following result.

Theorem 3. *Given a monotone polygon P with n vertices, a subsuming polygon of P with the minimum number of convex vertices can be computed in $O(n \log n)$ time.*

4. Computing Subsuming Polygons

In this section, we show that for any simple polygon P , a min-convex subsuming polygon P_{min} can be computed in polynomial time under the restriction that the subsuming chains are of constant length and lie on $\mathcal{A}_e(P)$. We first present definitions and preliminary results on outerstring graphs, which will be an important tool for computing subsuming polygons.

4.1. Independent Set in Outerstring Graphs

A graph G is a *string graph* if it is an intersection graph of a set of simple curves in the plane, i.e., each vertex of G is mapped to a curve (string), and two vertices are adjacent in G if and only if the corresponding curves intersect.

G is an *outerstring graph* if the underlying curves lie interior to a simple cycle C , where each curve intersects C at one of its endpoints. Figure 7(a) illustrates an outerstring graph and the corresponding arrangement of curves. Later in our algorithm, the polygon will correspond to the cycle of an outerstring graph, and some polygonal chains attached to the boundary of the polygon will correspond to the strings of that outerstring graph.

A set of strings is called *independent* if no two strings in the set intersect, the corresponding vertices in G are called an independent set of vertices. Let G be a weighted outerstring graph with a set \mathcal{T} of weighted strings. A *maximum weight independent set* $\text{MWIS}(\mathcal{T})$ (resp., $\text{MWIS}(G)$) is a set of independent strings $T \subseteq \mathcal{T}$ (resp., vertices) that maximizes the sum of the weights of the strings in T . By $|\text{MWIS}(G)|$ we denote the weight of $\text{MWIS}(G)$.

Let $\Gamma(G)$ be the arrangement of curves that corresponds to G ; e.g., see Figure 7(a). Let R be a geometric representation of $\Gamma(G)$, where C is represented as a simple polygon P , and each curve is represented as a simple polygonal chain inside P such that one of its endpoints coincides with a distinct vertex of P . Keil et al. [10] showed that given a geometric representation R of G , one can compute a maximum weight independent set of G in $O(s^3)$ time, where s is the number of line segments in R .

Theorem 4 (Keil et al. [10]). *Let G be a weighted outerstring graph. Given a geometric representation R of G , there exists a dynamic programming algorithm that computes a maximum weight independent set of G in $O(s^3)$ time, where s is the number of straight line segments in R .*

Figure 7(b) illustrates a geometric representation R of some G , where each string is represented with at most 4 segments. Keil et al. [10] observed that any maximum weight independent set of strings can be triangulated to create a triangulation P_t of P , and such a triangulation can be used to design a dynamic programming algorithm to compute a maximum weight independent set. See Figure 7(c) for an example of such a triangulation. We next describe very briefly how the triangulation of a maximum independent set can be used.

Let \mathcal{T} be the strings in R . Then, the problem of finding $\text{MWIS}(\mathcal{T})$ can be solved by dividing the problem into subproblems each of which described using only two vertices of R . We illustrate how the subproblems are computed using Figure 7(d). Let $P(v_1, v_2)$ be the problem of finding $\text{MWIS}(\mathcal{T}_{v_1, v_2})$, where \mathcal{T}_{v_1, v_2} consists of the strings that lie wholly to the left of $v_1 v_2$. The authors consider three types of subproblems depending on whether v_1 and v_2 are the vertices of P or the vertices of some strings that lie inside P . The authors show in each of these three cases that a third vertex w (either a vertex of P or a vertex of some string) can be chosen in such a way that $P(v_1, v_2)$ can be computed from the solutions to the subproblems $P(v_1, w)$ and $P(w, v_2)$. For example, in Figure 7(d), w is a vertex on some string d inside $P(v_1, v_2)$ and so $P(v_1, v_2)$ can be computed using the corresponding subproblems as shown in Figure 7(e). Note that if v_1 and v_2 are polygon vertices, then there cannot be more than one triangle on the border of $P(v_1, v_2)$ adjacent to the segment $v_1 v_2$. Otherwise, the

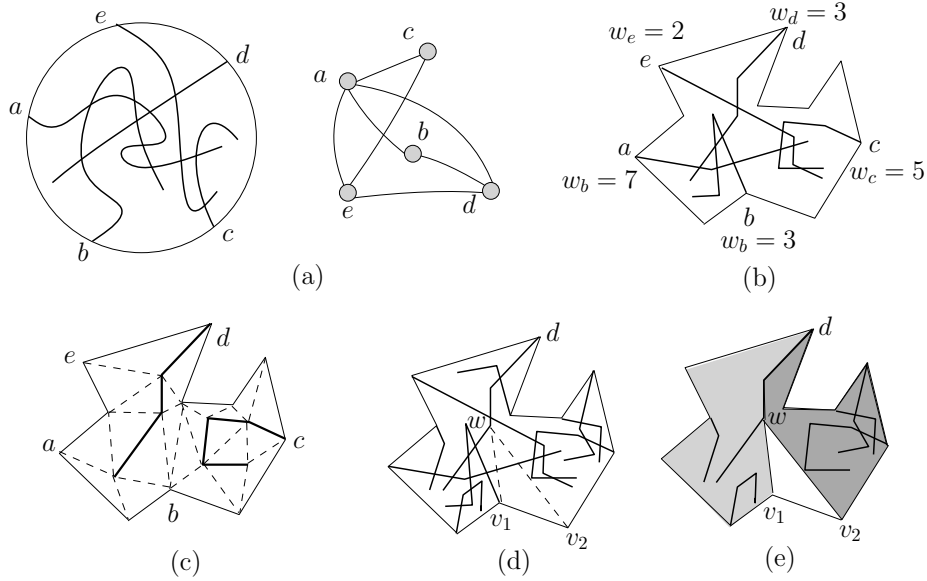


Figure 7: (a) Illustration for G and $\Gamma(G)$. (b) A geometric representation R of G . (c) A triangulated polygon obtained from an independent set of G . (d)–(e) Dynamic programming to find maximum weight independent set.

strings that contain v_1 and v_2 would define the boundary of the subproblem. Therefore, these two vertices suffice, even if there are more triangles in the boundary of $P(v_1, v_2)$ incident to the corresponding strings.

Keil et al. [10] showed that there are only a few different cases depending on whether the vertices describing the subproblems belong to the polygon or to the strings. We will use this idea of computing $\text{MWIS}(\mathcal{T})$ to compute subsuming polygons.

4.2. Subsuming Polygons via Outerstring Graphs

Let $P = (v_0, v_1, \dots, v_{n-1})$ be a simple polygon with n vertices, $r > 0$ of which are reflex vertices. A *convex chain* of P is a path $C_{ij} = (v_i, v_{i+1}, \dots, v_{j-1}, v_j)$ of strictly convex vertices, where the indices are considered modulo n .

Let $P' = (w_0, w_1, \dots, w_{m-1})$ be a subsuming polygon of P , where $\mathcal{A}_e(P') = \mathcal{A}_e(P)$, and the subsuming chains are of length at most t . Here, by “length”, we mean the number of edges (not the Euclidean length). Let $C'_{qr} = (w_q, \dots, w_r)$ be a subsuming chain of P' . Then, by the definition of a subsuming chain, there exist edges (v_i, v_{i+1}) and (v_{j-1}, v_j) in P that coincide with the edges (w_q, w_{q+1}) and (w_{r-1}, w_r) . We call the vertex v_i the *left support* of C'_{qr} . Since $\mathcal{A}_e(P') = \mathcal{A}_e(P)$, the chain C'_{qr} must lie on $\mathcal{A}_e(P)$. Moreover, since P' is a min-convex subsuming polygon, the number of vertices in C'_{qr} would be at most the number of vertices in C_{ij} .

We claim that the number of paths in $\mathcal{A}_e(P)$ from v_i to v_j is at most n^t , where $t = O(1)$ is an upper bound on the length of the subsuming chains. Thus

any subsuming chain can have at most $(t - 1)$ line segments. Since there are only $O(n)$ straight lines in the arrangement $\mathcal{A}_e(P)$, there can be at most n^j paths of j edges, where $1 \leq j \leq t - 1$. Consequently, the number of candidate chains that can subsume C_{ij} is $O(n^t)$.

Lemma 5. *Given a simple polygon P with n vertices, every convex chain C of P has at most $O(n^t)$ candidate subsuming chains in $\mathcal{A}_e(P)$, each of length at most t .*

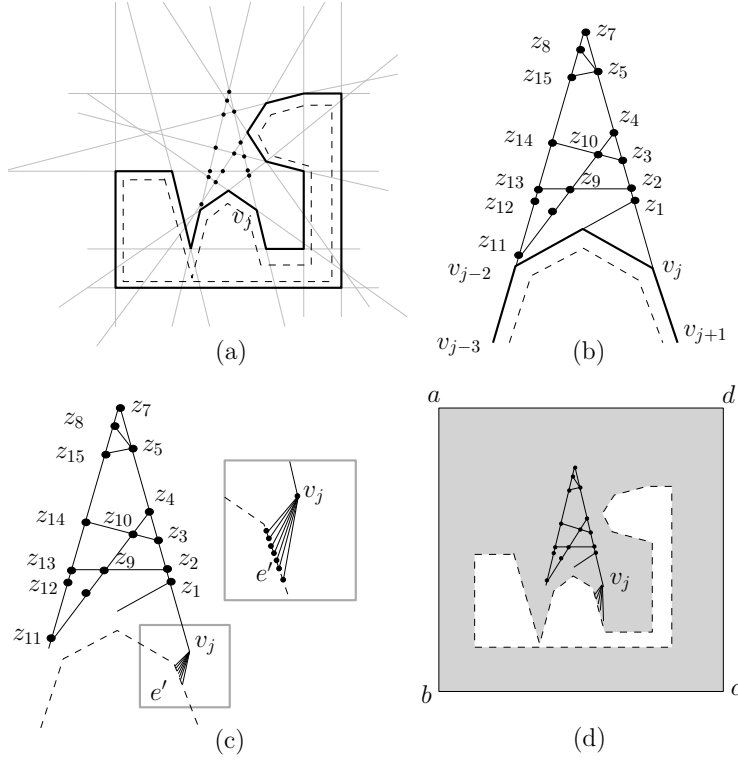


Figure 8: (a) Illustration for the polygon P (in bold), $\mathcal{A}_e(P)$ (in gray), and Q (in dashed lines). (b) Chains of v_j . (c) Attaching the strings to Q . (d) Dynamic programming inside the gray region.

In the following, we construct an outerstring graph using these candidate subsuming chains. We first compute a simple polygon Q interior to P such that for each edge e in P , there exists a corresponding edge e' in Q which is parallel to e and the perpendicular distance between e and e' is ϵ , as shown in dashed line in Figure 8(a). We choose ϵ sufficiently small⁴ such that for

⁴Choose $\epsilon = \delta/3$, where δ is the distance between the closest visible pair of boundary points.

each component w of P , Q contains exactly one component inside w . We now construct the strings. Let v_j be a convex corner of P . Let S_j be the set of candidate subsuming chains such that each chain contains (v_j, v_{j+1}) as the last edge while traversing the chain starting at its left support. For example, the subsuming chains that correspond to v_j are (v_{j-2}, z_1, v_{j+1}) , $(v_{j-3}, z_{13}, z_2, v_{j+1})$, $(v_{j-3}, z_{14}, z_3, v_{j+1})$, $(v_{j-3}, z_{11}, z_4, v_{j+1})$, $(v_{j-3}, z_{15}, z_5, v_{j+1})$, $(v_{j-3}, z_8, z_5, v_{j+1})$, (v_{j-3}, z_7, v_{j+1}) , as shown in Figure 8(b). For each of these chains, we create a unique endpoint on the edge e' of Q , where e' corresponds to the edge $v_j v_{j+1}$ in P , as shown in Figure 8(c). We then attach these chains to Q by adding a segment from v_j to its unique endpoint on Q .

We attach the chains for all the convex vertices of P to Q . Later we will use these chains as the strings of an outerstring graph. We then assign each chain a weight, which is the number of convex vertices of P it can remove. For example in Figure 8(b), the weight of the chain $(v_{j-3}, z_8, z_5, v_{j+1})$ is one.

Although the strings are outside of the simple cycle, it is straightforward to construct a representation with all the strings inside a simple cycle Q : Consider placing a dummy vertex at the intersection points of the arrangement, and then find a straight-line embedding of the resulting planar graph such that the boundary of Q corresponds to the outer face of the embedding. Consequently, Q and its associated strings correspond to an outerstring graph representation R . Let G be the underlying outerstring graph. We now claim that any $\text{MWIS}(G)$ corresponds to a min-convex subsuming polygon of P .

Lemma 6. *Let P be a simple polygon, where there exists a min-convex subsuming polygon that lies on $\mathcal{A}_e(P)$, and let G be the corresponding outerstring graph. Any maximum weight independent set of G yields a min-convex subsuming polygon of P .*

PROOF. Let T be a set of strings that correspond to a maximum weight independent set of G . Since T is an independent set, the corresponding subsuming chains do not create edge crossings. Moreover, since each subsuming chain is weighted by the number of convex corners it can remove, the subsuming chains corresponding to T can remove $|\text{MWIS}(G)|$ convex corners in total.

Assume now that there exists a min-convex subsuming polygon that can remove at least k convex corners. The corresponding subsuming chains would correspond to an independent set T' of strings in G . Since each string is weighted by the number of convex corners the corresponding subsuming chain can remove, the weight of T' would be at least k . \square

4.3. Time Complexity

To construct G , we first placed a dummy vertex at the intersection points of the chains, and then computed a straight-line embedding of the resulting planar graph such that all the vertices of Q are on the outerface. Therefore, the geometric representation used at most nt edges to represent each string. Since each convex vertex of P is associated with at most $O(n^t)$ strings, there are at most $n \times O(n^t)$ strings in G . Consequently, the total number of segments used

in the geometric representation is $O(tn^{2+t})$. A subtle point here is that the strings in our representation may partially overlap, and more than three strings may intersect at one point. Removing such degeneracy does not increase the asymptotic size of the representation. Finally, by Theorem 4, one can compute the optimal subsuming polygon in $O(t^3n^{6+3t})$ time.

The complexity can be improved further as follows. Let $abcd$ be a rectangle that contains all the intersection points of $\mathcal{A}_e(P)$. Then, every optimal solution can be extended to a triangulation of the closed region between $abcd$ and Q . Figure 8(d) illustrates this region in gray. We can now apply dynamic programming similar to Section 4.1 to compute the maximum weight independent string set, where each subproblem finds a maximum weight set inside some subpolygon. Each such subpolygon can be described using two points v_1, v_2 , each lying either on Q or on some string, and a subset of $\{a, b, c, d\}$ that helps enclosing the subpolygon.

Since there are $n \times O(n^t)$ strings, each containing at most t points, the number of vertices that correspond to the strings is $O(tn^{1+t})$. We will refer these as the *string vertices*. Note that the number of total vertices in the geometric representation is also $O(tn^{1+t})$. If the subproblem is bounded by two string vertices, or one string vertex and one polygon vertex, then similar to Keil et al. [10], we can use a pair of vertices to describe a subproblem. However, sometimes we need more information to describe a subproblem, e.g., assume that the subproblem is bounded from one side by the point a and some vertex v (corresponding to a string), and from the other side by the point d and some vertex v' (corresponding to a string). For these problems, we need a subset of $\{a, b, c, d\}$ to describe the problem boundary. Therefore, we define our dynamic programming table to be $\mathcal{D}[x, y, z]$, where x and y corresponds to the string or polygon vertices, and z corresponds to the constant size additional description of the boundary (whenever needed). Thus the size of the dynamic programming table is $O(tn^{1+t}) \times O(tn^{1+t}) \times O(1)$. Since there are at most $O(tn^{1+t})$ string vertices, there can be at most $O(tn^{1+t})$ candidate triangles v_1v_2w (e.g., Figure 7 (e)). Consequently, we can fill an entry of the table in $O(tn^{1+t})$ time. Hence the dynamic program takes at most $O(t^3n^{3+3t})$ time in total.

Theorem 7. *Given a simple polygon P with n vertices, one can compute in polynomial time, a min-convex subsuming polygon under the restriction that the subsuming chains must be of constant length and lie on $\mathcal{A}_e(P)$.*

4.4. Generalizations

We can further generalize the results for any given line arrangements. However, such a generalization may increase the time complexity. For example, consider the case when the given line arrangement is $\mathcal{A}_v(P)$, which is determined by the pairs of vertices of P . Since we now have $O(n^2)$ lines in the arrangement $\mathcal{A}_v(P)$, the time complexity increases to $O(t^3(n^2)^{3+3t})$, i.e., $O(t^3n^{6+6t})$.




Figure 9: Illustration for the case when the optimal subsuming polygon contains a subsuming chain of length $\Omega(n)$. The subsuming chain is shown in bold.

5. Conclusion

In this paper, we developed a polynomial-time algorithm that can compute optimal subsuming polygons for a given simple polygon in restricted settings. On the other hand, if the polygon contains holes, then we showed that the problem of computing an optimal subsuming polygon is NP-hard. Therefore, the question of whether the problem is polynomial-time solvable for simple polygons [7], remains open. Note that islands are crucial in our hardness proof. The complexity of the problem for polygons with holes (but without any island) is also open. Since the optimization in one hole is independent of the optimization in the other holes of the polygon, resolving the complexity for polygon with holes would readily give important insight about the complexity for simple polygons.

Our algorithm can find an optimal solution if the optimal subsuming polygon lies on some prescribed arrangement of lines, e.g., $\mathcal{A}_e(P)$ or $\mathcal{A}_v(P)$. The running time of our algorithm depends on the length of the subsuming chains, i.e., the running time is polynomial if the subsuming chains are of constant length. However, there exist polygons whose optimal subsuming polygons contain subsuming chains of length $\Omega(n)$. Figure 9 illustrates such an example optimal solution that is lying on $\mathcal{A}_e(P)$. An interesting research direction would be to examine whether there exists a good approximation algorithm for the general problem.

Recently, Lubiw et al. [11] showed that the problem of drawing a graph inside a polygonal region is hard for the existential theory of the reals. The subsuming polygon problem can also be viewed as a constrained graph drawing problem whereas the subsuming chains are modeled by edges that need to be drawn outside the polygon, possibly with bends. The goal is to find a crossing-free drawing of these edges that minimizes the total number of bends. It would be interesting to examine whether the problem is $\exists\mathbb{R}$ -hard in such a graph drawing model.

Acknowledgement. We thank anonymous reviewers for their feedback on improving the presentation of the paper. The research of Stephane Durocher and Debajyoti Mondal is supported in part by Natural Sciences and Engineering Research Council of Canada (NSERC). Saeed Mehrabi is supported by a

Carleton-Fields postdoctoral fellowship.

References

- [1] Y. Bahoo, S. Durocher, J. M. Keil, S. Mehrabi, S. Mehrpour, D. Mondal, Polygon simplification by minimizing convex corners, in: T. N. Dinh, M. T. Thai (Eds.), Proceedings of the 22nd International Conference on Computing and Combinatorics (COCOON), Vol. 9797 of LNCS, Springer, 2016, pp. 547–559.
- [2] W. A. Mackaness, A. Ruas, L. T. Sarjakoski, Generalisation of Geographic Information: Cartographic Modelling and Applications, Elsevier, 2011.
- [3] H. Ratschek, J. Rokne, Geometric Computations with Interval and New Robust Methods: Applications in Computer Graphics, GIS and Computational Geometry, Horwood Publishing, 2003.
- [4] L. Arge, L. Deleuran, T. Mølhave, M. Revsbæk, J. Truelssen, Simplifying massive contour maps, in: Proc. ESA, Vol. 7501 of LNCS, Springer, 2012, pp. 96–107.
- [5] D. H. Douglas, T. K. Peucker, Algorithm for the reduction of the number of points required to represent a line or its caricature, The Canadian Cartographer 10 (2) (1973) 112–122.
- [6] L. J. Guibas, J. Hershberger, J. S. B. Mitchell, J. Snoeyink, Approximating polygons and subdivisions with minimum link paths, International Journal of Computational Geometry & Applications 3 (4) (1993) 383–415.
- [7] O. Aichholzer, T. Hackl, M. Korman, A. Pilz, B. Vogtenhuber, Geodesic-preserving polygon simplification, International Journal of Computational Geometry & Applications 24 (4) (2014) 307–324.
- [8] M. de Berg, A. Khosravi, Optimal binary space partitions for segments in the plane, International Journal of Computational Geometry & Applications 22 (3) (2012) 187–206.
- [9] A. Darmann, J. Döcker, B. Dorn, On planar variants of the monotone satisfiability problem with bounded variable appearances, <https://arxiv.org/abs/1604.05588> (2016).
- [10] J. M. Keil, J. S. B. Mitchell, D. Pradhan, M. Vatshelle, An algorithm for the maximum weight independent set problem on outerstring graphs, Computational Geometry 60 (2017) 19–25.
- [11] A. Lubiw, T. Miltzow, D. Mondal, The complexity of drawing a graph in a polygonal region, CoRR abs/1802.06699.