Small Strictly Convex Quadrilateral Meshes of Point Sets

David Bremner* Ferran Hurtado† Suneeta Ramaswami‡ Vera Sacristán†

October 29, 2018

Abstract

In this paper, we give upper and lower bounds on the number of Steiner points required to construct a strictly convex quadrilateral mesh for a planar point set. In particular, we show that $3\lfloor \frac{n}{2} \rfloor$ internal Steiner points are always sufficient for a convex quadrilateral mesh of n points in the plane. Furthermore, for any given $n \geq 4$, there are point sets for which $\lceil \frac{n-3}{2} \rceil - 1$ Steiner points are necessary for a convex quadrilateral mesh.

1 Introduction

Discrete approximations of a surface or volume are necessary in numerous applications. Some examples are models of human organs in medical imaging, terrain models in GIS, or models of parts in a CAD/CAM system. These applications typically assume that the geometric domain under consideration is divided into small, simple pieces called *finite elements*. The collection of finite elements is referred to as a *mesh*. For several applications, *quadrilateral/hexahedral* mesh elements are preferred over triangles/tetrahedra owing to their numerous benefits, both geometric and numerical; for example, quadrilateral meshes give lower approximation errors in finite element methods for elasticity analysis [1, 3] or metal forming processes [13]. However, much less is known about quadrilateralizations and hexahedralizations and in general, high-quality quadrilateral/hexahedral (quad/hex) meshes are harder to generate than good triangular/tetrahedral (tri/tet) ones. Indeed, there are several important open questions, both combinatorial as well as algorithmic, about quad/hex meshes for sets of objects such as polygons, points, etc., even in two dimensions. Whereas triangulations of polygons and two-dimensional (2D) point sets and tetrahedralizations of three-dimensional (3D) point sets and convex polyhedra always exist (not so for non-convex polyhedra [24]), quadrilateralizations of 2D point sets do not. Hence it becomes necessary to add extra points, called *Steiner points*, to the geometric domain. This raises the issue of bounding the number of Steiner points, and hence the mesh complexity, while also providing guarantees on the quality of element shape. Such problems are especially relevant for applications in scattered data interpolation [8, 15, 16], which require quadrilateral meshes that modify the original data as little as possible, i.e., add few Steiner points.

* Faculty of Computer Science, University of New Brunswick, P.O. Box 4400, Fredericton, NB, E3B 5A3 Canada. bremner@unb.ca Partially supported by an NSERC Individual Research Grant

† Dep. Matemàtica Aplicada II, Universitat Politècnica de Catalunya, Pau Gargallo 5, 08028 Barcelona, Spain. { hurtado, vera }@ma2.upc.es Partially supported by CUR Gen. Cat. 1999SGR00356 and Proyecto DGES-MEC PB98-0933.

‡ Dep. Computer Science, 322 Business and Science Building, Rutgers University, Camden, NJ 08102, USA. rsuneeta@camden.rutgers.edu

A theoretical treatment of quadrilateral/hexahedral meshes has only recently begun [4, 7, 10, 19, 20, 21, 22]. Some work on quadrangulations[1] of restricted classes of polygons has been done in the computational geometry community [9, 14, 17, 23]. However, there are numerous unresolved questions. For example, even the fundamental question of deciding if a 2D set of points admits a convex quadrangulation without the addition of Steiner points, is unsolved. A survey of results on quadrangulations of planar sets appears in [25].

Any planar point set can be quadrangulated with at most one Steiner point, which is required only if the number of points on the convex hull is odd [7]. For planar simple n-gons, $\lfloor n/4 \rfloor$ internal Steiner points suffice to quadrangulate the polygon [22]. In both cases, the quadrilaterals of the resulting mesh will be, in general, non-convex. However, for many applications, an important requirement is that the quadrangulation be *strictly convex*, i.e., every quadrilateral of the mesh must have interior angles strictly less than 180°. A natural problem then is to construct strictly convex quadrilateral meshes for planar geometric domains, such as polygons or point sets, with a bounded number of Steiner points. Some results on convex quadrangulations of planar simple polygons are known. For example, it was shown in [11] that any simple n-gon can be decomposed into at most $5(n-2)/3$ strictly convex quadrilaterals and that $n-2$ are sometimes necessary. Furthermore, circle-packing techniques [4, 5, 18] have been used to generate, for a simple polygon, quadrilateral meshes in which no quadrilateral has angle greater than 120°. For planar point sets, experimental results on the use of some heuristics to construct quadrangulations with many convex quadrangles appear in [6]. In [12], it is shown that a related optimization problem, namely finding a minimum weight convex quadrangulation (i.e. where the sum of the edge lengths is minimized) can be found in polynomial time for point sets constrained to lie on a fixed number of convex layers.

In this paper, we study the problem of constructing a strictly convex quadrilateral mesh for a planar point set using a bounded number of Steiner points. We use "convex-quadrangulate" to mean "obtain a strictly convex quadrangulation for". If the number of extreme points of the set is even, it is always possible to convex-quadrangulate the set using Steiner points which are all internal to the convex hull. If the number of points on the convex hull is odd, the same is true, assuming that in the quadrangulation we are allowed to have exactly one triangle. We provide upper and lower bounds on the number of Steiner points required for a strictly convex quadrangulation of a planar point set. In particular, in Section 2, we prove that for any $n \geq 4$, $\lceil \frac{n-3}{2} \rceil - 1$ Steiner points may sometimes be necessary to convex-quadrangulate a set of n points. In Section 3, we prove that $3\lfloor \frac{n}{2} \rfloor$ internal Steiner points are always sufficient to convex-quadrangulate any set of n points.

2 Lower bound

In this section we describe a particular configuration of $m+3 \geq 4$ points which requires at least $\lceil \frac{m}{2} \rceil - 1$ Steiner points to be convex-quadrangulated. We also show a convex-quadrangulation of the set that uses close to that few Steiner points.

Description of the configuration of points: The configuration of $m+3$ points consists of $m+1$ points placed along a line ℓ, with one point above the line and another point below the line, such that

[1]In this paper, we use the term *quadrangulation* interchangeably with quadrilateralization. Both terms are common in the meshing literature.

the convex hull of the set has 4 vertices, namely the extreme points on the line and the top and bottom points (see Figure 1). We refer to the vertices on ℓ as *line vertices*. We will refer to the entire configuration as S.

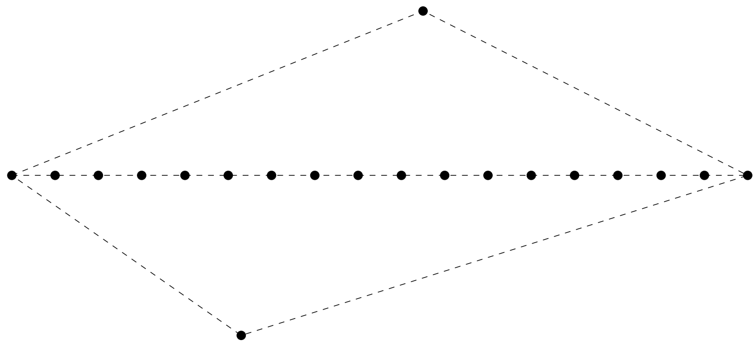

Figure 1: The point set S has $m+1$ points along the line, plus the top and the bottom points. Its convex hull is a quadrangle.

Consider any strictly convex quadrangulation $\mathcal{C}$ of the set. Since all the quadrangles in $\mathcal{C}$ are strictly convex, each point on ℓ must belong to at least one edge of the quadrangulation lying strictly above the line, and at least one edge lying strictly below the line. Quadrangulation edges incident on an input point and lying above (below) ℓ will be called *upward* (*downward*) edges.

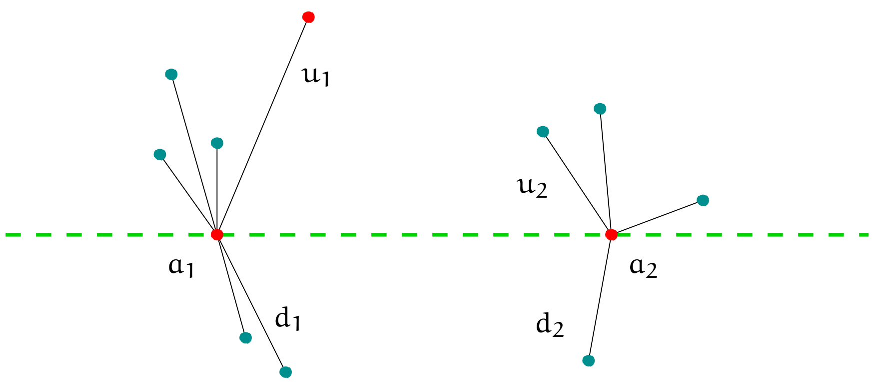

Figure 2: Relevant upward and downward edges.

Consider two consecutive points a_1 and a_2 on ℓ with a_1 to the left of a_2. Let u_1 be the clockwise last upward edge incident on a_1, and let u_2 be the counterclockwise last upward edge incident on a_2. Symmetrically, let d_1 be the counterclockwise last downward edge incident on a_1 and let d_2 be the clockwise last downward edge incident on a_2 (see Figure 2). If (a_1, a_2) is an edge of $\mathcal{C}$, then it must form one quadrangle of $\mathcal{C}$ together with u_1 and u_2, and another one with d_1 and d_2. We call these two faces *squares*. If (a_1, a_2) is not an edge of $\mathcal{C}$, u_1 and d_1 must belong to the same quadrangle, and so must also u_2 and d_2. If these two quadrangles are the same, we call it a *diamond*. If they are different, we call them a *pair of half-diamonds*. These three cases are illustrated in Figure 3.

2.1. Theorem. *The point set S requires at least $\lceil \frac{m}{2} \rceil - 1$ Steiner points to be convex-quadrangulated.*

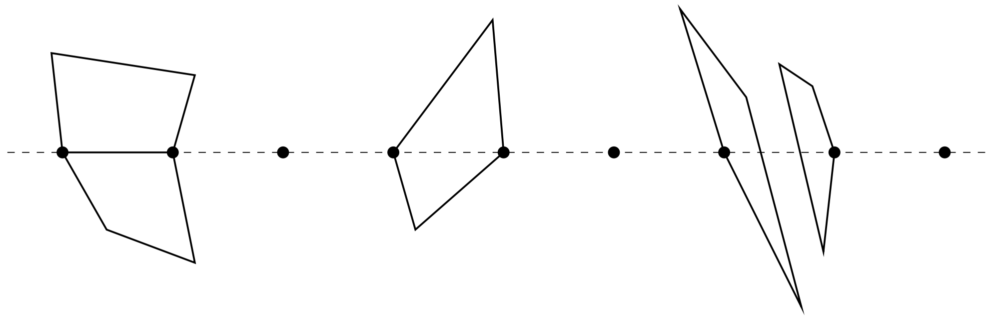

Figure 3: Squares, diamonds and half-diamonds.

Proof. Consider the graph $G = (V, E)$ formed by taking the union of all the squares, diamonds and half-diamonds, together with the convex hull edges. This graph, which is a subgraph of $\mathcal{C}$, is planar and its faces consist of the squares, the diamonds, the half-diamonds, and possibly some other faces that we will call "extra faces". Its edges are all square, diamond, half-diamond, or convex hull edges. Let q be the number of squares, d the number of diamonds and h the number of half-diamonds. We have

$$m = \frac{q}{2} + d + \frac{h}{2}. \tag{1}$$

Let v, e, f denote the number of vertices, edges and faces of G. Let s be the number of vertices that did not belong to the original set, i.e., the number of Steiner points in $\mathcal{C}$. Let x be the number of extra faces. We have $v \leq m+3+s$ (because not every Steiner point need be a vertex of G), and $f = q+d+h+x$. Since G is planar, we can apply Euler's formula and (1) as follows:

$$\begin{aligned} v + f &= e + 2 \\ (m+3+s) + (q+d+h+x) &\geq e + 2 \\ s &\geq e - \frac{3}{2}q - 2d - \frac{3}{2}h - x - 1 \end{aligned}$$

Now, if we can prove that

$$e \geq \frac{7}{4}q + \frac{5}{2}d + \frac{7}{4}h + x, \tag{2}$$

we will obtain that

$$\begin{aligned} s &\geq e - \frac{3}{2}q - 2d - \frac{3}{2}h - x - 1 \\ &\geq (\frac{7}{4} - \frac{3}{2})q + (\frac{5}{2} - 2)d + (\frac{7}{4} - \frac{3}{2})h - 1 \\ &= \frac{q}{4} + \frac{d}{2} + \frac{h}{4} - 1 = \frac{m}{2} - 1 \\ &\geq \left\lceil \frac{m}{2} \right\rceil - 1 \qquad \text{(because } s \text{ must be an integer).} \end{aligned}$$

The general scheme to establish (2) will be to partition the edges of (quadrangles in) G into three sets, and then charge each edge to the faces bounded by the edge. The classification of edges and the charging scheme are as follows:

- Line edges: edges with both endpoints on the line ℓ. Each such edge is shared by a pair of squares. Each square gets charged 1/2.
- Steiner edges: edges with neither endpoint on the line ℓ. Each such edge charges 1/2 to each of the faces that it bounds.
- Vertical edges: edges with exactly one endpoint on the line ℓ.
  - If a vertical edge is shared by two diamonds, each diamond gets charged 1/2.
  - If it is shared by a diamond and an extra face, the diamond gets charged 3/4 and the extra face gets charged 1/4.
  - If it belongs to a square or a half-diamond, the square or half-diamond gets charged 3/8, and the other face gets charged 5/8. Notice that in this last case the total charge is less than 1 when the edge is shared by squares and/or half-diamonds.

As a result, each face of G gets charged in the following way:

- Each extra face is charged at least 1, since it has at least four edges (recall that G is a subgraph of the quadrangulation $\mathcal{C}$) and is charged at least 1/4 from each edge.
- Each square is charged 7/4: 1/2 from its line edge, 1/2 from its Steiner edge, and 3/8 from each of its two vertical edges.
- Each half-diamond is charged 7/4: 1/2 from each of its two Steiner edges and 3/8 from each of its two vertical edges.
- Each diamond is charged at least 5/2. Notice that if a diamond α shares one upward vertical edge (a_1, u_1) with another diamond or half-diamond β, then it must share the coincident downward edge (a_1, d_1) with an extra face, since no square, diamond or half-diamond could share it. This is because *(i)* α obviously cannot share (a_1, d_1) with β because of strict convexity, and *(ii)* α cannot share (a_1, d_1) with any other diamond, square, or half-diamond face because such a face would intersect β. For the same reasons, if a diamond shares a downward edge (a_1, d_1) with another diamond or half-diamond, the coincident upward edge (a_1, u_1) must be shared with an extra face. So, if the diamond is adjacent to another diamond, it is charged at least $\frac{1}{2}+\frac{3}{4}=\frac{5}{4}$ from the two edges incident on that line vertex. If it is adjacent to a pair of squares, it is charged $\frac{5}{8}+\frac{5}{8}=\frac{5}{4}$. Any other combination would charge more. In total, the diamond gets charged at least 5/2.

This proves that

$$e \geq \sum \text{charges} = x + \frac{7}{4}q + \frac{7}{4}h + \frac{5}{2}d.$$

■

2.2. Theorem. *The point set S can be convex-quadrangulated with* $s \leq \left\lceil \frac{m+3}{2} \right\rceil$ *Steiner points.*

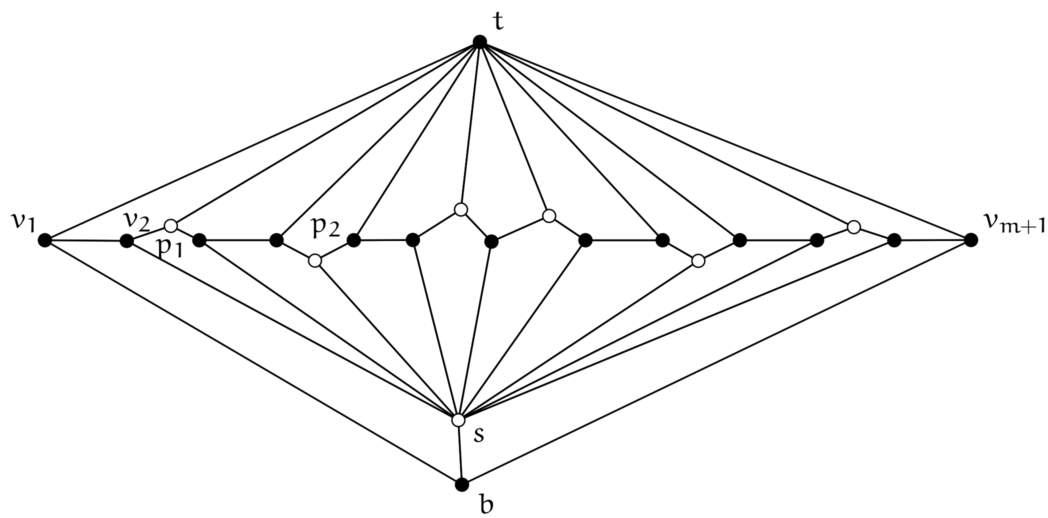

Figure 4: A convex quadrangulation of S using $\lceil \frac{m}{2} \rceil$ Steiner points.

Proof. It is possible to convex-quadrangulate the given point set configuration with s Steiner points, where

$$s = \left\{ \begin{array}{ll} \frac{m}{2} + 1, & \text{if } m \equiv 0 \pmod 4 \\ \frac{m+1}{2} + 1, & \text{if } m \equiv 1 \pmod 4 \\ \frac{m}{2} + 2, & \text{if } m \equiv 2 \pmod 4 \\ \frac{m+1}{2}, & \text{if } m \equiv 3 \pmod 4 \end{array} \right\} \leq \left\lceil \frac{m+3}{2} \right\rceil$$

A solution is presented in Figure 4 , where the original points are shown in black and the Steiner points in white. This solution can be described as follows. Let v_i, $i \in \{1, \ldots, m+1\}$ be the points on the line ℓ, and t and b the top and bottom points. Place one Steiner point s below ℓ, inside the convex hull and in $L(b, v_2) \cap R(b, v_m)$. Quadrangles bsv_2v_1 and $bv_{m+1}v_ms$, both of which are strictly convex, are part of the quadrangulation. We call the line segment v_iv_{i+1} (not necessarily part of the quadrangulation) the ith *virtual edge* e_i. Suppose $m = 4k + r$, $0 \leq r \leq 3$. Starting from both ends of l, 2k Steiner points p_i are placed alternately above and below every other virtual edge on ℓ. More precisely, for $1 \leq i \leq k$ place a Steiner point p_i above (resp. below) e_{2i} if i is odd (resp. even). In both cases ensure p_i is in the intersection of the wedges $v_{2i}tv_{2i+1}$ and $v_{2i}sv_{2i+1}$. Connect p_i to t (i odd) or s (i even), and to v_{2i} and v_{2i+1}. Connect v_{2i-1} to v_{2i}. Carry out the analogous procedure starting with the rightmost virtual edge. After placing 2k Steiner points, we are left with r "untreated" virtual edges $e'_1, e'_2, \ldots e'_r$ in the center. If $r \leq 2$, we place Steiner points as follows: one point above (resp. below) each e'_i if k is odd (resp. even). If $r = 3$ then we place point below (resp. above) e'_2 if k is odd (resp. even). In all cases we insure the the Steiner point is within the two wedges defined by the virtual edge, s and t. The strict convexity of the quadrangles created by this procedure is ensured by placing each Steiner point in the intersection of these two wedges. The fact that the number of Steiner points used in these quadrangulations is off by a small constant from the bound given by our charging scheme is explained by the charges on the extra faces (drawn shaded in Figure 5). In these cases, the extra faces actually get charged more than 1, whereas we count a charge of only 1 for any extra face of the quadrangulation when proving the lower bound. ∎

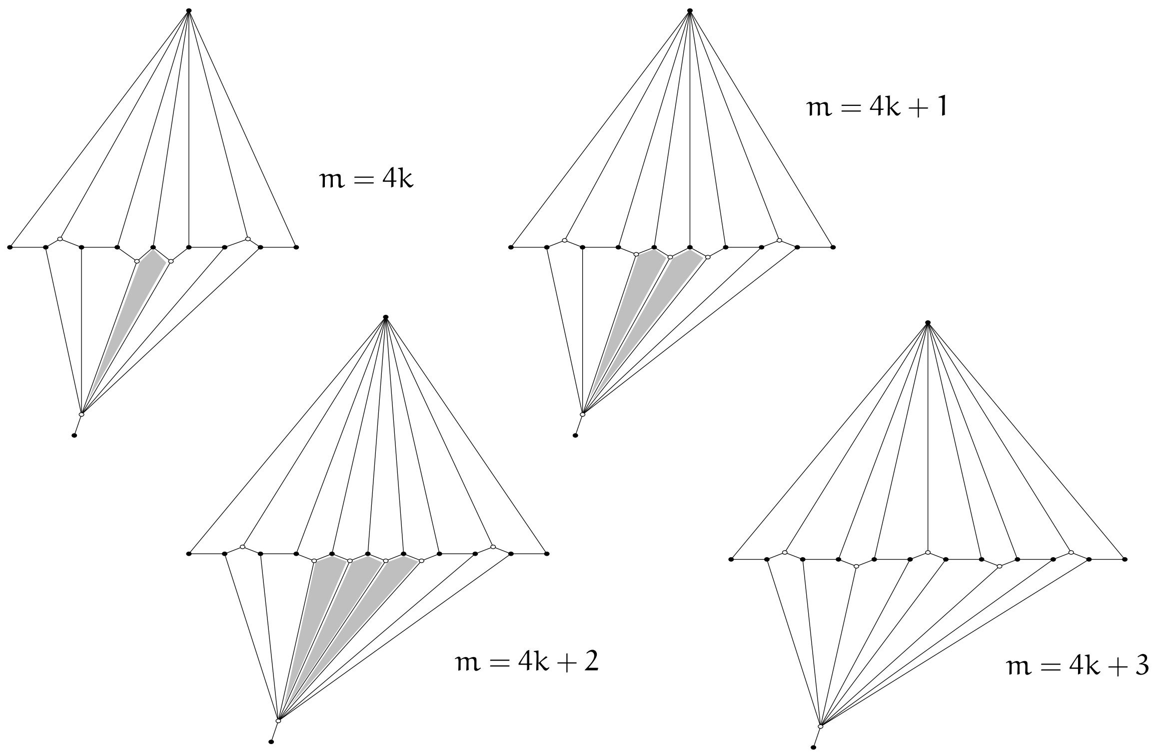

Figure 5: Convex quadrangulations for any parity of m (shaded faces are extra faces).

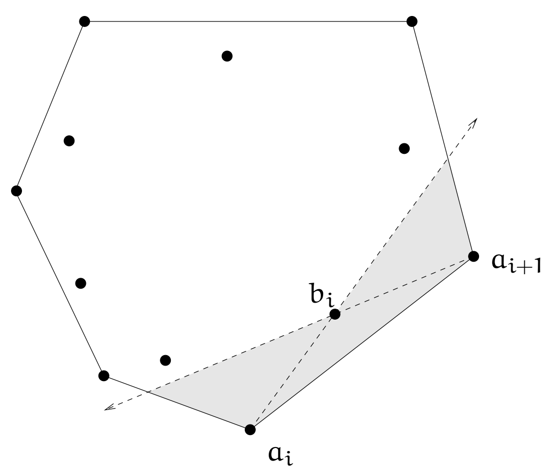

Figure 6: The point set P has $n/2$ points on the convex hull, and $n/2$ interior points lying very close to edges of the convex hull.

Theorem 2.1 uses a highly degenerate configuration, where most of the points lie on a straight line. It turns out that the same lower bound result cannot be obtained from this point configuration if it is perturbed: if the points do not lie on a straight line, then squares can be formed using only input points (i.e., without using Steiner points), and so can diamonds. It turns out that there exist analogous configurations with $m+1$ points on an arbitrary upward convex curve (instead of a straight line), where the point set can be convex-quadrangulated with a constant number of Steiner points. We now describe a perturbable (i.e. non-degenerate) point set configuration that requires at least $\frac{n}{4}$ Steiner points for a strictly convex quadrangulation.

Description of the perturbable configuration of points: Let $n = 2k$. Place k points in convex position. Place the remaining k points such that, for each edge e of the convex hull, there is one point lying in the interior of the convex hull, very close to the midpoint of e. To be more precise, if (a_i, a_{i+1}) is an edge of the convex hull, the new point b_i must be located so that $a_{i+2} \in L(a_i, b_i)$ and $a_{i-1} \in R(a_{i+1}, b_i)$, as illustrated in Figure 6. Call this point set P.

2.3. Theorem. *The point set P requires at least $\frac{n}{4}$ Steiner points to be convex-quadrangulated.*

Proof. By definition, each convex hull edge (a_i, a_{i+1}) must belong to one quadrangle Q_i. For Q_i to be convex and not contain any interior point, its remaining two vertices must belong to the region $G(i) = R(a_i, b_i) \cup L(a_{i+1}, b_i)$; one of these vertices may be b_i (see Figure 6). Hence, for every convex hull edge there is at least one Steiner point in region $G(i)$. Since only consecutive regions intersect, at least one Steiner point is needed for every pair of convex hull edges, i.e. at least $\frac{n}{4}$ Steiner points are needed. ■

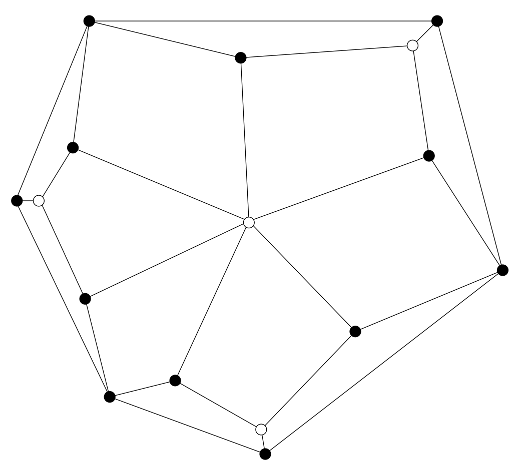

Figure 7: A convex-quadrangulation of P that uses $\frac{n}{4} + 1$ Steiner points.

2.4. Theorem. *P can be convex-quadrangulated with at most $\frac{n}{4} + 1$ Steiner points.*

Proof. Figure 7 shows a convex-quadrangulation of the set that uses $\frac{n}{4} + 1$ Steiner points. There is one quadrangle for every convex hull edge $a_i a_{i+1}$. It has b_i as a vertex and uses one Steiner point, which is shared by the adjacent convex hull edge. Finally, one central Steiner point is used to convex-quadrangulate the remaining interior face.

In this configuration n is always even. If $n \equiv 2 \pmod 4$, then we have an odd number of points on the convex hull, and $\lceil \frac{n}{4} \rceil$ Steiner points suffice to convex-quadrangulate with an extra triangle, formed by one of the convex hull edge and its corresponding interior point. ∎

3 Upper bound

Given a set S of n points in the plane, $\mathrm{conv}(S)$ is the convex hull of S. For a simple polygon P, $\mathrm{int}(P)$ denotes the interior of P and $\mathrm{kernel}(P)$ is the locus of points in P that can see all of P. A *convex quadrangulation* of S is a decomposition of $\mathrm{conv}(S)$ into strictly convex quadrangles and at most one triangle, such that no cell contains a point of S in its interior. The vertices of the quadrangulation that do not belong to S are called *Steiner points*. In what follows, angles greater than or equal to 180° are *reflex angles*.

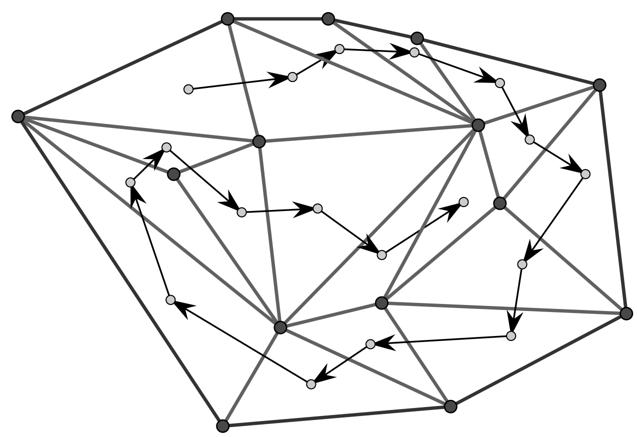

Figure 8: A path triangulation of a point set

3.1. Theorem. *Any set of n points can be convex-quadrangulated using at most $3\lfloor \frac{n}{2} \rfloor$ Steiner points.*

Proof. Any set S of n points has a path triangulation (a triangulation whose dual graph has a Hamiltonian path), which can be constructed in $O(n \log n)$ time [2, 7] (Figure 8 illustrates such a triangulation of a point set). Denote by t the number of triangles in any triangulation of n points with h extreme points ($t = 2n - 2 - h$). By pairing up the triangles along the path, we obtain a path quadrangulation of S with possibly one unpaired triangle (see Figure 9). We will prove in Section 3.1 that it is always possible to convex-quadrangulate a pair of consecutive quadrangles by using at most 3 internal Steiner points. At the end of the process we may have any of the following situations:

- There is no unpaired triangle (*i.e.*, h is even) and all the quadrangles have been paired up. In this case, a convex-quadrangulation has been obtained with no leftover triangle. Therefore,

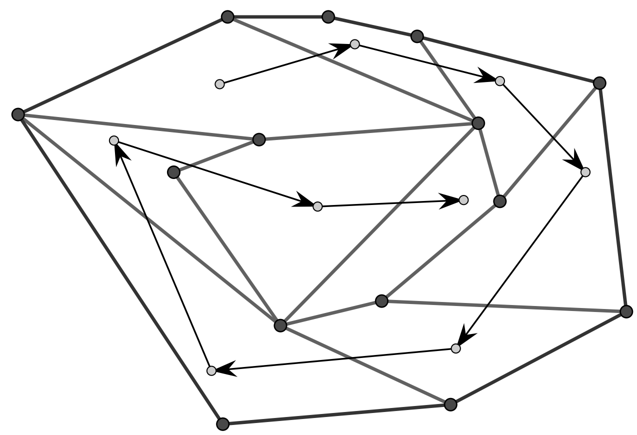

Figure 9: The triangles of Figure 8 paired to form a (not necessarily convex) quadrangulation

the number of quadrangles is $q = \frac{t}{2} = n - 1 - \frac{h}{2}$, and the total number of Steiner points used is $s = \frac{3}{2}q = \frac{3}{2}(n - 1 - \frac{h}{2}) \leq 3\lfloor\frac{n}{2}\rfloor$.

- There is one unpaired triangle (*i.e.*, h is odd), all the quadrangles have been paired up. In this case, the number of quadrangles is $q = \frac{t-1}{2}$, which is less than in the previous case. Once again, the number of Steiner points $s = \frac{3}{2}q < 3\lfloor\frac{n}{2}\rfloor$.

- There is no unpaired triangle and all the quadrangles except one have been paired up. In this case, the last quadrangle can be convex-quadrangulated, if it is not convex, by adding 4 internal Steiner points (see page 22 for details). Since all quadrangles except one have been paired up, the number of Steiner points used is $s = \frac{3}{2}(q-1) + 4 = \frac{3}{2}(\frac{t}{2} - 1) + 4 = \frac{3}{2}(n - 2 - \frac{h}{2}) + 4 \leq \frac{3}{2}n - \frac{3}{4}h + 1 < 3\lfloor\frac{n}{2}\rfloor$ since $h \geq 4$.

- There is one unpaired triangle, and all the quadrangles except one have been paired up. We can convex-quadrangulate the remaining quadrangle with 4 Steiner points as before, and leave the triangle as it is. In this case, we have $q = \frac{t-1}{2}$ and $s = \frac{3}{2}(q-1) + 4 < 3\lfloor\frac{n}{2}\rfloor$ (as argued in the previous case).

Note that the number of quadrilaterals in the quadrangulation is at most $5\lfloor\frac{n}{2}\rfloor - \frac{h}{2}$. Note also that the quadrangulation produced by our algorithm is strictly convex, even if the path quadrangulation contains degenerate quadrilaterals. ∎

3.1 Pairing up quadrangles

Before discussing the details of how to convex-quadrangulate a pair of adjacent quadrilaterals, we introduce some notation and mention a few useful facts about polygons. Given two points p and q, we will denote by $L(p, q)$ (resp. $R(p, q)$) the left (resp. right) open half-plane defined by the oriented line from p to q. Throughout this section, vertices of polygons will be enumerated counterclockwise. Given a vertex v of a polygon P, we denote its successor (resp. predecessor) by v^+ (resp. v^-), and we write $\text{wedge}(v)$ to mean $L(v^-, v) \cap R(v^+, v) \cap \text{int}(P)$. If v is reflex, $\text{wedge}(v)$ will denote the locus of points (inside P) that can be connected to v forming strictly convex angles at v. If v is convex, wedge(v) is the interior of the visibility region of v in P. Given three points p, q, and r, $\triangle(pqr)$ is the open triangle defined by the three points, i.e. $\triangle(pqr) = \operatorname{int}\operatorname{conv}(p, q, r)$. Note that

$$\operatorname{int}\operatorname{kernel}(P) = \cap_{v \in P} L(v, v^{+}) . \tag{3}$$

We can observe the following

$$\begin{aligned} \operatorname{wedge}(v_i) &= L(v_i^-, v_i) \cap R(v_i^+, v_i) \\ &= L(v_i^-, v_i) \cap L(v_i, v_i^+) \end{aligned}$$

It follows that

$$\operatorname{int}\operatorname{kernel}(P) = \bigcap_i \operatorname{wedge}(v_{2i}) . \tag{4}$$

Similarly, by noting that if $v_0 \ldots v_k$ form a reflex chain,

$$\bigcap_{0 \le i < k} L(v_i, v_{i+1}) = L(v_0, v_1) \cap L(v_{k-1}, v_k) ,$$

it follows that

$$\operatorname{int}\operatorname{kernel}(P) = \bigcap_{v \text{ convex}} \operatorname{wedge}(v) . \tag{5}$$

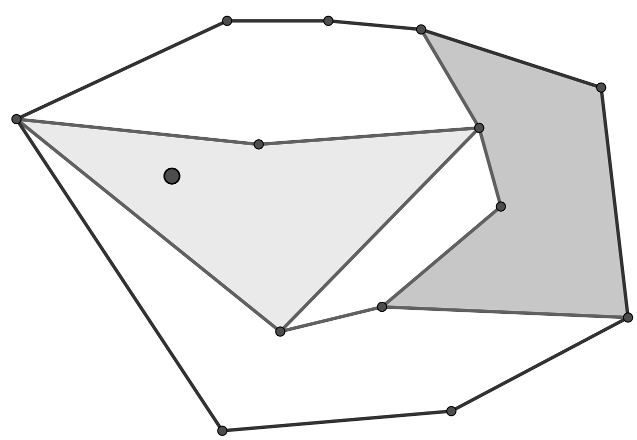

Figure 10: The union of a pair of adjacent quadrangles is either a hexagon or quadrangle with a fifth interior point.

Consider a pair of consecutive quadrangles in the path quadrangulation. They may share one edge or two edges. In the first case, their union is a hexagon, while in the second case it is a quadrangle containing a fifth point in its interior (see Figure 10). In the rest of this section we will examine in detail how to convex-quadrangulate the union of two quadrangles. The general scheme will be inductive, i.e. to reduce each case to one requiring fewer Steiner points by the addition of a single Steiner point. Table 1 provides a summary of all the cases and their interdependencies. Most

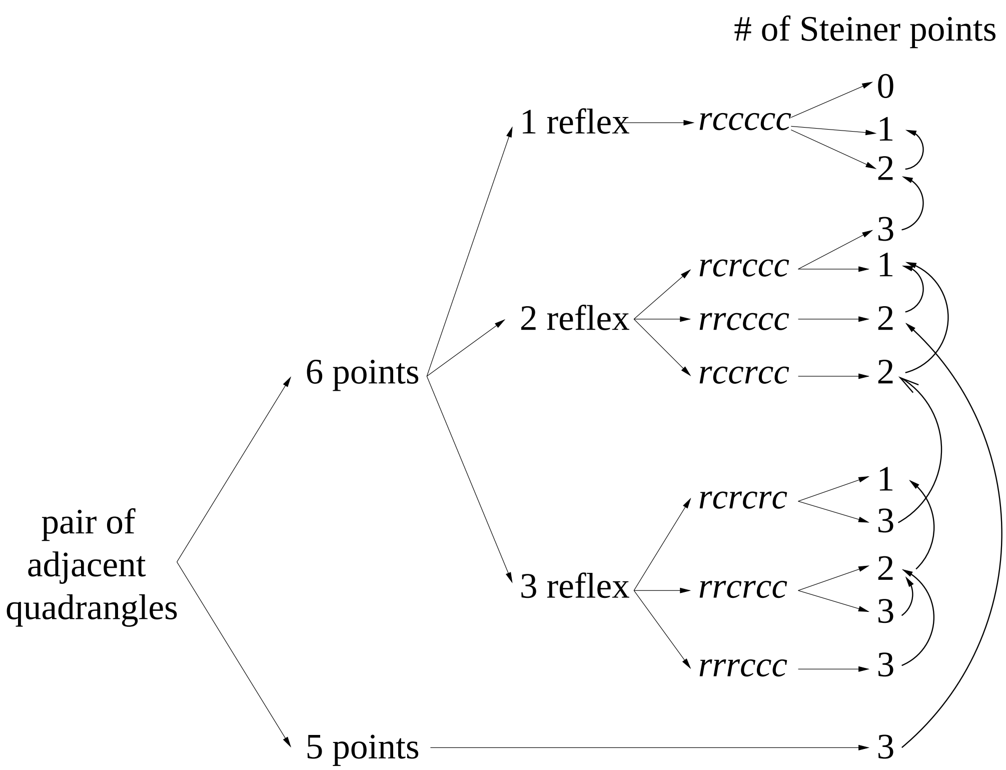

Table 1: Scheme of the proof.

of the cases are given a mnemonic label describing the cyclic order of reflex and convex vertices around the polygon boundary and the total number of Steiner points necessary (e.g. rcrcrc-1 describes the case where reflex and convex vertices alternate and one Steiner point suffices to convex-quadrangulate the hexagon). The last column reports the number of Steiner points used in each case. The arrows on the right indicate the reductions, after adding one Steiner point, from one case to another. As is suggested by Table 1, the majority of our effort in the remainder of this section will be devoted to proving the following theorem.

3.2. Theorem. *Any hexagon can be convex-quadrangulated by placing at most* 3 *Steiner points in its interior.*

The curious reader is referred to Figure 11 for a convex quadrangulation resulting from applying our techniques to the point set of Figures 8, 9, and 10. In the figure the white points are Steiner points.

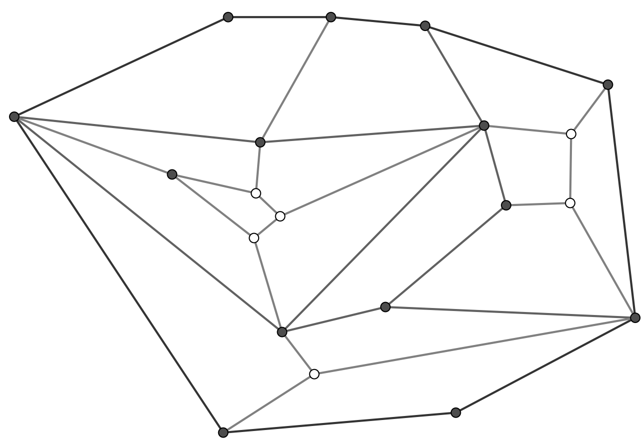

Figure 11: A convex quadrangulation resulting from applying our techniques to the point set of Figures 8, 9, and 10.

3.1.1 Independent Triples

We call a set of vertices of a polygon *independent* if no two of them are endpoints of the same edge. We start by establishing some useful properties of independent triples of vertices of a hexagon. All lemmas in this section hold even when reflex angles are exactly equal to 180°. Let $\{a, c, e\}$ be an independent triple for a hexagon $P = abcdef$.

3.3. Lemma. *If* $\triangle(ace) \subset P$ *then* $\triangle(ace) \cap \mathrm{wedge}(a) = \triangle(ac'e')$*, where* $c'e' \subseteq ce$ *and* $c' \neq e'$*.*

Proof. It suffices to establish that $ce \cap \mathrm{wedge}(a)$ is a non-trivial line segment. The result then follows by convexity. If a is convex, then c and e are in the visibility polygon of a, i.e. in $\mathrm{wedge}(a)$. Suppose then that a is reflex. If one of c or e is contained in $\mathrm{wedge}(a)$ then the lemma holds. If neither c nor e belongs to $\mathrm{wedge}(a)$, then they cannot both belong to $R(a^-, a)$ (resp. $L(a^+, a)$) because then a^- (resp. a^+) cannot see c or e, which is a contradiction because a^{--} (resp. a^{++})

must be c or e, since a, c and e are at distance two. Therefore, c and e must be on opposite sides of $\mathrm{wedge}(a)$ (*i.e.*, one in $R(a^-, a)$ and the other in $L(a^+, a)$), and the segment ce must have non-trivial intersection with $\mathrm{wedge}(a)$. ∎

3.4. Lemma. *If* $\triangle(ace) \subset P$ *then* $\mathrm{wedge}(a) \cap \mathrm{wedge}(c) \cap \triangle(ace) \neq \emptyset$*.*

Proof. This follows by applying Lemma 3.3 twice, and convexity. ∎

3.5. Lemma. *If* $\triangle(ace) \subset P$ *then* $\triangle(ace) \cap \mathrm{wedge}(a^-) \cap \mathrm{wedge}(a^+) \neq \emptyset$*.*

Proof. Note that a^+ and a^- must both be convex (or exactly 180°). It follows that in a neighborhood $N(a)$ of the point a, we have $N(a) \cap \mathrm{wedge}(a) = N(a) \cap \mathrm{wedge}(a^-) \cap \mathrm{wedge}(a^+)$, and Lemma 3.3 applies. ∎

3.6. Lemma. *If* P *is starshaped and* $\triangle(ace) \subset P$*, then one Steiner point suffices to convex-quadrangulate* P*.*

Proof. From (4), $\mathrm{int\,kernel}(P) = \mathrm{wedge}(a) \cap \mathrm{wedge}(c) \cap \mathrm{wedge}(e)$. Each pair of these wedges intersect $\triangle(ace)$, as a consequence of Lemma 3.4. Each pair of wedges intersects as a consequence of Lemma 3.5. In this case we consider the wedges extended to the entire plane, and not restricted to the polygon. Since both the triangle and the (extended) wedges are convex, Helly's theorem [26] applies. It follows they all intersect, i.e. the triangle $\triangle(ace)$ must intersect the interior of the kernel.

One Steiner point s can then be placed in the intersection of the triangle and the kernel, and connected to the three reflex vertices (see Figure 12). Since s belongs to the kernel, it belongs to the wedges of the three reflex vertices, hence a, c, and e are now strictly convex vertices in the quadrangulation. Since s belongs to $\triangle(ace)$, s is convex in all the quadrangles. ∎

3.7. Lemma. *If* c *does not see* e*, and* a *is the only reflex vertex other than possibly* c *or* e*, then*

(a) $\mathrm{wedge}(a) \cap \mathrm{wedge}(c) \neq \emptyset$*, and*

(b) $\mathrm{wedge}(a) \subset L(a, c)$*.*

Proof.

(a) Since b and f are convex, $a \in \triangle(cde)$, since otherwise nothing can block ce (notice that both abc and aef are ears of the polygon). Similarly, the (non-convex) quadrangle $acde$ must be empty (see Figure 13). It follows that a sees $\{c, d, e\}$. We can conclude that $a \in R(d, c)$ (by seeing d) and $a \in L(b, c)$ (by seeing c and d). In other words, $a \in \mathrm{wedge}(c)$. The claim then follows from the fact that for some neighborhood $N(a)$, $N(a) \cap \mathrm{wedge}(a) \subset \mathrm{wedge}(c)$.

(b) Since $a \in \triangle(cde)$, it follows that $e \in R(a, c)$. Again considering the fact that aef forms an ear of the polygon, we have $f \in R(a, c)$. It follows that both of the chords defining $\mathrm{wedge}(a)$ are contained in $L(a, c)$. ∎

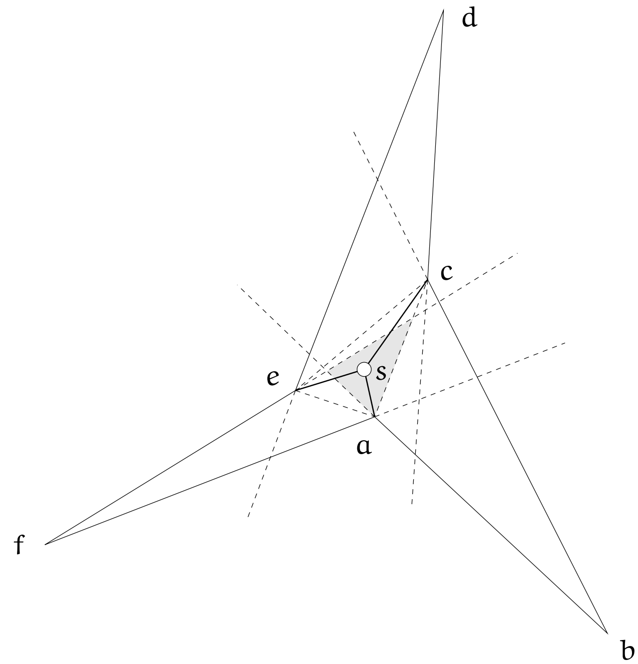

Figure 12: One Steiner point suffices if $\triangle(abc) \subset P$ and P is starshaped.

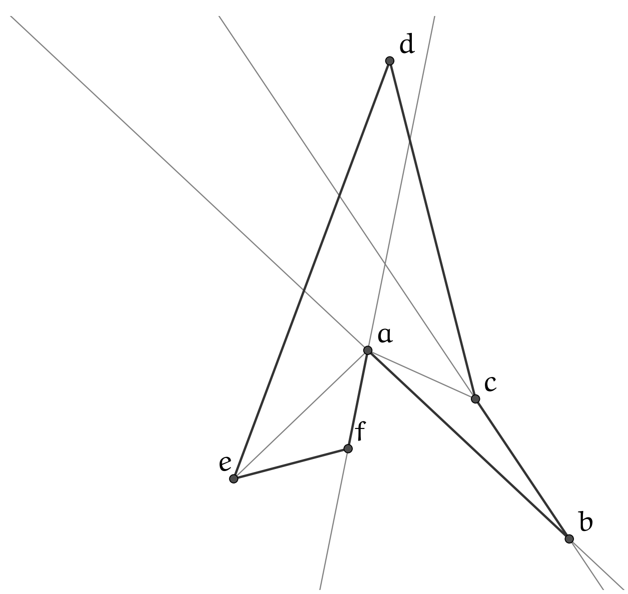

Figure 13: Illustrating the proof of Lemma 3.7

3.1.2 Proof of Theorem 3.2.

We are now ready to carry out the case analysis described in Table 1. A hexagon may have zero, one, two, or three reflex vertices; we consider each of these cases in turn. In the remainder of this section, we will use *convex* to mean *strictly convex*.

Hexagon with no reflex vertices. In this case, the hexagon can be trivially decomposed into two convex quadrangles without using any Steiner points.

Hexagon with one reflex vertex. Suppose w.l.o.g. that vertex a is reflex.

1. (*rccccc–0*) If $d \in \mathrm{wedge}(a)$ then no Steiner points are needed. Connecting d with a will produce a convex quadrangulation of the hexagon, as shown in Figure 14. Note that if vertex a is equal to 180°, this case must be satisfied.

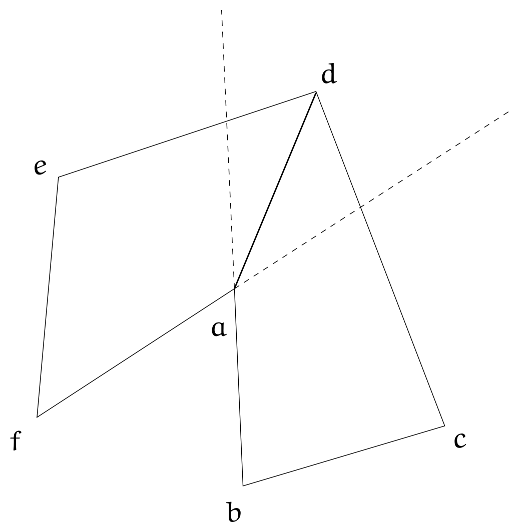

Figure 14: No Steiner points needed.

2. If $d \notin \mathrm{wedge}(a)$, then d must lie on one side of $\mathrm{wedge}(a)$ (i.e., $d \in L(b, a)$ or $d \in R(f, a)$), and at least one of e or c, w.l.o.g. e, must lie on the same side (since both e and c are convex).

2.1. (*rccccc–1*) If $ce \subset P$, by Lemma 3.6 one Steiner point is sufficient.

2.2. (*rccccc–2*) If c and e do not see each other, two Steiner points are enough. Placing a Steiner point s in $\mathrm{wedge}(a)$ and connecting it to a and c decomposes the hexagon into a quadrangle *abcs* and a hexagon *ascdef* (see Figure 15). The quadrangle is convex: a is convex because $s \in \mathrm{wedge}(a)$. The vertex s is convex because $c \in R(f, a)$ ($c \notin L(f, a)$, because then c and e would see each other) and hence $c \in R(a, s)$. The hexagon *ascdef* is as in the previous case rccccc-1 (d and f necessarily see each other because of our assumption that d and e lie on the same side of $\mathrm{wedge}(a)$) and hence can be quadrangulated with one additional Steiner point.

Hexagon with two reflex vertices. There are several different cases, depending on the relative positions of the two reflex vertices in the polygon boundary.

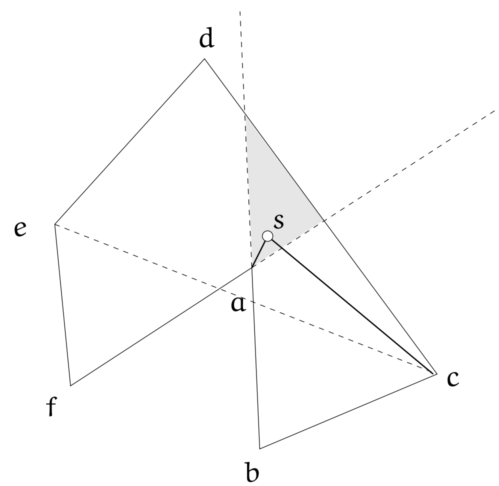

Figure 15: One Steiner point reduces case rcccccc-2 to case rcccccc-1.

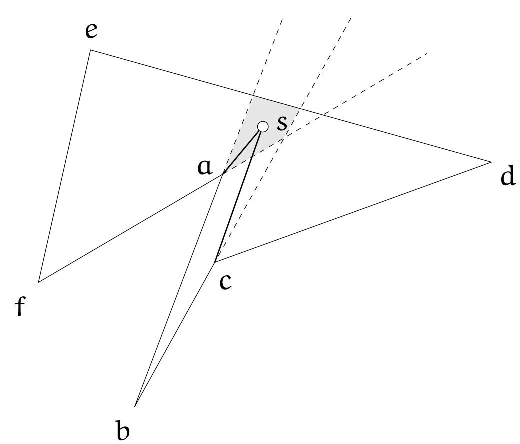

Figure 16: One Steiner point reduces the problem to the one reflex vertex case.

1. (*rcrccc*) Suppose that the two reflex vertices are separated by a convex vertex of the polygon. Let us assume that a and c are the reflex vertices of the hexagon *abcdef*. There are two sub-cases.

1.1. (*rcrccc—1*) If both a and c can see e, then one Steiner point is enough. Note that since e is convex $\triangle(ace) \subset \text{wedge}(e)$. By Lemma 3.4 $\text{wedge}(a) \cap \text{wedge}(c) \cap \triangle(ace) \neq \emptyset$. It follows from (4), that the hexagon is starshaped. We can then apply Lemma 3.6.

1.2. (*rcrccc—3*) Otherwise, one of the reflex vertices, w.l.o.g. a, obstructs the visibility from the other reflex vertex to e. We show that 3 Steiner points suffice. By Lemma 3.7 $\text{wedge}(a) \cap \text{wedge}(c) \cap L(a, c) \neq \emptyset$. Place a Steiner point s in this region and connect it to a and c (see Figure 16). The quadrangle *abcs* must be convex: a and c are convex because s belongs to their wedges. The vertex s is convex because it belongs to $L(a, c)$. The remaining hexagon has only one reflex vertex s, hence can be convex-quadrangulated with at most 2 additional Steiner points.

2. (*rrcccc—2*) If the two reflex vertices are consecutive, then two Steiner points are always sufficient. Let a and b be the two reflex vertices. Notice that since there are only two consecutive reflex vertices, a must necessarily see e, and b must see d. Place one Steiner point s in $\text{wedge}(a) \cap R(a, e) \cap L(b, d)$. This region is not empty because $de \in L(b, d)$, and $\text{wedge}(a) \cap R(a, e)$ contains a subset of the edge de (this can be seen by noting that either d or e belong to $\text{wedge}(a)$, or they lie on opposite sides of $\text{wedge}(a)$). Connect s to a and e (refer to Figure 17). The quadrangle *asef* must be convex: a is convex since $s \in \text{wedge}(a)$, and s is convex because $s \in R(a, e)$. The remaining hexagon has two reflex vertices, namely s and b, separated by a convex vertex a. Both s and b can see d, since $s \in L(b, d)$. This is the *rcrccc*-1 case, which requires one additional Steiner point.

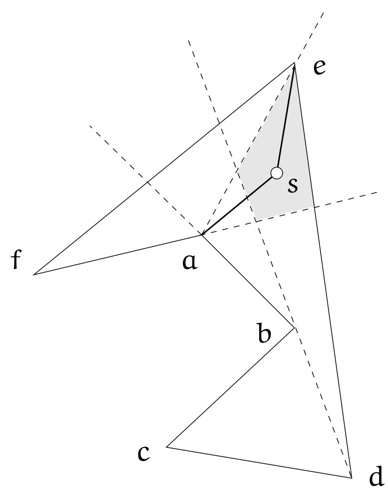

Figure 17: One Steiner point reduces the problem to the *rcrccc*-1 case.

3. (*rccrcc—2*) We are left with the case in which there are two convex vertices between the two reflex vertices, both clockwise and counterclockwise. In this case, two Steiner points suffice. Let a and d be the reflex vertices. We will use the fact that either the two diagonals ae and bd are internal to the polygon or ac and df are. The reason is that if ac is obstructed by d,

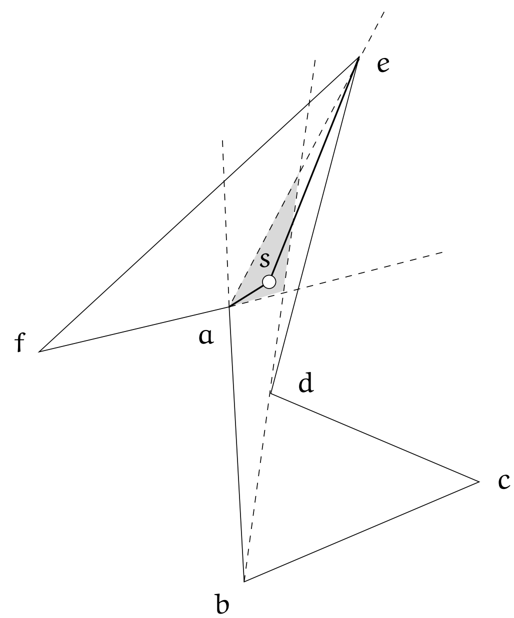

Figure 18: One Steiner point reduces rccrcc-2 to rcrccc-1 case.

then d belongs to $\triangle(abc)$ (recall that a, b, c, d are consecutive) and must see b, which implies that diagonal ae cannot be obstructed. A symmetric argument holds if df is obstructed by a. Let us assume that ae and bd are internal diagonals (see Figure 18). Then one Steiner point s can be placed in $\mathrm{wedge}(a) \cap R(a, e) \cap L(b, d)$. This region can be seen to be nonempty as follows: from the convexity of f, f must belong to $L(a, e)$ and hence $R(a, e) \cap \mathrm{wedge}(a) \neq \emptyset$. In fact $R(a, e) \cap \mathrm{wedge}(a)$ contains a neighborhood of a, which is in turn contained in $L(b, d)$. Connect s to a and e. The quadrangle *asef* is convex. The remaining polygon is the *rcrccc*-1 type: s and d are its reflex vertices, and they both see b, since $s \in L(b, d)$.

Hexagon with three reflex vertices. Again, there are different situations, depending on the relative positions of the reflex vertices along the polygon boundary.

1. (*rcrcrc*) We start with the case in which the reflex and the convex vertices alternate.

1.1. (*rcrcrc—1*) In the special case that $\triangle(ace)$ is inside the polygon *and* the polygon is star shaped, Lemma 3.6 implies that one Steiner point suffices.

1.2. (*rcrcrc—3*) Otherwise, we show that 3 Steiner points suffice. The region $\rho = \mathrm{wedge}(a) \cap \mathrm{wedge}(e) \cap R(a, e)$ must be non-empty for the following reason: If $\triangle(ace)$ is inside the polygon, then ρ is non-empty as a consequence of Lemma 3.4. If on the other hand one of the edges of $\triangle(ace)$, w.l.o.g. ac is obstructed, then ρ is non-empty by Lemma 3.7. Place a Steiner point s inside ρ. Connect s to a and e. The quadrangle $efas$ is convex. Vertices a and e are convex by virtue of s being in the appropriate wedges. The vertex s is convex because $s \in R(a, e)$. The hexagon $sabcde$ is of type rccrcc-2 (since $s \in \mathrm{wedge}(a) \cap \mathrm{wedge}(e)$) hence can be quadrangulated with two additional Steiner points.

2. (*rrcrcc*) We now study the case in which there are exactly two consecutive reflex vertices. These polygons are always star-shaped, for the following reasons: Suppose that a, b and d are the reflex vertices (refer to Figure 19). Consider the wedges of f and c. The point e must lie on the left ray of wedge(f), and to the right of (or on) the right ray of wedge(c) (since d is reflex). As a consequence, these two rays must intersect (inside P) in a point that we will call i. Since d is reflex, it must lie in the segment ci, and e cannot lie in the interior of segment fi. As a consequence, some portion of the edge ef must belong to $\text{wedge}(f) \cap \text{wedge}(c) \cap \text{wedge}(e) = \text{int kernel}(P)$, (see (5)). We have two cases depending on whether e sees at least one of a and b.

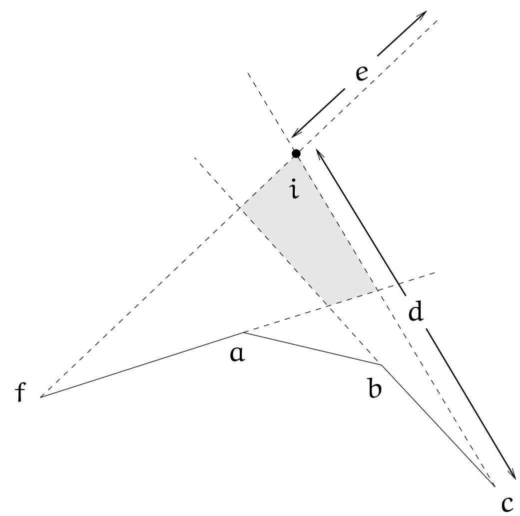

Figure 19: Proving that the *rrcrcc* polygons are starshaped.

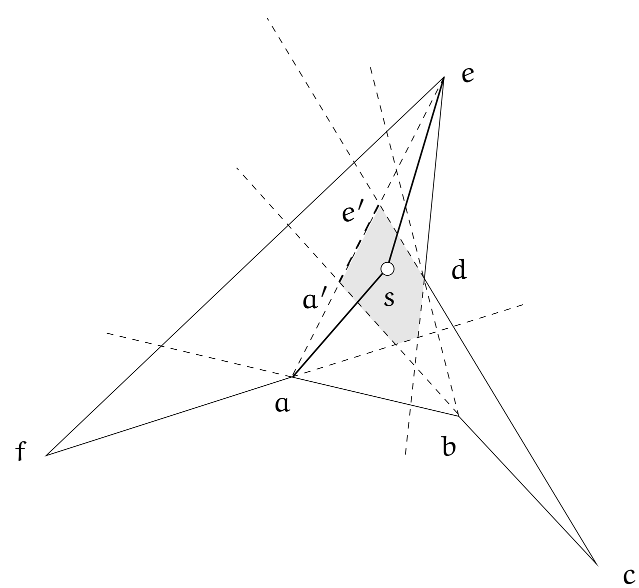

Figure 20: One Steiner point reduces the problem to the *rcrcrc*-1 case.

2.1. (*rrcrcc—2*) If e sees at least a, two Steiner points suffice. In particular the region $\text{kernel}(P) \cap R(a, e) \cap L(b, d)$ (see Figure 20) cannot be empty, for the following reason: The fact that e and

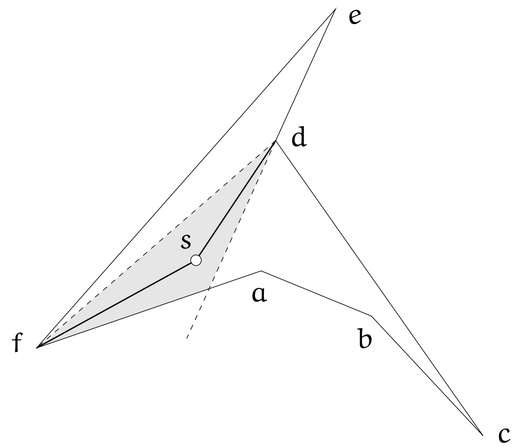

Figure 21: One Steiner point reduces the problem to the *rrcrcc*-2 case.

a see each other implies that $f \in L(a, e)$ and $b, c, d \in R(a, e)$. Hence $ae \subset \text{wedge}(f) \cap \text{wedge}(e)$. On the other hand, wedge(c) must intersect ae, since e lies to its right (because d is reflex) and similarly a lies to its left. Let $a'e'$ be the intersection of wedge(c) with ae (see Figure 20). Since $a'e' \subseteq \text{kernel}(P)$ and $a'e' \in L(b, d)$ (because b belongs to segment ca' and d belongs to segment ce'), it follows that $\text{kernel}(P) \cap R(a, e) \cap L(b, d) \neq \emptyset$. Place a Steiner point s in the region, and connect it to a and e. The quadrangle *asef* is convex: a is convex because $s \in \text{wedge}(a)$, and s is convex because $s \in R(a, e)$. The hexagon *abcdes* is of the *rcrcrc*-1 type because s, b, d are mutually visible (since $s \in L(b, d)$).

2.2. (*rrcrcc—3*) If e sees neither a nor b, then three Steiner points suffice.

In fact, we can reduce the problem to the previous one, after adding one Steiner point s in the region $\text{wedge}(e) \cap R(f, d)$ (see Figure 21), which must be non-empty. The point s can then be connected to f and d. The quadrangle *sdef* is convex: d is convex because $s \in \text{wedge}(d)$, and s is convex because $s \in R(f, d)$. The remaining hexagon is of the kind *rrcrcc*-2, since d can see both a and b, because $s \in \text{wedge}(e)$.

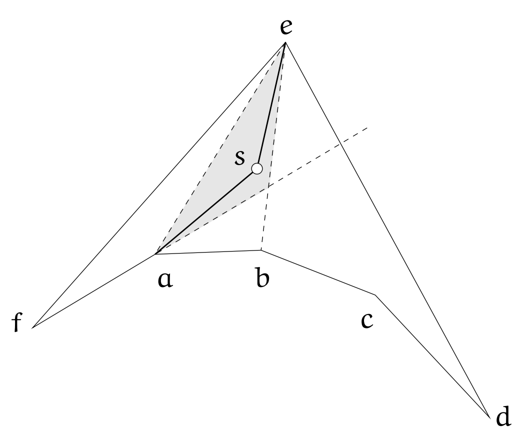

Figure 22: One Steiner point reduces the problem to the *rrcrcc*-2 case.

3. (*rrrccc*) We are left with the case in which the three reflex vertices are consecutive. This case can be solved with three Steiner points. In fact, it can be reduced to the *rrcrcc*-2 case after adding one Steiner point. Suppose that the three reflex vertices are a, b and c. Place a Steiner point s in the region $\text{wedge}(a) \cap R(a, e) \cap L(b, e)$, which is trivially non-empty. Connecting s with a and e gives rise to the convex quadrangle $asef$: a is convex because $s \in \text{wedge}(a)$, and s is convex because $s \in R(a, e)$. The remaining hexagon is of the *rrcrcc*-2 type, since e sees b and c, because $s \in L(b, e)$ (see Figure 22) .

This completes the proof of Theorem 3.2. It remains to consider the case when the union of two quadrangles is not a hexagon.

3.1.3 Quadrangle with one interior point.

As stated earlier, when two quadrangles share two edges, their union is a quadrangle which contains one of the vertices of the original quadrangles in its interior. We will show that three Steiner points suffice to convex-quadrangulate this polygon, thus establishing the following theorem:

3.8. Theorem. *Any union of two quadrangles can be convex-quadrangulated with at most three Steiner points.*

Proof. We consider here only the case where the union is not a hexagon. Let us call the four vertices of the union quadrangle r, a, b and c, where r is the only (possibly) reflex vertex. Let i be the interior point. Since only r may be reflex, i must see either a or c , because r cannot obstruct its view to both. Suppose that i sees a, as illustrated in Figure 23. Since $i \in \text{wedge}(a)$,

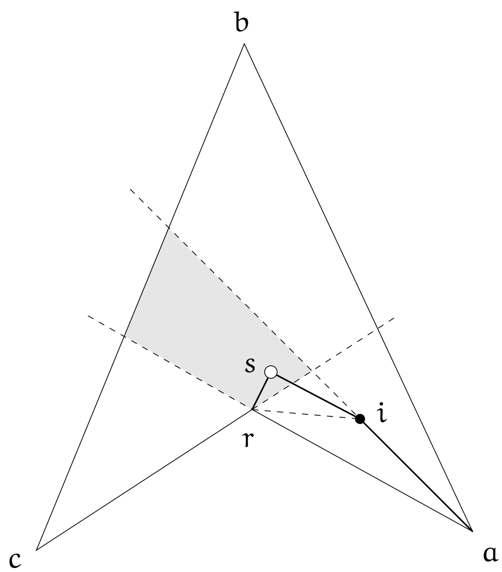

Figure 23: One Steiner point reduces the problem to the *rrcccc* case.

$\text{wedge}(r) \cap L(r, i) \cap L(a, i) \neq \emptyset$. Place one Steiner point s in the region. Then the quadrangle *rais* in convex: r is convex because $s \in \text{wedge}(r)$, i is convex because $s \in L(a, i)$, and s is convex because $s \in L(r, i) \cap \text{wedge}(r)$. On the other hand, the hexagon *siabcr* is a *rrcccc* hexagon, which can be convex-quadrangulated with two Steiner points. ∎

Each of the cases described in this section runs in constant time, thus:

3.9. Theorem. *A strictly convex quadrilateral mesh of* n *points using at most* $3\lfloor\frac{n}{2}\rfloor$ *Steiner points can be computed in* $O(n \log n)$ *time.*

4 Concluding Remarks

We have given upper and lower bounds on the number of Steiner points required to construct a convex quadrangulation for a planar set of points. Both bounds are constructive, and the upper bound yields a straightforward $O(n \log n)$ time algorithm. The obvious open problem is that of reducing the gap between the lower and upper bounds. One way to reduce the upper bound may be by constructing a convex quadrangulation of the point set directly, rather than by converting a triangulation (by combining triangles and then quadrangles) as we do now. Also, it would be interesting to explore the possibility of improving (raising) the lower bound for a non-degenerate point set by combining in some way the two point set configurations given in Section 2.

References

[1] D. J. Allman. A quadrilateral finite element including vertex rotations for plane elasticity analysis. *International Journal for Numerical Methods in Engineering*, 26:717–730, 1988.

[2] E. Arkin, M. Held, J. Mitchell, and S. Skiena. Hamiltonian triangulations for fast rendering. In J. van Leeuwen, editor, *Algorithms-ESA'94*, LNCS 855, pages 36–47, Utrecht, The Netherlands, September 1994.

[3] S. Benzley, E. Perry, K. Merkley, B. Clark, and K. Sjaardema. A comparison of all-hexahedral and all-tetrahedral finite element meshes for elastic and elasto-plastic analysis. In *4th Int. Meshing Roundtable*, pages 179–191, 1995.

[4] M. Bern and D. Eppstein. Quadrilateral meshing by circle packing. In *6th International Meshing Roundtable*, pages 7–19, 1997.

[5] M. Bern, S. A. Mitchell, and J. Ruppert. Linear-size nonobtuse triangulation of polygons. *Discrete & Computational Geometry*, 14:411–428, 1995.

[6] P. Bose, S. Ramaswami, A. Turki, and G. Toussaint. Experimental comparison of quadrangulation algorithms for sets of points. In *Twelfth European Workshop on Computational Geometry*, Münster, Germany, 1996.

[7] P. Bose and G. Toussaint. Characterizing and efficiently computing quadrangulations of planar point sets. *Computer Aided Geometric Design*, 14:763–785, 1997.

[8] C. K. Chui and M.-J. Lai. Filing polygonal holes using C^1 cubic triangular spline patches. *Computer Aided Geometric Design*, 17:297–307, 2000.

[9] H. Edelsbrunner, J. O'Rourke, and E. Welzl. Stationing guards in rectilinear art galleries. *Computer Vision, Graphics and Image Processing*, 27:167–176, 1984.

[10] D. Eppstein. Linear complexity hexahedral mesh generation. In *Proc. of the 12th ACM Symposium on Computational Geometry*, pages 58–67, 1996.

[11] H. Everett, W. Lenhart, M. Overmars, T. Shermer, and J. Urrutia. Strictly convex quadrilateralizations of polygons. In *Proc. of the 4th Canadian Conference on Computational Geometry*, pages 77–82, St. Johns, Newfoundland, 1992.

[12] T. Fevens, H. Meijer, and D. Rappaport. Minimum weight convex quadrilateralization of a constrained point set. In *Second CGC Workshop on Computational Geometry*, Durham, NC, USA, 1997.

[13] B. P. Johnston, J. M. Sullivan, and A. Kwasnik. Automatic conversion of triangular finite meshes to quadrilateral elements. *International Journal of Numerical Methods in Engineering*, 31(1):67–84, 1991.

[14] J. Kahn, M. Klawe, and D. Kleitman. Traditional galleries require fewer watchmen. *SIAM Journal of Algorithms and Discrete Methods*, 4(2):194–206, June 1983.

[15] M.-J. Lai. Convex preserving scattered data interpolation using bivariate C^1 cubic splines. *J. Comput. Applied Math.*, 119:249–258, 2000.

[16] M.-J. Lai and L. L. Schumaker. Scattered data interpolation using C^2 supersplines of degree six. *SIAM Journal on Numerical Analysis*, 34(3):905–921, 1997.

[17] A. Lubiw. Decomposing polygonal regions into convex quadrilaterals. In *Proc. of the 1st ACM Symposium on Computational Geometry*, pages 97–106, 1985.

[18] B. M. and D. Eppstein. Polynomial-size nonobtuse triangulation of polygons. *International Journal of Computational Geometry and Applications*, 2:241–255, 1992.

[19] S. Mitchell. A characterization of the quadrilateral meshes of a surface which admit a compatible hexahedral mesh of the enclosed volume. In *5th MSI Workshop on Computational Geometry*, 1995.

[20] S. Mitchell. Hexahedral mesh generation via the dual. In *Proc. of the 11th ACM Symposium on Computational Geometry*, pages C4–C5, 1995.

[21] M. Müller-Hannemann and K. Weihe. Quadrangular refinements of polygons with an application to finite-element meshes. In *Proc. of 13th ACM Symposium on Computational Geometry*, 1997.

[22] S. Ramaswami, P. Ramos, and G. Toussaint. Converting triangulations to quadrangulations. *Computational Geometry: Theory and Applications*, 9:257–276, 1998.

[23] J. R. Sack. An $O(n \log n)$ algorithm for decomposing simple rectilinear polygons into convex quadrilaterals. In *Proc. 20th Annual Allerton Conference*, pages 64–75, October 1982.

[24] E. Schönhardt. Über die Zerlegung von Dreieckspolyedern in Tetraeder. *Math. Annalen*, 98:309–312, 1928.

[25] G. Toussaint. Quadrangulations of planar sets. In *Workshop on Algorithms and Data Structures*, Lecture Notes in Computer Science, pages 218–227. Springer-Verlag, August 1995.

[26] R. Wenger. Helly type theorems and geometric transversals. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, Discrete Mathematics and its Applications, chapter 4. CRC Press, 1997.