

Minimum Feature Size in Proper Triangulations

Youssef Chouay

Thesis submitted to the University of Ottawa in partial fulfillment
of the requirements for the Master of Computer Science

Supervisor: Vida Dujmović

School of Electrical Engineering and Computer Science
Faculty of Engineering
University of Ottawa

© Youssef Chouay, Ottawa, Canada, 2026

Abstract

Triangulating a polygon can introduce distances between vertices and nonincident edges that are much smaller than those in the input. The *minimum feature size* of a crossing-free geometric graph is the smallest distance between a vertex and a nonincident edge. This measure applies to both a polygon and its triangulations, and the ratio of the input minimum feature size to that of the triangulation is the *degradation*. Even when Steiner vertices are allowed, some degradation remains unavoidable for proper triangulations, in which every bounded face has exactly three edges on its boundary.

We prove that, for every sufficiently large N , there is an N -vertex polygon P such that every proper triangulation G of P has degradation $\Omega(\log N)$. The bound holds even when any finite number of Steiner vertices may be added in the bounded face of P and on any of its edges. Together with the known $O(\log N)$ upper bound, it establishes that the worst-case optimal degradation for proper triangulations of N -vertex polygons is $\Theta(\log N)$. These polygons can also be chosen so that achieving $O(\log N)$ degradation requires $\Omega(N)$ Steiner vertices. This matches the known $O(N)$ upper bound on the number of Steiner vertices.

Acknowledgements

Although this work bears my name, it owes much to the people whose care, guidance, and encouragement have helped me along the way. I am especially thankful to my supervisor, Dr. Vida Dujmović, whose patience, thoughtful guidance, and detailed feedback have shaped this project and my growth as a researcher. To my parents, Driss and Fatiha, my sister, Jihane, and my two loving nieces, thank you for the love and steady support that have accompanied me throughout this journey. To my partner, Amani, your care and encouragement have sustained me through countless long nights. Having you by my side has meant more than I can express. Finally, to my dear friend Yassine, thank you for the camaraderie we share and for helping me grow, both in academia and in life. I would not be the researcher, student, or person I am today without you all.

Contents

Abstract	ii
Acknowledgements	iii
List of Figures	v
List of Tables	vi
1 Introduction	1
1.1 Other measures of triangulation quality	2
1.2 Bounds on feature-size degradation	2
1.3 The results of this thesis	4
2 Terminology and notation	6
2.1 Geometric graphs and faces	6
2.2 Polygons and subdivisions	7
2.3 Triangulations	7
2.4 Minimum feature size and degradation	8
2.5 Notation used in the proof	9
3 The logarithmic lower bound	11
3.1 The polygon family Q_n and Theorem 2	11
3.2 Proof of Theorem 2	13
3.3 The number of Steiner vertices	25
4 Conclusions	30
Bibliography	31

List of Figures

2.1	Classical, proper, and nonproper triangulations	8
3.1	The polygon family Q_n	12
3.2	The rectangular polygon and a constant-degradation triangulation	13
3.3	Tiles and their leftmost horizontal sides	14
3.4	The auxiliary rectangle in Q_n	17
3.5	A tile and its extension	18
3.6	Finding a tile in a face	19
3.7	Finding a tile upside down	20
3.8	Looking for a vertex in $\mathcal{E}(\mathcal{B})$	22
3.9	Looking for the new tile in $\mathcal{E}(\mathcal{B})$	23
3.10	Counting successive tiles	25
3.11	Bounding the number of Steiner vertices in Q_n	26
3.12	The ultimate lower bound	28

List of Tables

1.1 Degradation in the three triangulation models	5
---	---

1. Introduction

Triangulating polygons is a fundamental and widely used operation in discrete and computational geometry, serving as a building block in areas such as computer graphics, geometric modeling, mesh generation, and numerical simulation. In these applications, the quality of a triangulation often matters and is assessed using different geometric criteria, depending on the requirements of the application. Common quality criteria include avoiding very small or very large angles, bounding triangle aspect ratios, maximizing the minimum triangle height, and minimizing the total edge length.

This thesis focuses on the preservation of geometric separation as a criterion for triangulation quality. Triangulating a polygon can introduce distances between vertices and nonincident edges that are much smaller than those present in the input. We study how well this separation can be preserved, using minimum feature size to compare the polygon with its triangulation.

A geometric graph is a graph where vertices are points in the plane and edges are straight-line segments joining their endpoints. For a crossing-free geometric graph G , the *minimum feature size*, denoted by $\text{mfs}(G)$, is the smallest distance between a vertex and a nonincident edge. Both a polygon and a triangulation of that polygon are geometric graphs, so this measure applies to both the input and the output.

One motivation for preserving minimum feature size is to control sensitivity to coordinate errors. When a triangulation is stored or manipulated, its vertex coordinates may be rounded or perturbed. If a vertex lies very close to a nonincident edge, even a small perturbation can create an intersection and destroy the crossing-free representation. Preserving minimum feature size therefore helps ensure that triangulation does not make the geometry substantially more sensitive to coordinate errors than the input polygon.

Small vertex–edge distances can also make distinct parts of a drawing difficult to distinguish on a finite-resolution display. Aloupis et al. [1, 3] discuss these motivations and applications of feature-size-preserving triangulations to geometric dissection. Feature-size preservation is also closely related to the problem of controlling the minimum triangle height, studied by Bern, Dobkin, and Eppstein [4].

To quantify the loss of geometric separation, we define the *degradation* of a triangulation as the ratio of the input polygon’s minimum feature size to that of the triangulation. Smaller

degradation means better preservation of feature size. This measure concerns the whole geometric graph. The closest vertex and nonincident edge need not belong to the same triangular face. We ask how much degradation is unavoidable in the worst case, even when any finite number of additional vertices may be inserted.

1.1 Other measures of triangulation quality

A substantial body of work concerns bounds on triangle angles. Delaunay refinement algorithms add vertices to obtain angle guarantees under assumptions on the input geometry [9, 10, 11]. Bishop [7] showed that polygons admit triangulations with angles between 30° and 75° , except in one triangle at each input corner smaller than 30° . He also determined the optimal lower and upper angle bounds for triangulations with additional vertices [8].

Such angle guarantees do not by themselves imply a bound on feature-size degradation. For example, repeatedly subdividing an equilateral triangle into smaller equilateral triangles preserves every angle while making the minimum feature size arbitrarily small relative to the input.

The surveys of Bern and Eppstein [5] and Bern, Shewchuk, and Amenta [6] discuss these and other quality criteria and their applications. Of particular relevance here is the minimum triangle height. Relating this quantity to feature size requires distinguishing distances within the polygon from distances across its exterior.

1.2 Bounds on feature-size degradation

Vertices added when triangulating a polygon are called *Steiner vertices*.

We distinguish two triangulation models according to whether a side of a triangular face must be a single graph edge or may consist of several collinear graph edges. In a *proper triangulation*, each triangular face is bounded by exactly three graph edges. In a *nonproper triangulation*, the boundary of a face still has the shape of a triangle, but a side of that triangle may consist of several collinear graph edges. The latter model permits a vertex of one triangle to lie in the relative interior of a side of another. A proper triangulation using only input vertices is called *classical*. Chapter 2 makes these conventions precise.

Bern, Dobkin, and Eppstein [4, Section 2] give an example, attributed there to Mitchell, showing that the best achievable degradation in a classical triangulation is unbounded even for polygons with only six vertices. Thus Steiner vertices are necessary for any finite worst-case degradation bound depending only on the number of input vertices. Section 3.3 describes their rectangle example and uses it to obtain a lower bound on the number of Steiner vertices.

An important earlier result concerns distances measured within the polygon. Bern, Dobkin, and Eppstein [4, Theorem 1] proved that every polygon, possibly with holes, admits

a proper triangulation in which every triangle height is at least a constant multiple of the input’s *internal feature size*.¹ Here, internal feature size means the minimum length of a segment entirely contained in the polygonal region that joins an input vertex to a nonincident input edge. For a triangulation, internal feature size is defined in the same way using its vertices and edges, with segments still required to lie in the input polygonal region. For a proper triangulation, this quantity equals the smallest triangle height. Thus their theorem guarantees constant degradation of internal feature size. Their construction allows Steiner vertices both inside the polygon and on its edges.

Extending this result to 2-connected crossing-free geometric graphs, where every bounded face is to be triangulated, introduces two difficulties. Each face of such a graph is bounded by a polygon. First, preserving internal feature size within each face does not control vertex–edge distances when the connecting segment passes outside that face. Second, Steiner vertices inserted on a shared edge may reduce feature size in an adjacent face and prevent independently constructed triangulations from fitting together properly. Processing the faces successively does not resolve this difficulty. Each construction may introduce subdivision vertices that reduce the input feature size for subsequent constructions, allowing the degradation to accumulate.

These difficulties motivate seeking a polygon triangulation that preserves minimum feature size and introduces no vertices on the polygon’s boundary. The first property controls distances across the polygon’s exterior as well as its interior. The second leaves shared edges unchanged. Together, they allow the construction to be applied independently to the bounded faces of a 2-connected crossing-free geometric graph. Bern, Dobkin, and Eppstein [4, Section 6] posed the extension of their result to crossing-free geometric graphs as an open problem.

Aloupis et al. [3] obtained a construction with both properties. Specifically, every N -vertex polygon admits a proper triangulation with $O(\log N)$ minimum-feature-size degradation, using only interior Steiner vertices. Their construction therefore extends to 2-connected crossing-free geometric graphs. However, this upper bound is logarithmic, whereas Bern, Dobkin, and Eppstein’s guarantee for internal feature size is constant. This raises the question of whether the logarithmic loss is unavoidable or whether a different construction could guarantee smaller degradation, possibly by allowing Steiner vertices on the polygon’s boundary.

Aloupis et al. [1] proved an $\Omega(\log N / \log \log N)$ lower bound and later announced a matching $\Omega(\log N)$ bound without providing its proof [2]. This thesis gives a complete proof of the logarithmic lower bound. The bound holds even when Steiner vertices may be inserted on the polygon’s edges as well as in its bounded face. Thus allowing boundary subdivisions does not improve the worst-case asymptotic degradation.

¹Bern, Dobkin, and Eppstein call this quantity *feature size*.

For nonproper triangulations, Aloupis et al. [3, Theorem 1] proved that every N -vertex polygon admits a triangulation with $O(1)$ minimum-feature-size degradation. Both their proper and nonproper constructions use $O(N)$ Steiner vertices and leave the polygon's edges unsubdivided.

1.3 The results of this thesis

For a polygon P and a triangulation G of P , write $\text{de}_P(G) = \text{mfs}(P) / \text{mfs}(G)$ for the degradation of G relative to P . The following theorem is the main result of this thesis.

Theorem 1. *There exist an absolute constant $c > 0$ and an integer N_0 such that, for every integer $N \geq N_0$, there is a polygon P with exactly N vertices for which every proper triangulation G of P satisfies*

$$\text{de}_P(G) \geq c \log_2 N.$$

Steiner vertices may be added in the bounded face of P and on any of its edges, with no restriction on their number other than finiteness.

Let $\mathcal{T}(P)$ be the class of all proper triangulations of P , allowing Steiner vertices both in the bounded face of P and on its edges. For an integer $N \geq 3$, define the *worst-case optimal degradation* for N -vertex polygons by

$$\sup_{\substack{P \text{ a polygon} \\ |V(P)|=N}} \inf_{G \in \mathcal{T}(P)} \text{de}_P(G).$$

The infimum is taken over all proper triangulations of a fixed polygon, and the supremum over all polygons with N vertices.

Corollary 1.1. *The worst-case optimal minimum-feature-size degradation for proper triangulations of N -vertex polygons is $\Theta(\log N)$.*

Proof. Combine Theorem 1 with the $O(\log N)$ upper bound of Aloupis et al. [3, Theorem 2]. □

The worst-case optimal degradation for properly triangulating all bounded faces of N -vertex 2-connected crossing-free geometric graphs is also $\Theta(\log N)$. Aloupis et al. [3, Section 1] extend their upper bound to these inputs by applying their construction separately inside each bounded face. The matching lower bound already holds for polygons. The worst-case optimal degradation remains $\Theta(\log N)$ if input edges must remain unsubdivided, since the upper-bound construction leaves them unsubdivided and Theorem 1 allows their subdivision.

Triangulation	Degradation	Source
Classical	Unbounded even for $N = 6$	Bern, Dobkin, and Eppstein [4, Section 2], attributed to Mitchell
Proper	$\Theta(\log N)$ in the worst case	Lower bound: Theorem 1. Upper bound: [3]
Nonproper	$\Theta(1)$ in the worst case	Constant upper bound: [3]

Table 1.1: Worst-case optimal minimum-feature-size degradation in the three triangulation models.

In contrast, nonproper triangulations have worst-case optimal degradation $\Theta(1)$. Table 1.1 summarizes these bounds for polygons.

We also prove a matching worst-case lower bound on the number of Steiner vertices. Proposition 3.3 gives a family of N -vertex polygons for which every proper triangulation has degradation $\Omega(\log N)$, and every proper triangulation achieving $O(\log N)$ degradation requires $\Omega(N)$ Steiner vertices. Together with the $O(N)$ upper bound of Aloupis et al. [3], this shows that a linear number of Steiner vertices is necessary and sufficient in the worst case for logarithmic degradation.

Chapter 2 gives the definitions and notation used throughout the thesis. Chapter 3 constructs the polygon family and proves Theorem 1. It then bounds the number of Steiner vertices for this family and proves the stronger worst-case bound of Proposition 3.3. Chapter 4 summarizes the results and outlines a direction for further research.

2. Terminology and notation

We now give precise definitions for the objects and quantities discussed in the introduction. All graphs are finite, all distances are Euclidean unless stated otherwise, and all logarithms have base two. An input polygon may have consecutive collinear vertices. These count as input vertices. Unless stated otherwise, a triangulation may subdivide any input edge.

2.1 Geometric graphs and faces

For a point $p \in \mathbb{R}^2$, write $p = (p_x, p_y)$ for its horizontal and vertical coordinates. For two points p, q , let pq be the closed straight-line segment joining them and $|pq|$ its length. For a set $S \subseteq \mathbb{R}^2$, write $\overline{S}$ for its closure, $\text{int } S$ for its interior, and $\partial S = \overline{S} \setminus \text{int } S$ for its boundary.

A *geometric graph* $G = (V(G), E(G))$ consists of a finite set of points $V(G) \subset \mathbb{R}^2$ and a set $E(G)$ of straight-line segments joining pairs of distinct vertices. A vertex is *incident* to an edge if it is an endpoint of that edge. The graph is *crossing-free* if no vertex lies in the relative interior of an edge and any two distinct edges intersect only at a common endpoint, if they intersect at all.

All geometric graphs in this thesis are crossing-free. The *underlying abstract graph* of G retains its vertices and their adjacencies, disregarding the coordinates. Connectedness, 2-connectivity, paths, cycles, and vertex degrees refer to this underlying abstract graph. The vertices and edges of G together form the point set

$$V(G) \cup \bigcup_{e \in E(G)} e.$$

We also denote this set by G . Each edge includes its endpoints; including $V(G)$ separately accounts for isolated vertices. Different geometric graphs can have the same point set. For example, inserting a vertex in the interior of an edge and replacing that edge by the two resulting segments changes the graph's vertex and edge sets without changing the point set.

A *face* of a geometric graph G is a connected component of $\mathbb{R}^2 \setminus G$. Its unique unbounded face is denoted by $F_\infty(G)$ and is also called the *outer face*. For a face F , let $B_G(F)$ be the

geometric graph consisting of the vertices and edges of G on ∂F :

$$V(B_G(F)) = V(G) \cap \partial F, \quad E(B_G(F)) = \{e \in E(G) : e \subseteq \partial F\}.$$

A vertex or edge of G is *incident to* F if it belongs to $B_G(F)$. If the underlying abstract graph of $B_G(F)$ is a cycle, then $B_G(F)$ is the *boundary cycle* of F .

Faces are open, so $\partial F = \overline{F} \setminus F$. As point sets, $B_G(F) = \partial F$. For a general geometric graph, $B_G(F)$ need not be a cycle. For example, a graph consisting of a single edge and its endpoints is the boundary of its only face. The polygons and triangulations considered below have a boundary cycle for every face.

2.2 Polygons and subdivisions

A *polygon* is a geometric graph P whose underlying abstract graph is a cycle with at least three vertices. By the Jordan curve theorem, P has exactly two faces: a bounded face, denoted by F_P , and the unbounded face $F_\infty(P)$. Its closed *polygonal region* is $\Omega(P) = \overline{F_P}$. Both faces have P as their boundary cycle. In particular,

$$B_P(F_P) = P, \quad \partial\Omega(P) = P, \quad \text{int } \Omega(P) = F_P.$$

An edge of P is an edge of the geometric graph P . A *corner* is a vertex of P whose two incident edges are not collinear.

A *side* of a polygon P is a nondegenerate straight-line segment that is maximal, under inclusion, among the straight-line segments contained in P . A side is therefore a maximal union of consecutive collinear edges. Its endpoints are corners, and it can contain other vertices between them. An edge and a side need not be the same segment.

A geometric graph G' is a *geometric subdivision* of a geometric graph G if G' can be obtained from G by inserting finitely many additional vertices in the relative interiors of edges of G , with each edge replaced by the segments between consecutive vertices, including its original endpoints. The added vertices are called *subdivision vertices*. Inserting no vertices is allowed.

A subdivided edge is a path in the underlying abstract graph. Its new vertices have degree two in that path, but may have larger degree in a graph containing the path. Thus a subdivision vertex of an input polygon edge may be incident to several edges in a triangulation of the polygon.

2.3 Triangulations

A *triangulation* of a polygon P is a connected geometric graph G such that:

1. The outer face of G has a boundary cycle that is a geometric subdivision of P .
2. For every bounded face F of G , the boundary ∂F consists of the three sides of a nondegenerate triangle.

The vertices in $V(G) \setminus V(P)$ are the *Steiner vertices* of G relative to P .

These conditions imply that $G \subseteq \Omega(P)$ and $F_\infty(G) = F_\infty(P)$. The outer boundary of G is the same simple closed curve as P . The closures of the bounded faces of G cover $\Omega(P)$. A Steiner vertex on P is a subdivision vertex of an input edge. Every other Steiner vertex lies in F_P .

For a face F whose boundary is a polygon, a *side* of F means a side of $B_G(F)$. In a triangulation, each bounded face has three sides, but a side may consist of more than one graph edge. Following Aloupis et al. [2, 3], we distinguish:

- A *proper triangulation* is a triangulation in which $B_G(F)$ has exactly three edges for every bounded face F .
- A *classical triangulation* is a proper triangulation with $V(G) = V(P)$.
- A *nonproper triangulation* is any triangulation, with no requirement that a side of a bounded face be a single graph edge. This class includes proper triangulations.

Figure 2.1 illustrates the three models. In the nonproper example, the vertex v lies on a side of the upper-left triangular face without being a corner of that face.




Figure 2.1: All input edges remain unsubdivided. Black points are input vertices. Blue points are Steiner vertices in the bounded face of the input polygon.

Proper triangulations may subdivide input edges. Several triangular faces may be incident to consecutive edges of a subdivided input edge, while each of those faces still has exactly three graph edges. Conversely, a vertex in the relative interior of a side of a nonproper triangular face need not be a Steiner vertex. It may be a collinear vertex of P .

2.4 Minimum feature size and degradation

For a point v and a closed segment or line S , let $\text{dist}(v, S) = \min_{q \in S} |vq|$.

The *minimum feature size* of a geometric graph G is

$$\text{mfs}(G) = \min_{\substack{v \in V(G), e \in E(G) \\ v \text{ not incident to } e}} \text{dist}(v, e).$$

If there is no nonincident vertex–edge pair, let $\text{mfs}(G) = +\infty$. Recall that the degradation of a triangulation G relative to its input polygon P is

$$\text{de}_P(G) = \frac{\text{mfs}(P)}{\text{mfs}(G)}.$$

For polygons and their triangulations, minimum feature size is strictly positive and finite. Strict positivity follows from the finiteness of the graph and the fact that no vertex belongs to a nonincident edge. The distance in the definition is measured in the plane. The segment realizing a vertex–edge distance need not lie in the polygonal region.

Every triangulation G of P satisfies $\text{mfs}(G) \leq \text{mfs}(P)$, even when input edges are subdivided. To see this, let v and e be a nonincident input pair realizing $\text{mfs}(P)$, and let $q \in e$ be closest to v . Since e is a union of edges of G , there exists an edge e' of G such that $e' \subseteq e$ and $q \in e'$. Since $v \notin e$, this edge is nonincident to v , and

$$\text{mfs}(G) \leq \text{dist}(v, e') = \text{dist}(v, e) = \text{mfs}(P).$$

Consequently $\text{de}_P(G) \geq 1$.

2.5 Notation used in the proof

The following table summarizes notation used in Chapter 3. The geometric graph G and polygon P in the definitions above are arbitrary. In Subsection 3.2.2, G is fixed to be a proper triangulation of one polygon Q_n . The parameters and auxiliary regions in the last rows are introduced during the proof. The rectangle W is chosen in Subsection 3.2.4 and remains fixed throughout the counting argument.

Symbol	Meaning
p_x, p_y	Horizontal and vertical coordinates of a point p .
$\overline{S}, \text{int } S$	Closure and interior of a set $S \subseteq \mathbb{R}^2$.
$G, V(G), E(G)$	A geometric graph, its vertices, and its edges.
$F, \overline{F}$	A face of G and its closure $F \cup \partial F$.
$F_\infty(G)$	The unbounded (outer) face of G .
∂F	Boundary of the face F .
Q_n, h	Notched polygon and the parameter $h = n^2$.

Symbol	Meaning
$F_{Q_n}, \Omega(Q_n)$	Bounded face of Q_n and its closure.
N	Number of vertices of the input polygon.
$\text{dist}(v, S)$	Distance from the point v to the closed segment or line S .
$\text{mfs}(G)$	Minimum feature size of the graph G .
$\text{de}_P(G)$	Degradation $\text{mfs}(P)/\text{mfs}(G)$.
s, δ	Minimum feature size $s = \text{mfs}(G)$ and length $\delta = s/3$ of each horizontal side of a tile.
$\mathcal{B}, \text{ht}(\mathcal{B})$	A tile and its height.
$x_{\min}(\mathcal{B})$	Least horizontal coordinate of a tile.
$\mathcal{E}(\mathcal{B}), \rho^-, \rho^+$	Extension of a tile and its boundary rays.
W, a, b, m	Rectangle $W = [0, 1] \times [a, b]$ and its height $m = b - a$.
μ, d	Minimum feature size of the current input polygon and degradation of its triangulation.

3. The logarithmic lower bound

We prove Theorem 1 using a family of notched polygons. Section 3.1 defines the family Q_n , computes its minimum feature size, and states its logarithmic lower bound as Theorem 2. Section 3.2 proves this theorem. After the sketch in Subsection 3.2.1, Subsection 3.2.2 states the main counting proposition and uses it to prove Theorems 2 and 1. The remaining subsections establish the proposition. Section 3.3 then bounds the number of Steiner vertices needed to achieve logarithmic degradation on this family and proves a stronger bound for general polygons.

3.1 The polygon family Q_n and Theorem 2

For each integer $n \geq 8$, let $h = n^2$ and let Q_n be the polygon with corners

$$\begin{aligned} A &= (-1, -h), & V &= (0, 0), & B &= (0, -h), \\ C &= (1, -h), & D &= (1, 1), & E &= (-1, 1) \end{aligned}$$

in cyclic order, and the vertices

$$q_j = (1, 1 - j), \quad j = 1, \dots, n,$$

on its right vertical side DC .


Figure 3.1 illustrates Q_n and its exterior triangular notch, formed by the sides AV and VB .

The vertices on DC divide that side into the edges

$$Dq_1, q_1q_2, \dots, q_{n-1}q_n, q_nC.$$

The first n edges have length one. The last edge q_nC has length $h + 1 - n$. Each of the other five sides is a single input edge. Thus Q_n has six sides, six corners, and $n + 6$ vertices and edges.

(a) Entire polygon



(b) Upper part of the polygon




Figure 3.1: The polygon Q_n for $n = 8$ and $h = 64$. Figure 3.1(a) shows the entire polygon with its vertical scale compressed. Figure 3.1(b) enlarges the upper part. The shaded region is the polygon's interior, and the white triangular notch is exterior. Black dots are input vertices. Consecutive marked vertices near the top of DC are one unit apart.

Lemma 3.1. *The polygon Q_n satisfies*

$$\text{mfs}(Q_n) = \frac{h}{\sqrt{h^2 + 1}}, \quad \frac{1}{2} \leq \text{mfs}(Q_n) < 1.$$

Proof. Triangle ABV is right-angled at B , so the distance from B to its hypotenuse AV is the height with respect to AV . Since $|AB| = 1$, $|BV| = h$, and $|AV| = \sqrt{h^2 + 1}$,

$$\text{dist}(B, AV) = \frac{|AB| |BV|}{|AV|} = \frac{h}{\sqrt{h^2 + 1}}.$$

We show that all other nonincident vertex–edge pairs of Q_n have distance at least one. Every edge except AV is horizontal or vertical, all input vertices have integer coordinates, and no input vertex lies on a nonincident edge. For such an axis-parallel edge and a nonincident vertex, either the perpendicular distance to its supporting line is a positive integer, or a closest endpoint differs from the vertex by at least one in an integer coordinate. Thus the distance is at least one.

It remains to consider the other vertices nonincident to AV . The vertices C , D , and $q_1, \dots, q_n$ have x -coordinate one, whereas every point of AV has $x \leq 0$. Their distances to AV are therefore at least one. The only remaining vertex is $E = (-1, 1)$. Since every point of AV has $y \leq 0$, its distance to AV is also at least one.

Thus the pair B, AV realizes $\text{mfs}(Q_n)$. The bounds $1/2 \leq \text{mfs}(Q_n) < 1$ follow from $h \geq 1$. $\square$

Theorem 1 follows immediately from the following theorem, since Q_n has $n + 6$ vertices. Most of the remainder of this chapter is devoted to proving Theorem 2.

Theorem 2. *For every integer $n \geq 8$ and every proper triangulation G of Q_n ,*

$$\text{de}_{Q_n}(G) \geq \frac{\log_2 n}{36}.$$

3.2 Proof of Theorem 2

3.2.1 Proof sketch

We start by giving a sketch of the proof of Theorem 2. If one zooms in on the upper-right part of Q_n , near the vertices $q_1, \dots, q_n$, a rectangle of height n and width one fits inside the polygonal region. This is the region in which most of the proof takes place. For ease of explanation, we first sketch the argument for a simpler family of polygons W_n . Each W_n has four sides forming a rectangle of height n and width one. Its right side consists of n unit edges, and each of its other sides is a single edge. See Figure 3.2(a). Its minimum feature size is one.

Unlike Q_n , this polygon admits a triangulation with constant degradation if vertices may be added on its left side. Add vertices there at unit intervals, join them horizontally to the vertices on the right side, and triangulate each resulting unit square with a diagonal. See Figure 3.2(b). This triangulation has minimum feature size $1/\sqrt{2}$, and hence degradation $\sqrt{2}$. For the purpose of the sketch, we therefore forbid additional vertices on the boundary of W_n .




Figure 3.2: The polygon W_n and a proper triangulation with degradation $\sqrt{2}$, shown for $n = 8$ with the vertical scale compressed. Black dots are vertices of W_n . Blue dots are Steiner vertices on its left side.

Now let H be any proper triangulation of W_n with no additional vertices on its boundary, and let $\sigma = \text{mfs}(H)$. We find a sequence of auxiliary regions, called tiles. Each tile is a parallelogram with horizontal top and bottom sides of length $\sigma/3$, contained in the closure of a triangular face of H . Figure 3.3 shows three tiles and highlights each tile's leftmost horizontal side. For a tile with vertical sides, either horizontal side may be chosen.

Our construction ensures that the leftmost horizontal side of each subsequent tile lies entirely to the right of the preceding tile's leftmost horizontal side. Specifically, the left endpoint of the new tile's leftmost horizontal side has x -coordinate at least that of the right endpoint of the preceding tile's leftmost horizontal side. We show that the first tile can be found with height proportional to n . Whenever a tile has height greater than two, we further show that another tile can be found whose height is at least one eighth that of the preceding tile. Since only finitely many such horizontal sides of length $\sigma/3$ fit across the rectangle, the sequence must end with a tile of height at most two. Reaching a tile of height at most two requires $\Omega(\log n)$ tiles. Their leftmost horizontal sides, each of length $\sigma/3$, occupy horizontal intervals with disjoint interiors in a rectangle of width one, so $\sigma = O(1/\log n)$ and the degradation is $\Omega(\log n)$.




Figure 3.3: A proper triangulation of W_8 with one interior Steiner vertex and no additional vertices on its boundary. Each coloured tile lies in the closure of a triangular face. The bold segments indicate each tile's leftmost horizontal side. The vertical scale is compressed.

For Q_n , we allow subdivisions of every input edge. The notch ensures that, unless the degradation is already at least n , a long portion of VB remains in a single graph edge. Opposite it, subdividing the unit edges on DC only makes them shorter. These are the properties needed to carry out the same argument in a rectangle inside Q_n . We use the original triangular faces even when they extend beyond the top or bottom of that rectangle. We now make this argument precise for Q_n .

3.2.2 Proofs of Theorems 1 and 2 via the main proposition

For the proof of Theorem 2, fix an integer $n \geq 8$ and an arbitrary proper triangulation G of Q_n . Recall that $h = n^2$. We refer to the bounded faces of G as its *triangular faces*. Let

$$s = \text{mfs}(G), \quad \delta = s/3, \quad \mu = \text{mfs}(Q_n), \quad d = \text{de}_{Q_n}(G) = \mu/s.$$

Most of the remainder of this chapter is devoted to proving the following proposition, which implies Theorem 2.

Proposition 3.1. *Under the assumption $d < n$,*

$$n - 2 < 2 \cdot 8^{3/s}.$$

We first deduce the two theorems from Proposition 3.1.

Proof of Theorem 2. If $d \geq n$, the claimed bound is immediate. If $d < n$, Proposition 3.1 gives

$$n - 2 < 2 \cdot 8^{3/s} \implies \frac{1}{s} > \frac{1}{9} \log_2 \left(\frac{n-2}{2} \right).$$

Combining this with $\mu \geq 1/2$ from Lemma 3.1 yields

$$d = \frac{\mu}{s} \geq \frac{1}{18} \log_2 \left(\frac{n-2}{2} \right) \geq \frac{\log_2 n}{36}.$$

The last inequality follows from $(n-2)/2 \geq \sqrt{n}$ for integers $n \geq 8$. □

Proof of Theorem 1. Let $N \geq 14$ and $n = N - 6 \geq 8$. The polygon Q_n has exactly N vertices. By Theorem 2, every proper triangulation G of Q_n satisfies

$$\text{de}_{Q_n}(G) \geq \frac{\log_2 n}{36} \geq \frac{\log_2 N}{72},$$

where the last inequality follows from $N = n + 6 \leq n^2$ for $n \geq 8$. Thus Theorem 1 holds with $N_0 = 14$ and $c = 1/72$. □

To prove Proposition 3.1, we first bound the possible locations of subdivision vertices on VB in Subsection 3.2.3. Subsection 3.2.4 then uses this bound to choose an auxiliary rectangle W contained in $\Omega(Q_n)$. Subsection 3.2.5 defines tiles and shows how to construct one in the closure of a triangular face. Subsection 3.2.6 shows how to construct each subsequent tile, and Subsection 3.2.7 constructs the first tile. Finally, Subsection 3.2.8 combines these results to prove Proposition 3.1.

Every vertex, edge, and face in this construction belongs to the fixed triangulation G unless another graph is explicitly named. The rectangle W and its parameters are chosen

once in Subsection 3.2.4 and remain fixed throughout the count. The tiles are auxiliary regions. Constructing them does not change G .

3.2.3 Subdivision vertices on the side VB

Recall that $h = n^2$ and $d = \text{de}_{Q_n}(G)$ is the degradation of G relative to Q_n .

Lemma 3.2. *Let $p = (0, -t)$, $0 < t \leq h$, be a vertex of G on $VB \setminus \{V\}$. Then*

$$d \geq \frac{h}{t}, \quad \text{equivalently} \quad t \geq \frac{h}{d}.$$

Proof. Since p lies on the leg VB of the right triangle ABV , its orthogonal projection onto the hypotenuse AV lies on AV . Using $|Vp| = t$ and $|VB| = h$, similar triangles give

$$\text{dist}(p, AV) = \frac{t}{h} \text{dist}(B, AV) = \frac{t}{\sqrt{h^2 + 1}}.$$

Let q be the closest point on AV . Since G preserves each input edge as a union of graph edges, some edge $e \subseteq AV$ of G contains q . The vertex p is not on AV and hence is not incident to e . Consequently

$$s \leq \text{dist}(p, e) = \text{dist}(p, AV) = \frac{t}{\sqrt{h^2 + 1}}.$$

Using Lemma 3.1, we obtain

$$d = \frac{\text{mfs}(Q_n)}{s} \geq \frac{h/\sqrt{h^2 + 1}}{t/\sqrt{h^2 + 1}} = \frac{h}{t}.$$

□

3.2.4 The rectangle containing the tiles

For the tile construction and count, from this subsection through Subsection 3.2.8, assume $d < n$. By Lemma 3.2, any vertex of G on VB other than V has the form $p = (0, -t)$ with $t \geq h/d > n$. Thus the segment from $V = (0, 0)$ to $(0, -n)$ contains no subdivision vertices of VB .

Choose $1 - n < a < b < 0$ so that neither horizontal line $y = a$ nor $y = b$ contains a vertex of G , and so that $m := b - a > n - 2$. Such a choice is possible because G is finite. Let

$$W = [0, 1] \times [a, b].$$

Figure 3.4 shows W . Its left side lies on VB . The absence of subdivision vertices proved above implies that it is contained in a single edge of G . Its right side lies on the side DC of


Q_n , between q_1 and q_n . Each edge of G along the right side of W is contained in a unit edge of Q_n , so it has length at most one.

The coordinates of Q_n imply

$$[0, 1] \times [-h, 0] \subseteq \Omega(Q_n), \quad (0, 1) \times (-h, 0) \subseteq F_{Q_n}.$$

Thus the choice $1 - n < a < b < 0$ ensures that $W \subseteq \Omega(Q_n)$ and $\text{int } W \subseteq F_{Q_n}$. Moreover, $\Omega(Q_n) \subseteq \{(x, y) : x \leq 1\}$ implies that every point of G and of the closure of each triangular face has $x \leq 1$.

(a) Entire polygon



(b) Upper part of the polygon




Figure 3.4: The polygon Q_n and the auxiliary rectangle W (orange), shown for $n = 8$ and $h = 64$ with different horizontal and vertical scales. Black dots are input vertices, and the blue sides form the exterior notch. Figure 3.4(b) enlarges the upper part of the polygon.

3.2.5 Triangle heights and tiles

Recall that $s = \text{mfs}(G)$ is the minimum feature size of the fixed proper triangulation G .

Lemma 3.3. *Every height of every triangular face of G is at least s .*

Proof. In a nondegenerate triangle, the smallest height corresponds to a longest side. The orthogonal projection of the opposite vertex onto the line containing that side lies on the side itself. Thus the smallest height is the distance between a vertex of G and a nonincident edge of G , and is at least s . $\square$

Let $e \in E(G)$ be nonhorizontal, with supporting line $x = \alpha(y)$, where α is affine. A *tile* on e is a parallelogram with horizontal top and bottom sides

$$\mathcal{B} = \{(\alpha(y) + \tau, y) : y_0 \leq y \leq y_1, 0 \leq \tau \leq \delta\}, \quad y_0 < y_1,$$

whose left side is contained in e , such that $\mathcal{B} \subseteq \Omega(Q_n)$ and $\text{int } \mathcal{B} \cap G = \emptyset$. Neither line $y = y_0$ nor line $y = y_1$ may contain a vertex of G . The *height* and least horizontal coordinate of the tile are

$$\text{ht}(\mathcal{B}) = y_1 - y_0, \quad x_{\min}(\mathcal{B}) = \min_{y_0 \leq y \leq y_1} \alpha(y),$$

so $x_{\min}(\mathcal{B})$ is the x -coordinate of the left endpoint of the tile's leftmost horizontal side, as used in Subsection 3.2.1. The *extension* is

$$\mathcal{E}(\mathcal{B}) = \{(x, y) : y_0 \leq y \leq y_1, x \geq \alpha(y) + \delta\}.$$

Its lower and upper boundary rays are $\rho^- = \mathcal{E}(\mathcal{B}) \cap \{y = y_0\}$ and $\rho^+ = \mathcal{E}(\mathcal{B}) \cap \{y = y_1\}$. See Figure 3.5.

It follows from the definition of the extension that every point of $\mathcal{E}(\mathcal{B})$ has x -coordinate at least $x_{\min}(\mathcal{B}) + \delta$.




Figure 3.5: The tile $\mathcal{B}$ is shown in blue and its extension $\mathcal{E}(\mathcal{B})$ in green.

Lemma 3.4. *Let vr be a nonhorizontal edge of G , and let u be the third vertex of a triangular face F to the right of the supporting line of vr . Let $z \in vr \setminus \{v\}$. Suppose that either*

1. $v_y < r_y$ and $u_y \leq z_y$ (as depicted in Figure 3.6).
2. $v_y > r_y$ and $u_y \geq z_y$.

Then $\overline{F}$ contains a tile $\mathcal{B}$ on vr whose left side is contained in vz and whose height satisfies

$$\text{ht}(\mathcal{B}) = \frac{|z_y - v_y|}{4}.$$

Proof. First suppose that $v_y < r_y$ and $u_y \leq z_y$, and let $H = z_y - v_y$. Let p be any point in the middle third of vz . The horizontal cross-section of $\overline{F}$ through p is a segment pq , whose right endpoint q lies on ru or vu . We show that $|pq| \geq s/3$. See Figure 3.6.

If $q \in ru$, let ℓ be the line through ru . Since $|vp| \leq 2|vz|/3 \leq 2|vr|/3$, it follows that $|pr| \geq |vr|/3$. Let p_1 and v_1 be the orthogonal projections of p and v onto ℓ , respectively. If vr is not perpendicular to ℓ , the right triangles pp_1r and vv_1r are similar because they have the same angle at r . Thus $|pp_1|/|vv_1| = |pr|/|vr|$. Since $q \in \ell$,

$$|pq| \geq |pp_1| = \frac{|pr|}{|vr|} |vv_1| \geq \frac{s}{3},$$

where $|vv_1| \geq s$ by Lemma 3.3. If vr is perpendicular to ℓ , the same equality holds directly because $p_1 = v_1 = r$.

If $q \in vu$, then $v_y < p_y \leq u_y \leq z_y$. Let t be the point of vz with $t_y = u_y$. Since u is a vertex of G and vr is a nonincident edge, $|tu| \geq \text{dist}(u, vr) \geq s$. The triangles vpq and vtu are similar because pq and tu are horizontal. Hence

$$\frac{|pq|}{|tu|} = \frac{|vp|}{|vt|} \geq \frac{|vz|/3}{|vz|} = \frac{1}{3},$$

so $|pq| \geq s/3$ in this case as well.




Figure 3.6: The horizontal cross-section pq ends on ru in (a) and on vu in (b). In (a), the dashed green segments join p and v to their orthogonal projections p_1 and v_1 on ℓ . In (b), the dashed green segment tu is parallel to pq . The middle third of vz is shown in blue, and each blue tile has horizontal side length $\delta = s/3$ and height $(z_y - v_y)/4$.

For every p in the middle third of vz , the horizontal segment of length $\delta = s/3$ starting at p and extending rightward is therefore contained in $\overline{F}$. These segments form a parallelogram of height $H/3$. Choose a closed interval J of length $H/4$ strictly inside $(v_y + H/3, v_y + 2H/3)$,

with neither endpoint equal to the y -coordinate of a vertex of G . Such a choice is possible because G has only finitely many vertices. Restricting the parallelogram to points whose y -coordinates lie in J gives the required tile. Its interior lies in F and is therefore disjoint from G .

Now suppose that $v_y > r_y$ and $u_y \geq z_y$. Reflect the configuration across the x -axis. Explicitly, each point $p = (p_x, p_y)$ is sent to

$$(p_x, p_Y), \quad p_Y = -p_y.$$

Apply this reflection to the whole graph G , including the triangle and the point z on vr . See Figure 3.7. We retain the vertex labels and use the subscript Y for their reflected vertical coordinates. Then

$$v_Y < r_Y, \quad u_Y \leq z_Y.$$

Reflection preserves lengths and the direction of increasing x . The reflected graph therefore has the same minimum feature size s , and the reflected face lies to the right of the reflected edge vr . The construction in the first case applies to this reflected triangle, with Y in place of y . It produces a parallelogram with left side on the reflected segment vz and height

$$\frac{z_Y - v_Y}{4} = \frac{v_y - z_y}{4} = \frac{|z_y - v_y|}{4}.$$

Reflect this parallelogram back to obtain the required tile in $\overline{F}$, with left side on the original segment vz . $\square$




Figure 3.7: The downward case of Lemma 3.4 and its reflection across the x -axis (dashed).

3.2.6 Constructing successive tiles

Let F be a triangular face of G . Imagine moving a horizontal line upward through the closed triangle $\overline{F}$. For each line $y = c$ intersecting $\overline{F}$, the intersection is a horizontal segment, possibly a single point. Keep only the right endpoint of that segment (or the point itself if the intersection is a single point). As c ranges over a specified interval, these points form the *right part of ∂F* . This set is contained in at most two sides of the triangle.

Note that, for a tile $\mathcal{B}$, the inclusion $\mathcal{B} \subseteq W$ implies that every point of $\mathcal{E}(\mathcal{B})$ has $x \geq 0$ and $a \leq y \leq b$. Consequently, any point of $\mathcal{E}(\mathcal{B})$ with $x \leq 1$ lies in W . This applies in particular to points of G and of triangular face closures that lie in the extension.

Lemma 3.5. *Suppose that there exists a tile $\mathcal{B}$ on an edge of G such that $\mathcal{B} \subseteq W$. Suppose further that no graph edge along the right side of W intersects both boundary rays ρ^- and ρ^+ of $\mathcal{E}(\mathcal{B})$. Then there exists a triangular face of G with a vertex $v \in \mathcal{E}(\mathcal{B})$ strictly between these rays and such that one of the two edges of this face incident to v intersects ρ^- , and the other intersects ρ^+ .*

Proof. Let the lower and upper horizontal sides of $\mathcal{B}$ lie on the lines $y = y_0$ and $y = y_1$, respectively, with $y_0 < y_1$. Start with the face F containing $\text{int } \mathcal{B}$. This triangular face exists because the tile interior is connected, is contained in F_{Q_n} , and is disjoint from G . Its closure contains $\mathcal{B}$. Thus, on each line $y = c$ with $c \in [y_0, y_1]$, the right part of ∂F lies at or to the right of the right side of $\mathcal{B}$. Since this right part lies in $\mathcal{E}(\mathcal{B})$ and every point of G has x -coordinate at most 1, it lies in $\mathcal{E}(\mathcal{B}) \cap W$.

If the right part of ∂F between the rays contains a vertex v , then $v \in \mathcal{E}(\mathcal{B})$ and v lies strictly between the rays, since neither boundary line contains a vertex of G . Because G is proper, this right part consists of portions of the two edges of F incident to v . One intersects ρ^- and the other intersects ρ^+ , so the conclusion follows. See Figure 3.8(a).

Otherwise, this right part is a segment contained in a single edge e intersecting both rays. One endpoint of e lies strictly below the line $y = y_0$, and the other lies strictly above the line $y = y_1$, since neither line contains a vertex of G . The portion of e between the rays lies in W and strictly to the right of the left side of W . By hypothesis, the right side of W contains a vertex of G strictly between the rays. Hence e cannot lie along that side. The open segment between the intersections of e with ρ^- and ρ^+ is therefore contained in $\text{int } W$.

The edge e is therefore incident to a triangular face on its right. Cross e into that face, and let F now denote it. On every horizontal line between the rays, the right endpoint of the intersection with $\overline{F}$ lies strictly to the right of the intersection with e . Consequently, the new right part also lies in $\mathcal{E}(\mathcal{B}) \cap W$. Repeat the argument with this face.

At each step, on every horizontal line between the rays, the right part of the new face's boundary lies strictly to the right of the right part of the preceding face's boundary. Thus no face is visited twice. Since G has finitely many faces, the process cannot continue indefinitely. It cannot end at a single edge along the right side of W , because that side contains a vertex




Figure 3.8: In (a), the starting face already has the required vertex v . In (b) and (c), several edges are crossed before reaching F . The last crossed edge is denoted by e . The vertex v is in the interior of W in (b) and on its right side in (c). The three closely spaced edges are less than $\delta = s/3$ apart within W . Their endpoints lie outside W and are not shown. In each figure, the tile lies on e_0 . Only the relevant graph edges are shown.

of G strictly between ρ^- and ρ^+ . The process therefore reaches a triangular face whose right part contains a vertex strictly between the rays. That face and vertex satisfy the conclusion of the lemma. See Figure 3.8(b) and (c). $\square$

Lemma 3.6. *Suppose that there exists a tile $\mathcal{B}$ on an edge of G such that $\mathcal{B} \subseteq W$ and $\text{ht}(\mathcal{B}) > 2$. Then there exists a tile $\mathcal{B}^+$ on an edge of G such that*

$$\mathcal{B}^+ \subseteq \mathcal{E}(\mathcal{B}) \cap W, \quad \text{ht}(\mathcal{B}^+) \geq \text{ht}(\mathcal{B})/8.$$

Proof. Each graph edge along the right side of W has length at most one, so none intersects both boundary rays of $\mathcal{E}(\mathcal{B})$. Lemma 3.5 gives a vertex v between the rays with one incident edge intersecting each ray. Choose a boundary ray ρ^* at vertical distance at least $\text{ht}(\mathcal{B})/2$ from v .

Among all graph edges incident to v that intersect ρ^* , choose vr with the rightmost intersection z . See Figure 3.9. At least one such edge exists by the choice of v . The observation preceding Lemma 3.5 implies that v and z lie in W . By convexity,

$$vz \subseteq \mathcal{E}(\mathcal{B}) \cap W, \quad |z_y - v_y| \geq \text{ht}(\mathcal{B})/2 > 1.$$

The edge vr cannot lie on the left side of W because $v_x > 0$, or on its right side because $|z_y - v_y| > 1$. Its portion between v and z therefore contains a point in the interior of W . Thus vr is not incident to $F_\infty(G)$. Let u be the third vertex of the triangular face incident to vr and lying to the right of the line through vr .


Suppose that u lies on the opposite side of the line $y = z_y$ from v . The horizontal cross-section of the closed triangle $vr u$ on this line then starts at z and ends on vu , strictly to the

right of z . That endpoint belongs to ρ^* , contradicting the choice of vr . Thus $u_y \leq z_y$ if the chosen ray is ρ^+ , and $u_y \geq z_y$ if it is ρ^- .

The corresponding case of Lemma 3.4 now gives a tile $\mathcal{B}^+$ on vr whose left side is contained in vz and whose height satisfies

$$\text{ht}(\mathcal{B}^+) = \frac{|z_y - v_y|}{4} \geq \frac{\text{ht}(\mathcal{B})}{8}.$$

It remains to check containment. Each point of $\mathcal{B}^+$ is a rightward horizontal translate of a point of vz , so it lies in $\mathcal{E}(\mathcal{B})$. The tile is also contained in the closure of a triangular face. The observation preceding Lemma 3.5 therefore implies that $\mathcal{B}^+ \subseteq W$. $\square$



$$|z_y - v_y| \geq \frac{\text{ht}(\mathcal{B})}{2}, \quad \text{ht}(\mathcal{B}^+) = \frac{|z_y - v_y|}{4} \geq \frac{\text{ht}(\mathcal{B})}{8}.$$

Figure 3.9: The orange face F is supplied by Lemma 3.5. The edge vr has the rightmost intersection z on the chosen ray. In (a), that ray is ρ^+ and r lies above it. In (b), it is ρ^- and r lies below it. The other endpoints of the additional edges incident to v are omitted. The third vertex u may lie outside the region between the rays, as shown here. Filled dots are graph vertices, and hollow circles mark intersections of graph edges with the rays. Only the relevant graph edges are shown.

3.2.7 The first tile

Recall from Subsection 3.2.4 that W has height m and its left side is contained in a single edge of G along VB . We now construct the first tile in W .

Lemma 3.7. *There exists a tile $\mathcal{B}_1$ on an edge of G such that*

$$\mathcal{B}_1 \subseteq W, \quad \text{ht}(\mathcal{B}_1) \geq m/8.$$

Proof. Among the graph edges intersecting both horizontal sides of W , choose the rightmost one, e . At least one exists because the left side of W is contained in a single graph edge. Such edges have a consistent left-to-right order because G is crossing-free. The edge e cannot lie on the right side of W , since $m > n - 2 \geq 6$ and the edges there have length at most one. Thus e has a triangular face F on its right.

The right part of ∂F between the lines $y = a$ and $y = b$ lies to the right of e and has $x \leq 1$, so it lies in W . It cannot be contained in a single edge throughout this interval. Such an edge would intersect both horizontal sides farther right than e . It therefore consists of portions of two edges intersecting at a vertex v strictly between the two horizontal lines, with $v_x > 0$.

Choose a horizontal side of W at vertical distance at least $m/2$ from v . Among the graph edges incident to v intersecting that side, choose vr with the rightmost intersection z . Then $vz \subseteq W$ and $|z_y - v_y| \geq m/2 > 1$. As in Lemma 3.6, this edge has a triangular face lying to the right of the line through vr . Let u be its third vertex. If u were on the opposite side of $y = z_y$ from v , the cross-section of the closed triangle $vr u$ on this line would end on vu strictly to the right of z , as in that lemma. This endpoint has $x \leq 1$ because it lies on an edge of G , so it belongs to the chosen side of W , contradicting the choice of vr . Thus $u_y \leq z_y$ if the chosen side is the upper one, and $u_y \geq z_y$ if it is the lower one.

The corresponding case of Lemma 3.4 gives a tile on vr of height $|z_y - v_y|/4 \geq m/8$. Its points have y -coordinates between v_y and z_y . Each point lies on or to the right of vz and is in the closure of that triangular face. Consequently its points satisfy $0 \leq x \leq 1$ and $a \leq y \leq b$, so the tile is contained in the fixed rectangle W . $\square$

3.2.8 Counting the tiles

Proof of Proposition 3.1. Start with the tile $\mathcal{B}_1$ of Lemma 3.7, whose height is at least $m/8$. While $\text{ht}(\mathcal{B}_i) > 2$, apply Lemma 3.6. Lemma 3.6 ensures that $\mathcal{B}_{i+1} \subseteq \mathcal{E}(\mathcal{B}_i)$. Recall from the definition of the extension that every point of $\mathcal{E}(\mathcal{B}_i)$ has x -coordinate at least $x_{\min}(\mathcal{B}_i) + \delta$. Thus every point of $\mathcal{B}_{i+1}$ satisfies the same bound. Taking the minimum over its points gives

$$x_{\min}(\mathcal{B}_{i+1}) \geq x_{\min}(\mathcal{B}_i) + \delta.$$

The projections onto the x -axis of the tiles' leftmost horizontal sides are intervals of length δ contained in $[0, 1]$. By the preceding inequality, their interiors are pairwise disjoint. Thus, if k tiles have been constructed,

$$k\delta \leq 1, \quad k \leq 3/s.$$

The construction terminates at a tile $\mathcal{B}_k$ with $\text{ht}(\mathcal{B}_k) \leq 2$. Otherwise Lemma 3.6 would produce arbitrarily many tiles, contradicting this bound. At each of the $k - 1$ steps, the new

tile has height at least one eighth that of the preceding tile. Thus

$$\frac{m}{8^k} \leq \text{ht}(\mathcal{B}_k) \leq 2, \quad n - 2 < m \leq 2 \cdot 8^k \leq 2 \cdot 8^{3/s}. \quad \square$$

Figure 3.10 illustrates the sequence of tiles.




Figure 3.10: Four successive tiles in W . The bold intervals on $y = a$ are the horizontal projections of the tiles' leftmost horizontal sides. Each has length δ , and their interiors are pairwise disjoint. Graph vertices and edges are omitted.

3.3 The number of Steiner vertices

We first derive a lower bound on the number of Steiner vertices needed for small degradation on Q_n . We then prove a stronger worst-case bound for general polygons.

Proposition 3.2. *Let $n \geq 8$ be an integer, and let G be a proper triangulation of Q_n with degradation $d = \text{de}_{Q_n}(G)$. Then*

$$|V(G) \setminus V(Q_n)| \geq \frac{n}{10d} - 4.$$

Proof. Let $s = \text{mfs}(G)$. Since triangulation cannot increase minimum feature size, Lemma 3.1 gives $0 < s \leq \text{mfs}(Q_n) < 1$ and $s = \text{mfs}(Q_n)/d \geq 1/(2d)$.

Let $Q = \{q_1, \dots, q_n\}$. Every $q \in Q$ has a neighbour not on the segment DC . Indeed, consider an edge of G along DC incident to q . Since G is proper, the triangular face incident to that edge has a third vertex not on the segment DC adjacent to q .

Consider an arbitrary vertex v of G not on the segment DC , and let ℓ be the number of neighbours of v in Q . We show that $\ell \leq 4/s + 1$. Since $0 < s < 1$, the claim is immediate

if $\ell \leq 2$. Otherwise, choose a neighbour $q \in Q$ maximizing $L = |v_y - q_y|$. The y -coordinates of distinct neighbours differ by at least one, so

$$L \geq \frac{\ell - 1}{2} \geq 1.$$

The vertex q has the smallest or the largest y -coordinate among the neighbours of v in Q . Let r be the next vertex of Q from q towards the other neighbours in Q . Then $|qr| = 1$. By the choice of q and the inequality $L \geq 1$, the y -coordinate of r lies between q_y and v_y , possibly equal to v_y . The horizontal line through r intersects qv at a point z . Since $|q_y - r_y| = 1$ and $|q_y - v_y| = L$, similar triangles give $|rz| = (1 - v_x)/L$. See Figure 3.11. Hence

$$s \leq \text{dist}(r, qv) \leq |rz| = \frac{1 - v_x}{L} \leq \frac{2}{L}.$$

Here r is not incident to qv , and $v_x \geq -1$ because $G \subseteq \Omega(Q_n)$. Hence $L \leq 2/s$, which gives $\ell \leq 4/s + 1$.




Figure 3.11: The similar triangles in the proof of Proposition 3.2. The small triangle qrz is similar to the larger shaded right triangle, so $|rz| = (1 - v_x)/L$. Only part of G is shown, and the horizontal scale is enlarged.

Let $k = |V(G) \setminus V(Q_n)|$ be the number of Steiner vertices. The only input vertices not on the segment DC are A, V, B, E . Thus at most $k + 4$ vertices not on the segment DC can be adjacent to a vertex of Q . Since each vertex in Q has at least one neighbour not on the segment DC , and each vertex of G not on that segment has at most $4/s + 1$ neighbours in Q ,

$$n \leq (k + 4) \left(\frac{4}{s} + 1 \right) \leq \frac{5(k + 4)}{s},$$

so

$$k \geq \frac{ns}{5} - 4 \geq \frac{n}{10d} - 4. \quad \square$$

In particular, for every fixed $C > 0$, a proper triangulation of Q_n with degradation at most $C \log_2 n$ requires $\Omega(n/\log n)$ Steiner vertices. Indeed, Proposition 3.2 gives the lower bound $n/(10C \log_2 n) - 4$.

For general polygons, the worst-case lower bound on the number of Steiner vertices can be strengthened to $\Omega(N)$. We use the rectangle example of Bern, Dobkin, and Eppstein [4, Section 2], who attribute it to Mitchell. For $M \geq 2$, let R_M be the polygon whose corners are

$$(0, 0), \quad (1, 0), \quad (1, 2M), \quad (0, 2M),$$

with two additional input vertices $p = (1, M)$ and $r = (1, M + 1)$ on its right side. Thus R_M has six vertices and $\text{mfs}(R_M) = 1$. Bern et al. observed that the degradation of every classical triangulation of this rectangle grows without bound as M increases. The following lemma gives the bound and extends the argument to a rectangle attached along its bottom side to a larger polygon.

Lemma 3.8. *Let P be a polygon with $T = \Omega(R_M) \subseteq \Omega(P)$. Suppose that the top and vertical sides of T lie on P and $V(P) \cap T = V(R_M)$. If a proper triangulation G of P has no Steiner vertex in T , then*

$$\text{mfs}(G) \leq \frac{1}{M}.$$

In particular, every classical triangulation of R_M has degradation at least M .

Proof. The edge pr is not subdivided. Let prv be its incident triangular face, with v strictly to the left of pr . If $v \in T$, then v is either $(0, 0)$ or $(0, 2M)$, since there are no Steiner vertices in T . If $v \notin T$, the edge rv can leave T only through its bottom side, so $v_y \leq 0$.

If $v = (0, 2M)$, the horizontal line through r intersects pv at distance $1/M$ from r . Otherwise $v_y \leq 0$, and rv intersects the bottom side of T . Similar triangles then show that the horizontal line through p intersects rv at distance at most $1/(M + 1)$ from p . In each case, a vertex is at distance at most $1/M$ from a nonincident edge. The final statement follows by taking $P = R_M$ and using $\text{mfs}(R_M) = 1$. $\square$

We now combine $\Theta(N)$ copies of R_M and one copy of Q_n , where $n = \Theta(N)$, into a comb-shaped polygon, as illustrated in Figure 3.12. The construction gives the following lower bounds on degradation and the number of Steiner vertices.

Proposition 3.3. *There exist absolute constants $c_1, c_2 > 0$ and a family of polygons P_N , with exactly N vertices for each sufficiently large integer N , having the following properties.*

1. *Every proper triangulation G of P_N satisfies*

$$\text{de}_{P_N}(G) \geq c_1 \log_2 N.$$

2. For every constant $C > 0$ and all sufficiently large N , every proper triangulation G of P_N with $\text{de}_{P_N}(G) \leq C \log_2 N$ has at least $c_2 N$ Steiner vertices.

Steiner vertices may be added in the bounded face of P_N and on any of its edges.

Proof. The construction consists of a rectangular base, a connector to Q_n , and rectangular teeth. For a sufficiently large integer N , let

$$k = \left\lfloor \frac{N-11}{12} \right\rfloor, \quad n = N - 6k - 11, \quad h = n^2, \quad M = N^2.$$

Both n and k are $\Theta(N)$. Let $y_* = -h - 2$, and define the base S , the connector K , and the teeth T_j by

$$S = [-3k - 2, 1] \times [y_* - 1, y_*],$$

$$K = [0, 1] \times [y_*, -h],$$

$$T_j = [-3j - 1, -3j] \times [y_*, y_* + 2M], \quad 1 \leq j \leq k.$$

Let P_N be the polygon with region

$$\Omega(P_N) = \Omega(Q_n) \cup K \cup S \cup \bigcup_{j=1}^k T_j$$

and with the input vertices specified below. The connector K shares its top side BC with $\Omega(Q_n)$ and its bottom side with S . Each tooth T_j intersects S exactly along its bottom side. There are no other intersections between these regions. These attachments produce a simple polygon contained in $x \leq 1$. See Figure 3.12.




Figure 3.12: The combined polygon P_N . Each tooth contains a unit input edge, shown in red. The polygon Q_n is attached to the base through BC . Dashed segments indicate attachments. The drawing is schematic, and only some vertices on the right side of Q_n are shown.

Take as vertices of P_N all vertices of Q_n , the four corners of each T_j together with $(-3j, y_* + M)$ and $(-3j, y_* + M + 1)$, and the four corners of S together with $(0, y_*)$. Retain

collinear vertices. These disjoint groups give

$$|V(P_N)| = (n + 6) + 6k + 5 = N.$$

The boundary of each tooth, with its six input vertices, is a translated copy of R_M . The edges AV , VB , and the unit edges on DC remain unchanged.

All input vertices have integer coordinates, and AV is the only non-axis-parallel input edge. Added vertices satisfy $x \leq -3$ or $y \leq -h - 2$, so their distance from AV is at least two. By the argument of Lemma 3.1,

$$\mu := \text{mfs}(P_N) = \text{mfs}(Q_n) = \frac{h}{\sqrt{h^2 + 1}} \geq \frac{1}{2}.$$

Let G be a proper triangulation of P_N , and let $s = \text{mfs}(G)$ and $d = \text{de}_{P_N}(G) = \mu/s$. If $d \geq n$, the logarithmic lower bound follows immediately. Otherwise, the argument of Lemma 3.2 shows that VB has no subdivision vertex between $(0, 0)$ and $(0, -n)$. We can therefore choose W as in Subsection 3.2.4. Its left side lies on a single edge of G , and the graph edges along its right side have length at most one. Moreover, $\Omega(P_N) \subseteq \{(x, y) : x \leq 1\}$. These are the properties used in the tile construction and count. Applying that argument to the triangular faces of G , including faces that extend across the attachment, gives $n - 2 < 2 \cdot 8^{3/s}$. As in the proof of Theorem 2, it follows that $d \geq (\log_2 n)/36 = \Omega(\log N)$.

Finally, fix $C > 0$ and suppose that G has degradation at most $C \log_2 N$. For all sufficiently large N ,

$$s \geq \frac{1}{2C \log_2 N} > \frac{1}{M}.$$

By Lemma 3.8, applied after a translation, each tooth must contain a Steiner vertex. Since the teeth are disjoint, G has at least $k = \Omega(N)$ Steiner vertices. $\square$

Together with the $O(N)$ -Steiner-vertex upper bound of Aloupis et al. [3, Theorem 2], this shows that $\Theta(N)$ Steiner vertices are necessary and sufficient in the worst case for $O(\log N)$ degradation.

4. Conclusions

This thesis proves that logarithmic minimum-feature-size degradation is unavoidable in proper triangulations of some polygons. For every sufficiently large N , the family Q_n , with $N = n + 6$ vertices, contains a polygon for which every proper triangulation has degradation $\Omega(\log N)$. The lower bound holds even when arbitrarily many Steiner vertices may be added in the polygon's bounded face and on its edges.

Aloupis et al. [3, Theorem 2] previously proved that every N -vertex polygon admits a proper triangulation with $O(\log N)$ degradation using $O(N)$ Steiner vertices. Their construction leaves all input edges unsubdivided. Together with the lower bound proved here, this establishes worst-case optimal degradation $\Theta(\log N)$, whether or not Steiner vertices are allowed on input edges. The same order holds for triangulating the bounded faces of 2-connected crossing-free geometric graphs. The upper-bound construction extends to these inputs, and the lower bound already holds for polygons.

The linear upper bound on the number of Steiner vertices is also optimal in the worst case. Proposition 3.3 gives a family of N -vertex polygons for which logarithmic degradation is unavoidable and achieving $O(\log N)$ degradation requires $\Omega(N)$ Steiner vertices.

A direction for further research is to study whether different criteria for triangulation quality can be satisfied simultaneously. For example, can logarithmic feature-size degradation be achieved while keeping triangle angles bounded away from 0 and π , with angle bounds depending on the smallest interior angle of the input polygon? More concretely, Bishop [8, Corollary 1.4] proved that a polygon with smallest interior angle α admits a proper triangulation in which every triangle angle θ satisfies

$$\min\{\alpha, 36^\circ\} \leq \theta \leq 90^\circ - \frac{1}{2} \min\{\alpha, 36^\circ\}.$$

Does every N -vertex polygon admit a proper triangulation satisfying both these angle bounds and an $O(\log N)$ bound on feature-size degradation? More generally, one can ask what tradeoffs arise between preserving feature size, controlling triangle shape, and limiting the number of Steiner vertices.

Bibliography

- [1] G. Aloupis, E. D. Demaine, M. L. Demaine, V. Dujmović, and J. Iacono. *Minimum feature size preserving decompositions*. arXiv:0908.2493, 2009. <https://arxiv.org/abs/0908.2493>.
- [2] G. Aloupis, E. D. Demaine, M. L. Demaine, V. Dujmović, and J. Iacono. Meshes preserving minimum feature size. In *Abstracts from the 14th Spanish Meeting on Computational Geometry (EGC 2011)*, CRM Documents 8, pp. 205–208, 2011. https://erikdemaine.org/papers/FeatureSize_EGC2011/paper.pdf.
- [3] G. Aloupis, E. D. Demaine, M. L. Demaine, V. Dujmović, and J. Iacono. Meshes preserving minimum feature size. In *Computational Geometry: EGC 2011, Revised Selected Papers*, Lecture Notes in Computer Science 7579, pp. 258–273. Springer, 2012. https://doi.org/10.1007/978-3-642-34191-5_25.
- [4] M. Bern, D. Dobkin, and D. Eppstein. Triangulating polygons without large angles. *International Journal of Computational Geometry & Applications* 5(1–2):171–192, 1995. <https://doi.org/10.1142/S0218195995000106>.
- [5] M. Bern and D. Eppstein. Mesh generation and optimal triangulation. In *Computing in Euclidean Geometry*, second edition, Lecture Notes Series on Computing 4, pp. 47–123. World Scientific, 1995. https://doi.org/10.1142/9789812831699_0003.
- [6] M. Bern, J. R. Shewchuk, and N. Amenta. Triangulations and mesh generation. In J. E. Goodman, J. O’Rourke, and C. D. Tóth, editors, *Handbook of Discrete and Computational Geometry*, third edition, pp. 763–785. CRC Press, 2017. <https://people.eecs.berkeley.edu/~jrs/papers/trimesh.pdf>.
- [7] C. J. Bishop. Uniformly acute triangulations of polygons. *Discrete & Computational Geometry* 70(4):1571–1592, 2023. <https://doi.org/10.1007/s00454-023-00525-w>.
- [8] C. J. Bishop. Optimal triangulation of polygons. *Discrete & Computational Geometry* 76:1426–1473, 2026. <https://doi.org/10.1007/s00454-026-00831-z>.

- [9] J. Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *Journal of Algorithms* 18(3):548–585, 1995. <https://doi.org/10.1006/jagm.1995.1021>.
- [10] S. P. Sastry. A 2D advancing-front Delaunay mesh refinement algorithm. *Computational Geometry: Theory and Applications* 97:101772, 2021. <https://doi.org/10.1016/j.comgeo.2021.101772>.
- [11] J. R. Shewchuk. Delaunay refinement algorithms for triangular mesh generation. *Computational Geometry: Theory and Applications* 22(1–3):21–74, 2002. [https://doi.org/10.1016/S0925-7721\(01\)00047-5](https://doi.org/10.1016/S0925-7721(01)00047-5).