

Recognizing Weakly Simple Polygons*

Hugo A. Akitaya[†] Greg Aloupis[†] Jeff Erickson[‡] Csaba D. Tóth^{†§}

Abstract

We present an $O(n \log n)$ -time algorithm that determines whether a given n -gon in the plane is weakly simple. This improves upon an $O(n^2 \log n)$ -time algorithm by Chang, Erickson, and Xu [6]. Weakly simple polygons are required as input for several geometric algorithms. As such, recognizing simple or weakly simple polygons is a fundamental problem.

Keywords: simple polygon, combinatorial embedding, perturbation

MSC: 05C10, 05C38, 52C45, 68R10.

1 Introduction

A polygon is *simple* if it has distinct vertices and interior-disjoint edges that do not pass through vertices. Geometric algorithms are often designed for simple polygons, but many also work for degenerate polygons that do not “self-cross.” A polygon with at least three vertices is *weakly simple* if for every $\varepsilon > 0$, the vertices can be perturbed within a ball of radius ε to obtain a simple polygon. Such polygons arise naturally in numerous applications, e.g., for modeling planar networks or as the geodesic hull of points within a simple polygon (Figure 1).




Figure 1: (a) A simple polygon P with 16 vertices. (b) Eight points in the interior of P (solid dots); their geodesic hull is a weakly simple polygon P' with 14 vertices. (c) A perturbation of P' into a simple polygon.

*A preliminary version of this paper appeared in the *Proceedings of the 32nd International Symposium on Computational Geometry (SoCG 2016)*, doi:10.4230/LIPIcs.SoCG.2016.8.

[†]Department of Computer Science, Tufts University, Medford, MA. Email: hugo.alves.akitaya@tufts.edu, aloupis.greg@gmail.com, cdtoth@eecs.tufts.edu

[‡]Department of Computer Science, University of Illinois, Urbana-Champaign, IL. Email: jeffe@illinois.edu

[§]Department of Mathematics, California State University Northridge, Los Angeles, CA.

17 Several alternative definitions have been proposed for weakly simple polygons, formalizing the
 18 intuition that such polygons do not self-cross. Some of these definitions were unnecessarily restric-
 19 tive or incorrect; see [6] for a detailed discussion and five equivalent definitions for weak simplicity
 20 of a polygon. Among others, a result by Ribó Mor [16, Theorem 3.1] implies an equivalent defini-
 21 tion in terms of Fréchet distance, in which a polygon is perturbed into a simple closed curve (see
 22 Section 2). This definition is particularly useful for recognizing weakly simple polygons, since it
 23 allows transforming edges into polylines (by subdividing the edges with Steiner points, which may
 24 be perturbed). With suitable Steiner points, the perturbation of a vertex incurs only local changes.
 25 (In other words, we do not need to worry about stretchability of the perturbed configuration.)

26 We can decide whether an n -gon in the plane is simple in $O(n \log n)$ time by a sweepline
 27 algorithm [17]. Chazelle’s polygon triangulation algorithm also recognizes simple polygons (in
 28 $O(n)$ time), because it only produces a triangulation if the input is simple [7]. Recognizing weakly
 29 simple polygons, however, is more subtle. Skopenkov [18] gave a combinatorial characterization of
 30 the topological obstructions to weak simplicity in terms of line graphs. Cortese et al. [10] gave an
 31 $O(n^6)$ -time algorithm to recognize weakly simple n -gons. Chang et al. [6] improved the running
 32 time to $O(n^2 \log n)$ in general; and to $O(n \log n)$ in several special cases. They identified two
 33 features that are difficult to handle: A *spur* is a vertex whose incident edges overlap, and a *fork*
 34 is a vertex that lies in the interior of an edge. (A vertex may be both a fork and a spur.) They
 35 gave an easy algorithm for polygons that have neither forks nor spurs, and two more involved ones
 36 for polygons with spurs but no forks and for polygons with forks but no spurs, all three running
 37 in $O(n \log n)$ time. In the presence of both forks and spurs, they presented an $O(n^2 \log n)$ time
 38 algorithm that eliminates forks by subdividing all edges that contain vertices in their interiors,
 39 potentially creating a quadratic number of vertices.

40 We show how to manage both forks and spurs efficiently, while building on ideas from [6, 10]
 41 and from Arkin et al. [2], and obtain the following main results.

42 **Theorem 1.** *Deciding whether a polygon P with n vertices in the plane is weakly simple takes*
 43 *$O(n \log n)$ time.*

44 **Theorem 2.** *Given a weakly simple polygon P with n vertices and a constant $\varepsilon > 0$, a simple*
 45 *polygon with $2n$ vertices within Fréchet distance ε from P can be computed in $O(n \log n)$ time.*

46 Our decision algorithm is detailed in Sections 3–5. It consists of three phases, simplifying
 47 the input polygon by a sequence of reduction steps. First, the *preprocessing* phase rules out
 48 edge crossings in $O(n \log n)$ time and applies known reduction steps such as *crimp reductions* and
 49 *node expansions* (Section 3). Second, the *bar simplification* phase successively eliminates all forks
 50 (Section 4). Third, the *spur elimination* phase eliminates all spurs (Section 5). When neither forks
 51 nor spurs are present, we can decide weak simplicity in $O(n)$ time [10]. Finally, by reversing the
 52 sequence of operations, we can also perturb any weakly simple polygon into a simple polygon in
 53 $O(n \log n)$ time (Section 6).

54 2 Preliminaries

55 In this section, we review previously established definitions and known methods from [6] and [10].

56 **Polygons and weak simplicity.** An *arc* in \mathbb{R}^2 is a continuous function $\gamma : [0, 1] \rightarrow \mathbb{R}^2$. A *closed*
 57 *curve* is a continuous function (map) $\gamma : \mathbb{S}^1 \rightarrow \mathbb{R}^2$. A closed curve γ is *simple* (also known as

58 a *Jordan curve*) if it is injective. A (*simple*) *polygon* is the image of a piecewise linear (*simple*)
59 closed curve. Thus a polygon P can be represented by a cyclic sequence of points (p_0, \dots, p_{n-1}) ,
60 called *vertices*, where the image of γ consists of line segments $p_0p_1, \dots, p_{n-2}p_{n-1}$, and $p_{n-1}p_0$ in this
61 cyclic order. Note that a nonsimple polygon may have repeated vertices and overlapping edges [14].
62 Similarly, a *polygonal chain* (alternatively, *path*) is the image of a piecewise linear arc, and can be
63 represented by a sequence of points $[p_0, \dots, p_{n-1}]$.

64 A polygon $P = (p_0, \dots, p_{n-1})$ is *weakly simple* if $n = 2$, or if $n > 2$ and for every $\varepsilon > 0$ there is a
65 simple polygon (p'_0, \dots, p'_{n-1}) such that $|p_i, p'_i| < \varepsilon$ for all $i = 0, \dots, n-1$. This definition is difficult
66 to work with because a small perturbation of a vertex modifies the two incident edges, which may be
67 long, and the effect of a perturbation is not localized. Combining earlier results from [9], [10], and
68 [16, Theorem 3.1], an equivalent definition was formulated by Chang et al. [6] in terms of Fréchet
69 distance: A polygon given by $\gamma : \mathbb{S}^1 \rightarrow \mathbb{R}^2$ is weakly simple if for every $\varepsilon > 0$ there is a simple
70 closed curve $\gamma' : \mathbb{S}^1 \rightarrow \mathbb{R}^2$ such that $\text{dist}_F(\gamma, \gamma') < \varepsilon$, where dist_F denotes the Fréchet distance
71 between two closed curves. The curve γ' can approximate an edge of the polygon by a polyline,
72 and any perturbation of a vertex can be restricted to a small neighborhood. With this definition,
73 recognizing weakly simple polygons becomes a combinatorial problem, as explained below. Note
74 that in topology, the broader question of *isotopic embeddability* has been considered [15, 18]: Given
75 a continuous map $f : A \rightarrow \mathbb{R}^d$ for a simplicial complex A , is it isotopic to some *injective* continuous
76 map (i.e., *embedding*) $g : A \rightarrow \mathbb{R}^d$?

77 **Bar decomposition and image graph.** Two edges of a polygon P *cross* if their interiors intersect
78 at precisely one point; we call this an *edge crossing*. Weakly simple polygons cannot have edge
79 crossings. In the remainder of this section, we assume that such crossings have been ruled out. Two
80 edges of P *overlap* if their intersection is a (nondegenerate) line segment. The transitive closure of
81 the overlap relation is an equivalence relation on the edges of P ; see Figure 2(a) where equivalence
82 classes are represented by purple regions. The union of all edges in an equivalence class is called a
83 *bar*.¹ All bars of a polygon can be computed in $O(n \log n)$ time [6]. The bars are open line segments
84 that are pairwise disjoint. There are at most n bars, since the bars are unions of disjoint subsets
85 of edges.

86 The vertices and bars of P define a planar straight-line graph G , called the *image graph* of P .
87 We call the vertices and edges of G *nodes* and *segments*¹ to distinguish them from the vertices and
88 edges of P . Every node that is not in the interior of a bar is called *sober*¹. The set of nodes in G
89 is $\{p_0, \dots, p_{n-1}\}$ (note that P may have repeated vertices that correspond to the same node); two
90 nodes are connected by a segment in G if they are consecutive nodes along a bar; see Figure 2(b).
91 Hence G has $O(n)$ nodes and segments, and it can be computed in $O(n \log n)$ time [6]. Note,
92 however, that up to $O(n)$ edges of P may pass through a node of G , and there may be $O(n^2)$
93 edge-node pairs such that an edge of P passes through a node of G . An $O(n \log n)$ -time algorithm
94 cannot afford to compute these pairs explicitly.

95 **Operations.** We use certain elementary operations that successively modify a polygon and ul-
96 timately eliminate forks and spurs. An operation that produces a weakly simple polygon if and
97 only if it is performed on a weakly simple polygon is called *ws-equivalent*. Several such operations
98 are already known (e.g., crimp reduction, node expansion, bar expansion). We shall use these and
99 introduce several new operations in Sections 3.3–5.

100 **Combinatorial characterization of weak simplicity.** To show that an operation is ws-

¹We adopt terminology from [6].




Figure 2: (a) The bar decomposition for a weakly simple polygon P with 16 vertices (P is perturbed into a simple polygon for clarity). (b) The image graph of P . (c) A perturbation in a strip system of P .

101 equivalent, it suffices to provide suitable simple ε -perturbations for all $\varepsilon > 0$. We use a combi-
 102 natorial representation of an ε -perturbation (independent of ε or any specific embedding). When a
 103 weakly simple polygon P is perturbed into a simple polygon, overlapping edges in P are perturbed
 104 into interior-disjoint near-parallel edges, which define an ordering. It turns out that these orderings
 105 over all segments of the image graph are sufficient to encode an ε -perturbation and to (re)construct
 106 an ε -perturbation.

107 We rely on the notion of “strip system” introduced in [6, Appendix B]. Similar concepts have
 108 previously been used in [9, 10, 11, 15, 18]. Let P be a polygon and G its image graph. Without
 109 loss of generality, we assume that no bar is vertical (so that the above-below relationship is defined
 110 between disjoint segments parallel to a bar). For every $\varepsilon > 0$, the ε -strip-system of P consists of
 111 the following regions:

- 112 • For every node u of G , let D_u be a disk of radius ε centered at u .
- 113 • For every segment uv , let the *corridor* N_{uv} be the set of points at distance at most ε^2 from
 114 uv , outside of the disks D_u and D_v , that is, $N_{uv} = \{p \in \mathbb{R}^2 : \text{dist}(p, uv) \leq \varepsilon^2, p \notin D_u \cup D_v\}$.

115 Denote by U_ε the union of all these disks and corridors. There is a sufficiently small $\varepsilon_0 = \varepsilon_0(P) > 0$,
 116 depending on P , such that the disks D_u are pairwise disjoint, the corridors N_{uv} are pairwise disjoint,
 117 and every corridor N_{uv} of a segment intersects only the disks at its endpoints D_u and D_v . These
 118 properties hold for all ε , $0 < \varepsilon < \varepsilon_0$.

119 A polygon is *in the ε -strip-system* of P if its edges alternate between an edge that connects the
 120 boundaries of two disks D_u and D_v and whose interior is contained in N_{uv} ; and an edge between
 121 two points on the boundary of a disk. In particular, the edges of P that lie in a disk D_u or a
 122 corridor N_{uv} form a perfect matching. See Figure 2(c) for an example, where the edges within
 123 the disk D_u are drawn with circular arcs for clarity. Let $\Phi(P)$ be the set of simple polygons in
 124 the ε -strip-system of P that cross the disks and corridors in the same order as P traverses the
 125 corresponding nodes and segments of G . It is clear that every $Q \in \Phi(P)$ is within Fréchet distance
 126 ε from P . By [6, Theorem B.2], P is weakly simple if and only if $\Phi(P) \neq \emptyset$.

127 **Combinatorial representation by signatures.** Let Q be a polygon in the strip system of P .
 128 For each segment uv , the above-below relationship of the edges of Q in N_{uv} is a total order. We
 129 define the *signature* of $Q \in \Phi(P)$, denoted $\sigma(Q)$, as the collection of these total orders for all
 130 segments of G .

131 Given the signature $\sigma(Q)$ of a polygon Q in the strip system of P , we can easily (re)construct a
 132 simple polygon Q' with the same signature in the ε -strip-system of P for any $0 < \varepsilon < \varepsilon_0$. For every
 133 segment uv of G , let the *volume* $\text{vol}(uv)$ be the number of edges of P that lie on uv . Place $\text{vol}(uv)$
 134 parallel line segments between ∂D_u and ∂D_v in N_{uv} of the ε -strip-system of P . Finally, for every
 135 disk D_u , construct a straight-line perfect matching between the endpoints of these edges that lie
 136 in ∂D_u : connect the endpoints of two edges if they correspond to adjacent edges of P . It is easily
 137 verified that the Fréchet distance between Q and Q' is at most 2ε . Furthermore, $Q \in \Phi(P)$ implies
 138 $Q' \in \Phi(P)$, since Q and Q' determine the same perfect matching between corresponding endpoints
 139 on ∂D_u at every node u .

140 **Remark 1.** The construction above has two consequences: (1) To prove weak simplicity, it is
 141 enough to find a signature that defines a simple perturbation. In other words, the signature can
 142 witness weak simplicity (independent of the value of ε). (2) Weak simplicity of a polygon depends
 143 only on the *combinatorial embedding* of the image graph G (i.e., the counterclockwise order of
 144 edges incident to each vertex), as long as G is a planar graph. Consequently, when an operation
 145 modifies the image graph, it is enough to maintain the combinatorial embedding of G (the precise
 146 coordinates of the nodes do not matter).

147 In the presence of spurs, the size of a signature is $O(n^2)$, and this bound is the best possible.
 148 We use this simple combinatorial representation in our proofs of correctness, but our algorithm
 149 does not maintain it explicitly. In Section 6, we introduce another combinatorial representation of
 150 $O(n)$ size that uses the ordering of the edges in each bar (rather than each segment) of the image
 151 graph.

152 **Combinatorially different perturbations.** In the absence of spurs, a polygon P determines a
 153 unique noncrossing perfect matching in each disk D_u , hence a unique noncrossing 2-regular graph
 154 in the ε -strip-system of P [6, Section 3.3]. Consequently, to decide whether P is weakly simple it is
 155 enough to check whether this graph is connected. The uniqueness no longer holds in the presence of
 156 spurs. In fact, it is not difficult to construct weakly simple n -gons that admit $2^{\Theta(n)}$ perturbations
 157 into simple polygons that are combinatorially different (i.e., have different bar-signatures); see
 158 Figure 3.



Figure 3: Two perturbations of a weakly simple polygon on 6 vertices (all of them spurs) that alternate between two distinct points in the plane.

159 3 Preprocessing

160 We are given a polygon $P = (p_0, \dots, p_{n-1})$ in the plane. By a standard line sweep [17], we can test
 161 whether any two edges properly cross; if they do, the algorithm halts and reports that P is not
 162 weakly simple. We then simplify the polygon, using some known steps from [2, 6], and some new
 163 ones. All of this takes $O(n \log n)$ time.

164 **3.1 Crimp reduction**

165 Arkin et al. [2] gave an $O(n)$ -time algorithm for recognizing weakly simple n -gons in the special case
 166 where all edges are collinear (in the context of flat foldability of a polygonal linkage). They defined
 167 the ws-equivalent crimp-reduction operation. A *crimp* is a chain of three consecutive collinear edges
 168 $[a, b, c, d]$ such that both the first edge $[a, b]$ and the last edge $[c, d]$ contain the middle edge $[b, c]$
 169 (the containment need not be proper). The operation $\text{crimp-reduction}(a, b, c, d)$ replaces the crimp
 170 $[a, b, c, d]$ with edge $[a, d]$; see Figure 4.




Figure 4: A crimp reduction replaces $[a, b, c, d]$ with $[a, d]$. Top: image graph. Bottom: polygon.

171 **Lemma 1.** *The crimp-reduction operation is ws-equivalent.*

172 *Proof.* Let P_1 and P_2 be two polygons such that P_2 is obtained from P_1 by the operation crimp-
 173 $\text{reduction}(a, b, c, d)$. Without loss of generality, assume that ad is horizontal with a on the left and
 174 d on the right.

175 First assume that P_1 is weakly simple. Then there exists a simple polygon $Q_1 \in \Phi(P_1)$. We
 176 modify Q_1 to obtain a simple polygon $Q_2 \in \Phi(P_2)$. Without loss of generality, assume that edge
 177 $[a, b]$ is above $[b, c]$ (consequently, $[c, d]$ is below $[b, c]$) in Q_1 . The modification involves the perfect
 178 matchings at the disks D_b and D_c , and all disks and corridors along the line segment bc . Denote by
 179 W_{top} the set of maximal paths that lie in the convex hull of $D_b \cup D_c$, below $[a, b]$ and above $[b, c]$;
 180 similarly, let W_{bot} be the set of maximal paths that lie in the convex hull of $D_b \cup D_c$, below $[b, c]$ and
 181 above $[c, d]$. We proceed in two steps; refer to Figure 5. First, replace the path $[a, b, c, d]$ with the
 182 path $[a, c, b, d]$ such that the new edge $[a, c]$ replaces the old $[a, b]$ in the edge ordering of segment
 183 ac , the new $[c, b]$ replaces $[b, c]$ in the segments contained in bc , and finally the new $[b, d]$ replaces
 184 $[c, d]$ in bd . Second, exchange W_{top} and W_{bot} such that the top-to-bottom order within each set of
 185 paths remains the same. Since the top-to-bottom order within W_{top} and W_{bot} is preserved, and the
 186 paths in W_{top} (resp., W_{bot}) lie below (resp., above) the new path $[a, c, b, d]$, no edge crossings have
 187 been introduced. We obtain a simple polygon $Q_2 \in \Phi(P_2)$, which shows that P_2 is weakly simple.




Figure 5: The operation crimp-reduction replaces a crimp $[a, b, c, d]$ with an edge $[ad]$.




Figure 6: The reversal of crimp-reduction replaces edge $[ad]$ with a crimp $[a, b, c, d]$.

188 Next assume that P_2 is weakly simple. Then, there exists a simple polygon $Q_2 \in \Phi(P_2)$. We
 189 modify Q_2 to obtain a simple polygon $Q_1 \in \Phi(P_1)$; refer to Figure 6. Replace edge $[a, d]$ by $[a, b, c, d]$

190 also replacing $[a, d]$ in the ordering of the affected segments by $[c, d]$, $[b, c]$, and $[a, b]$, in this order.
 191 The new ordering produces a polygon Q_1 in the strip system of P . Because Q_2 is simple, by
 192 construction the new matchings do not interact with the preexisting edges in the disks. Hence,
 193 $Q_1 \in \Phi(P_1)$, which shows that P_1 is weakly simple. \square

194 Given a chain of two edges $[a, b, c]$ such that $[a, b]$ and $[b, c]$ are collinear but do not overlap, the
 195 merge operation replaces $[a, b, c]$ with a single edge $[a, c]$. The merge operation (as well as its inverse,
 196 subdivision) is ws-equivalent by the definition of weak simplicity in terms of Fréchet distance [6]. If
 197 we greedily apply crimp-reduction and merge operations, in linear time we obtain a polygon with
 198 the following two properties:

- 199 (A1) Every two consecutive collinear edges overlap (i.e., form a spur).
- 200 (A2) No three consecutive collinear edges form a crimp.

201 Assuming properties (A1) and (A2), we can characterize a chain of collinear edges with the
 202 sequence of their edge lengths.

203 **Lemma 2.** *Let $C = [e_i, \dots, e_k]$ be a chain of collinear edges in a polygon with properties (A1)*
 204 *and (A2). Then the sequence of edge lengths $(|e_i|, \dots, |e_k|)$ is unimodal (all local maxima are*
 205 *consecutive); and no two consecutive edges have the same length, except possibly the maximal edge*
 206 *length that can occur at most twice.*

207 *Proof.* For every j such that $i < j < k$, consider $|e_j|$. If $|e_{j-1}|$ and $|e_{j+1}|$ are at least as large
 208 as $|e_j|$, then the three edges form a crimp, by (A1). However, this contradicts (A2). This proves
 209 unimodality, and that no three consecutive edges can have the same length. In fact if $|e_j|$ is not
 210 maximal, one neighbor must be strictly smaller, to avoid the same contradiction. \square

211 The operations introduced in Section 4 maintain properties (A1)–(A2) for all maximal paths
 212 inside an elliptical disk D_b .

213 3.2 Node expansion

214 Compute the bar decomposition of P and its image graph G (defined in Section 2, see Figure 2).
 215 For every sober node of the image graph, we perform the ws-equivalent node-expansion operation,
 216 described in [6, Section 3] (Cortese et al. [10] call this a *cluster expansion*). Let u be a sober node
 217 of the image graph. Let D_u be the disk centered at u with radius $\delta > 0$ sufficiently small so that D_u
 218 intersects only the segments incident to u . For each segment ux incident to u , create a new node
 219 u^x at the intersection point $ux \cap \partial D_u$. Then modify P by replacing each subpath $[x, u, y]$ passing
 220 through u by $[x, u^x, u^y, y]$; see Figure 7. If a node expansion produces an edge crossing, report that
 221 P is not weakly simple.

222 3.3 Bar expansion

223 Chang et al. [6, Section 4] define a bar expansion operation. In this paper, we refer to it as **old-bar-**
 224 **expansion**. For a bar b of the image graph, draw a long and narrow ellipse D_b around the interior
 225 nodes of b , create subdivision vertices at the intersection of ∂D_b with the edges, and replace each
 226 maximal path in D_b by a straight-line edge. If b contains no spurs, old-bar-expansion is known to




Figure 7: Node expansion. (Left) Changes in the image graph. (Right) Changes in P (the vertices are perturbed for clarity). New nodes are shown as squares.




Figure 8: The old-bar-expansion converts a non-weakly simple polygon to a weakly simple one.

227 be ws-equivalent [6]. Otherwise, it can produce false positives, hence it is not ws-equivalent; see
 228 Figure 8 for an example.

229 **New bar expansion operation.** Let b be a bar in the image graph with at least one interior node;
 230 see Figure 9. Without loss of generality, assume that b is horizontal. Let D_b be an ellipse whose
 231 major axis is in b such that D_b contains all interior nodes of b (nodes in b except its endpoints),
 232 but does not contain any other node of the image graph and does not intersect any segment that
 233 is not incident to some node inside D_b .

234 Similar to old-bar-expansion, the operation new-bar-expansion introduces subdivision vertices on
 235 ∂D_b , however we keep all interior vertices of a bar at their original positions. In Section 4, we apply
 236 a sequence of new operations to eliminate all vertices on b sequentially while creating new nodes in
 237 the vicinity of D_b . Our bar expansion operation can be considered as a preprocessing step for this
 238 subroutine.

239 For each segment ux between a node $u \in b \cap D_b$ and a node $x \notin b$, create a new node u^x at the
 240 intersection point $ux \cap \partial D_b$ and subdivide every edge $[u, x]$ to a path $[u, u^x, x]$. For each endpoint
 241 v of b , create two new nodes, v' and v'' , as follows. Node v is adjacent to a unique segment $vw \subset b$,
 242 where $w \in b \cap D_b$. Create a new node $v' \in \partial D_b$ sufficiently close to the intersection point $vw \cap \partial D_b$,
 243 but strictly above b ; and create a new node v'' in the interior of segment $vw \cap D_b$. Subdivide every
 244 edge $[v, y]$, where $y \in b$, into a path $[v, v', v'', y]$. Since the new-bar-expansion operation consists of
 245 only subdivisions (and slight perturbations of the edges passing through the end-segments of the
 246 bars), it is ws-equivalent.




Figure 9: The changes in the image graph caused by new-bar-expansion.

247 **Crossing paths.** Apart from node-expansion and old-bar-expansion, none of our operations creates
 248 edge crossings. In some cases, our bar simplification algorithm (Section 4) detects whether two
 249 subpaths cross. Crossings between overlapping paths are not easy to identify (see [6, Section 2] for
 250 a discussion). We rely on the following simple condition to detect some (but not all) crossings.

251 **Lemma 3.** *Let P be a weakly simple polygon parameterized by a curve $\gamma_1 : \mathbb{S}^1 \rightarrow \mathbb{R}^2$; and let
 252 $\gamma_2 : \mathbb{S}^1 \rightarrow \mathbb{R}^2$ be a closed Jordan curve that does not pass through any vertices of P and intersects
 253 every edge of P transversely. Suppose that q_1, \dots, q_4 are distinct points in $\gamma_2(\mathbb{S}^1)$ in counterclockwise
 254 order. Then there are no two disjoint arcs $I_1, I_2 \subset \mathbb{S}^1$ such that $\gamma_1(I_1)$ and $\gamma_1(I_2)$ connect q_1 to q_3
 255 and q_2 to q_4 , each passing through the interior of $\gamma_2(\mathbb{S}^1)$.*

256 *Proof.* Suppose, to the contrary, that there exist two disjoint arcs $I_1, I_2 \subset \mathbb{S}^1$ such that $\gamma_1(I_1)$
 257 and $\gamma_1(I_2)$ respectively connect q_1 to q_3 and q_2 to q_4 , passing through the interior of $\gamma_2(\mathbb{S}^1)$. (See
 258 Figure 10.) Since P is weakly simple, then γ_1 can be perturbed to a closed Jordan curve γ'_1 with
 259 the same properties as γ_1 . Let U denote the interior of $\gamma_2(\mathbb{S}^1)$, and note that U is simply connected.
 260 Consequently, $U \setminus \gamma'_1(I_1)$ has two components, which are incident to q_2 and q_4 , respectively. The
 261 Jordan arc $\gamma'_1(I_2)$ connects q_2 to q_4 via U , so it must intersect $\gamma'_1(I_1)$, contradicting the assumption
 262 that γ'_1 is a Jordan curve. \square




Figure 10: Forbidden configuration described by Lemma 3.

263 We show that a weakly simple polygon cannot contain certain configurations, outlined below.

264 **Corollary 1.** *A weakly simple polygon cannot contain a pair of paths of the following types:*

- 265 1. $[u_1, u_2, u_3]$ and $[v, u_2, w]$, where u_2u_1 , u_2v , u_2u_3 , and u_2w are nonoverlapping segments in
 266 this cyclic order around u_2 (node crossing; see Figure 11(a)).
- 267 2. $[u_1, u_3, w]$ and $[v, u_2, u_4]$, where u_1 , u_2 , u_3 , and u_4 are on a line in this order, and nodes v
 268 and w lie in an open halfplane bounded by this line (Figure 11(b)).
- 269 3. $[u_1, u_2, u_3]$ and $[v_1, v_2, \dots, v_{k-1}, v_k]$ where $v_2 \in \text{int}(u_2u_3)$, $v_3, \dots, v_{k-1} \in \{u_2\} \cup \text{int}(u_2u_3)$,
 270 nodes u_1 and v_1 lie in an open halfplane bounded by the supporting line of u_2u_3 , and node v_k
 271 lies on the other open halfplane bounded by this line (Figure 11(c)).

272 *Proof.* In all four cases, Lemma 3 with a suitable Jordan curve γ_2 completes the proof. In case 1,
 273 let γ_2 be a small circle around u_2 . In case 2, let γ_2 be a small neighborhood of segment u_1u_2 . In
 274 case 3, let γ_2 be a small neighborhood of the convex hull of $\{v_2, \dots, v_{k-1}\}$. \square

275 **Terminology.** We classify the maximal paths in D_b . All nodes $u \in \partial D_b$ lie either above or
 276 below b . We call them *top* and *bottom* nodes, respectively. Let \mathcal{P} denote the set of maximal paths
 277 $p = [u_1^x, u_1, \dots, u_k, u_k^y]$ in D_b . The paths in \mathcal{P} are classified based on the position of their endpoints.
 278 A path p can be labeled as follows:




Figure 11: Three pairs of incompatible paths.

- 279 • *cross-chain* if u_1^x and u_k^y are top and bottom nodes respectively,
- 280 • *top chain* (resp., *bottom chain*) if both u_1^x and u_k^y are top nodes (resp., bottom nodes),
- 281 • *pin* if $p = [u_1^x, u_1, u_1^x]$ (note that every pin is a top or a bottom chain),
- 282 • *V-chain* if $p = [u_1^x, u_1, u_1^y]$, where $x \neq y$ and p is a top or a bottom chain.

283 Finally, let $\mathcal{Pin} \subset \mathcal{P}$ be the set of pins, and $\mathcal{V} \subset \mathcal{P}$ the set of V-chains.

284 3.4 Clusters

285 As a preprocessing step for spur elimination (Section 5), we group all nodes that do not lie inside a
 286 bar into *clusters*. After **node-expansion** and **new-bar-expansion**, all such nodes lie on a boundary of
 287 a disk (circular or elliptical). For every sober node u , we create $\deg(u)$ clusters as follows. Refer to
 288 Figure 12. The node expansion has replaced u with new nodes on ∂D_u . Subdivide each segment
 289 in D_u with two new nodes. For each node $v \in \partial D_u$, form a cluster $C(v)$ that consists of v and
 290 all adjacent (subdivision) nodes inside D_u . For each node u on the boundary of an elliptical disk
 291 D_b , subdivide the unique edge outside D_b incident to u with a node u^* . Form a cluster $C(u^*)$
 292 containing u and u^* . Every cluster maintains the following invariants.

293 **Cluster Invariants.** For every cluster $C(u)$:

- 294 (I1) $C(u)$ induces a tree $T[u]$ in the image graph rooted at u .
- 295 (I2) Every maximal path of P in $C(u)$ is of one of the following two types:
 - 296 (a) both endpoints are at the root of $T[u]$ and the path contains a single spur;
 - 297 (b) one endpoint is at the root, the other is at a leaf, and the path contains no spurs.
- 298 (I3) Every leaf node ℓ satisfies one of the following conditions:
 - 299 (a) ℓ has degree one in the image graph of P (and every vertex at ℓ is a spur);
 - 300 (b) ℓ has degree two in the image graph of P and there is no spur at ℓ .
- 301 (I4) No edge passes through a leaf ℓ (i.e., there is no edge $[a, b]$ such that $\ell \in ab$ but $\ell \notin \{a, b\}$).

302 Initially, every cluster trivially satisfies (I1)–(I2) and every leaf node satisfies (I3)–(I4) since it
 303 was created by a subdivision.

304 **Dummy vertices.** Although the operations described in Sections 4 and 5 introduce new nodes in
 305 the clusters, the image graph will always have $O(n)$ nodes and segments. A vertex at a cluster node
 306 is called a *benchmark* if it is a spur or if it is at a leaf node; otherwise it is called a *dummy vertex*.
 307 Paths traversing clusters may jointly contain $\Theta(n^2)$ dummy vertices in the worst case, however we




Figure 12: Formation of new clusters around (left) a sober node and (right) a node on the boundary of an elliptical disk. The roots of the induced trees are colored blue.

308 do not store these explicitly. By (I1), (I2), and (I3) a maximal path in a cluster can be uniquely
 309 encoded by one benchmark vertex: if it goes from a root to a spur at an interior node s and back,
 310 we record only $[s]$; and if it traverses $T[u]$ from the root to a leaf ℓ , we record only $[\ell]$.

311 4 Bar simplification

312 In this section we introduce three new ws-equivalent operations and show that they can eliminate
 313 all vertices from each bar independently (thus eliminating all forks). The bar decomposition is
 314 pre-computed, and the bars remain fixed during this phase (even though all edges along each bar
 315 are eliminated).

316 We give an overview of the overall effect of the operations (Section 4.1), define them and show
 317 that they are ws-equivalent (Sections 4.2–4.3), and then show how to use these operations to
 318 eliminate all vertices from a bar (Section 4.4).

319 4.1 Overview

320 After preprocessing in Section 3, we may assume that P has no edge crossings and satisfies (A1)–
 321 (A2). We summarize the overall effect of the bar simplification subroutine for a given expanded
 322 bar.

323 **Changes in the image graph G .** Refer to Figure 13. All nodes in the interior of the ellipse
 324 D_b are eliminated. Some spurs on b are moved to new nodes in the clusters along ∂D_b . Segments
 325 inside D_b connect two leaves of trees induced by clusters.




Figure 13: The changes in the image graph caused by a bar simplification.

326 **Changes in the polygon P .** Refer to Figure 14. Consider a maximal path p in P that lies in
 327 D_b . The bar simplification replaces $p = [u, \dots, v]$ with a new path p' . By (I3)–(I4), only nodes
 328 u and v in p lie on ∂D_b . If p is the concatenation of a path p_1 and p_1^{-1} (the path formed by the
 329 vertices of p_1 in reverse order), then p' is a spur in the cluster containing u (Figure 14 (a)). If p
 330 has no such decomposition, but its two endpoints are at the same node, $u = v$, then p' is a single

331 edge connecting two leaves in the cluster containing u (Figure 14 (b)). If the endpoints of p are at
 332 two different nodes, p' is an edge between two leaves of the clusters containing u and v respectively
 333 (Figure 14 (c) and (d)).




Figure 14: The changes in the polygon caused by a bar simplification.

334 4.2 Primitives

335 The operations in Section 4.3 rely on two basic steps, **spur-reduction** and **node-split** (see Figure 15).
 336 Together with **merge** and **subdivision**, these operations are called *primitives*.

337 **spur-reduction**(u, v). Assume that every vertex at node u has at least one incident edge
 338 $[u, v]$. While there exists a path $[u, v, u]$, replace it with a single-vertex path $[u]$. (See
 339 Figure 15, left.)

340 **node-split**(u, v, w). Assume that segments uv and vw are consecutive in radial order
 341 around v , node v is not in the interior of any edge that contains uv or vw ; and P has
 342 no spurs of the form $[u, v, u]$ or $[w, v, w]$. Create node v^* in the interior of the wedge
 343 $\angle uvw$ sufficiently close to v ; replace every path $[u, v, w]$ with $[u, v^*, w]$. (See Figure 15,
 344 right.)




Figure 15: Left: Spur-reduction(u, v). Right: Node-split(u, v, w).

345 The following two lemmas are generalizations of the results in [6, Section 5].

346 **Lemma 4.** *Operation spur-reduction is ws-equivalent.*

347 *Proof.* Let P' be obtained from applying spur-reduction(u, v) to P . First suppose that P is weakly
 348 simple. Then, there exists a simple polygon $Q \in \Phi(P)$ represented by its signature. Successively
 349 replace any path $[u, v, u]$ by $[u]$ and delete these two edges from the ordering. The new signature
 350 defines a polygon Q' in the strip system of P' . By the assumption in the operation, every edge of
 351 Q in D_u is adjacent to an edge in N_{uv} , which has another endpoint in ∂D_v . Since Q is simple, the
 352 counterclockwise order of the endpoints of the deleted edges in ∂D_v is the same as the clockwise
 353 order of the endpoints of the new edges in ∂D_u . Thus, the new matching in D_u produces no
 354 crossings, $Q' \in \Phi(P')$, and P' is weakly simple.

355 Now suppose P' is weakly simple. Then, there exists a simple polygon $Q' \in \Phi(P')$ represented by
 356 its signature. Let H'_u be the set of all vertices in the node u in P' . Each vertex in H'_u corresponds to
 357 an edge in Q' that lies in the disk D_u ; these edges are noncrossing chords of the circle ∂D_u . We define
 358 a partial ordering on H'_u : For two vertices $u_1, u_2 \in H'_u$, let $u_1 \prec u_2$ if the chord corresponding to u_1
 359 separates the chord of u_2 from N_{uv} within the disk D_u . Intuitively, we have $u_1 \prec u_2$ if u_1 blocks u_2
 360 from the corridor N_{uv} . Note that if $u_1 \prec u_2$, then neither endpoint of the chord corresponding to u_1
 361 is on the boundary of N_{uv} ; consequently u_1 was obtained from a path $[u, v, u]$ or $[u, v, u, v, u, \dots, u]$
 362 in P after removing one or more spurs. We expand the paths $u_i \in H'_u$ incrementally, in an order
 363 determined by any linear extension of the partial ordering \prec . Replace the first vertex $u_1 \in H'_u$ by
 364 $[p, u, v, u]$ (or $[p, u, v, u, v, u, \dots, u, q]$ if needed), and modify the signature by inserting consecutive
 365 new edges into the total order of the edges along uv at any position that is not separated from
 366 the chord in D_u that corresponds to u_1 . The resulting polygon P'' and the new signature define a
 367 polygon Q'' in the strip system of P'' . By construction, the new edges in D_v connect consecutive
 368 endpoints in counterclockwise order around v , thus the new matching in D_v is noncrossing. In the
 369 disk D_u , the operation replaces the chord corresponding to u_1 by noncrossing new chords. Each
 370 new edge in D_u has at least one endpoint in N_{uv} ; consequently, none of them blocks access to $N_{u,v}$.
 371 Then, the new matching in D_u has no crossing and $Q'' \in \Phi(P'')$. By repeating this procedure we
 372 obtain P and a simple polygon $Q \in \Phi(P)$, hence P is weakly simple. \square

373 **Lemma 5.** *Operation node-split is ws-equivalent.*

374 *Proof.* Let P' be obtained from P via node-split(u, v, w). First assume that P is weakly simple.
 375 Then there is a simple polygon $Q \in \Phi(P)$. Consider the clockwise order of edges around v . Since
 376 Q is simple, the order of the edges $[u, v]$ of paths $[u, v, w]$ must be the reverse order of its adjacent
 377 edges $[v, w]$ (the paths must be nested as shown in Figure 15(right)). Because P has no spurs of
 378 the form $[u, v, u]$ or $[w, v, w]$, and the edges of P that pass through v avoid both uv and vw , every
 379 edge between a pair of adjacent edges $[u, v]$ and $[v, w]$ is also part of a path $[u, v, w]$. Replace the

380 paths $[u, v, w]$ by $[u, v^*, w]$ and set the order of edges at segments uv^* and v^*w to be the same order
381 of the removed edges at uv and vw . This defines a polygon $Q' \in \Phi(P')$, which is simple because
382 the circular order of endpoints around D_u and D_w remains unchanged and the matching in D_{v^*} is
383 a subset of the matching in D_v .

384 Now, assume that P' is weakly simple. Since the face in the image graph bounded by u, v, w, v^* is
385 empty, we can change the embedding of the graph by bringing v^* arbitrarily close to v , maintaining
386 weak simplicity. Let δ be the distance between v^* and v . Let $Q' \in \Phi(P')$ be a simple polygon
387 defined on disks of radius ε . Then, Q' is within $\varepsilon + \delta$ Fréchet distance from P and therefore P is
388 weakly simple. \square

389 4.3 Operations

390 We describe three complex operations: pin-extraction, V-shortcut, and L-shortcut. In Section 4.4, we
391 show how to use them to eliminate spurs along any given bar b . The pin-extraction and V-shortcut
392 operations eliminate pins and V-chains. Chains in \mathcal{P} with two or more vertices in the interior of
393 D_b are simplified incrementally, removing one vertex at a time, by the L-shortcut operation.

394 Since the image graph is determined by the polygon, it would suffice to describe how the
395 operations modify the polygon. However, it is sometimes more convenient to first define new nodes
396 and segments in the image graph, and use them to describe the changes in the polygon. In the last
397 step of these operations, we remove any node (segment) that contains no vertex (edge), to ensure
398 that the image graph is consistent with the polygon.

399 **pin-extraction**(u, v). Assume that P satisfies (I1)–(I4) and contains a pin $[v, u, v] \in \mathcal{P}in$.
400 By (I3), node v is adjacent to a unique node w outside of D_b . Perform the following three
401 primitives: (1) **subdivision** of every path $[v, w]$ into $[v, w^*, w]$; (2) **spur-reduction**(v, u).
402 (3) **spur-reduction**(w^*, v). (4) Update the image graph. See Figure 16 for an example.

403 **V-shortcut**(v_1, u, v_2). Assume that P satisfies (I1)–(I4) and $[v_1, u, v_2] \in \mathcal{V}$. Furthermore,
404 P contains no pin of the form $[v_1, u, v_1]$ or $[v_2, u, v_2]$, and no edge $[u, q]$ such that segment
405 uq is in the interior of the wedge $\angle v_1uv_2$. By (I3), nodes v_1 and v_2 are each adjacent
406 to unique nodes w_1 and w_2 outside of D_b , respectively.

407 The operation executes the following primitives sequentially: (1) **node-split**(v_1, u, v_2),
408 which creates a temporary node u^* ; (2) **node-split**(u^*, v_1, w_1) and **node-split**(u^*, v_2, w_2);
409 which create $v_1^*, v_2^* \in \partial D_b$, respectively; (3) **merge** every path $[v_1^*, u^*, v_2^*]$ to $[v_1^*, v_2^*]$. (4)
410 Update the image graph. See Figure 17 for an example.

411 **Lemma 6.** *pin-extraction and V-shortcut are ws-equivalent and maintain (A1)–(A2) in D_b and*
412 *(I1)–(I4) in adjacent clusters.*

413 *Proof.* **pin-extraction.** By construction, the operation maintains (A1)–(A2) in D_b and (I1)–(I4) in
414 adjacent clusters. Also, (I3)–(I4) ensure that **spur-reduction**(v, u) in step (2) satisfies its precondi-
415 tions. Consequently, all three primitives are ws-equivalent.

416 **V-shortcut.** By construction, the operation maintains (A1)–(A2) in D_b and (I1)–(I4) in adjacent
417 clusters. The first two primitives are ws-equivalent by Lemma 5. The third step is ws-equivalent
418 because triangle $\Delta(u^*v_1^*v_2^*)$ is empty of nodes and segments, by assumption. \square




Figure 16: pin-extraction. Changes in the image graph (top), changes in the polygon (bottom).




Figure 17: V-shortcut. Changes in the image graph (top), changes in the polygon (bottom).

419 **L-shortcut operation.** The purpose of this operation is to eliminate a vertex of a path that has
 420 an edge along a given bar. Before describing the operation, we introduce some notation; refer to
 421 Figure 18. For a node $v \in \partial D_b$, let L_v be the set of paths $[v, u_1, u_2]$ in \mathcal{P} such that $u_1, u_2 \in \text{int}(D_b)$.
 422 Each path in \mathcal{P} is either in $\mathcal{P}in$, in \mathcal{V} , or has two subpaths in some L_v . Let M_{cr} be the set of
 423 longest edges of cross-chains in \mathcal{P} . Denote by $\widehat{L}_v \subset L_v$ the set of paths $[v, u_1, u_2]$, where $[u_1, u_2]$ is
 424 *not* in M_{cr} .




Figure 18: Paths in $\mathcal{P}in$, \mathcal{V} , L_v^{TR} , L_v^{TL} , L_v^{BR} , and L_v^{BL} .

425 We partition L_v into four subsets (refer to Figure 18): a path $[v, u_1, u_2] \in L_v$ is in

- 426 1. L_v^{TR} (*top-right*) if v is a *top* vertex and $x(u_1) < x(u_2)$;
- 427 2. L_v^{TL} (*top-left*) if v is a *top* vertex and $x(u_1) > x(u_2)$;
- 428 3. L_v^{BR} (*bottom-right*) if v is a *bottom* vertex and $x(u_1) < x(u_2)$;

429 4. L_v^{BL} (*bottom-left*) if v is a *bottom* vertex and $x(u_1) < x(u_2)$.

430 We partition \widehat{L}_v into four subsets analogously. We define the operation **L-shortcut** for paths in L_v^{TR} ;
 431 the definition for the other subsets can be obtained by suitable reflections.

432 **L-shortcut**(v, TR). Assume that P satisfies (I1)–(I4), $v \in \partial D_b$ and $L_v^{TR} \neq \emptyset$. By (I3), v
 433 is adjacent to a unique node $u_1 \in b$ and to a unique node $w \notin D_b$. Let U denote the set
 434 of all nodes u_2 for which $[v, u_1, u_2] \in L_v^{TR}$. Let $u_{\min} \in U$ and $u_{\max} \in U$ be the leftmost
 435 and rightmost node in U , respectively. Further assume that P satisfies:

- 436 (B1) there is no pin of the form $[v, u_1, v]$;
- 437 (B2) no edge $[p, u_1]$ such that segment pu_1 is in the interior of the wedge $\angle vu_1u_{\min}$;
- 438 (B3) no edge $[p, q]$ such that $p \in \partial D_b$ is a top vertex and $q \in b$, $x(u_1) < x(q) < x(u_{\max})$.

439 Do the following (see Figure 19 for an example).

- 440 (0) Create a new node $v^* \in \partial D_b$ to the right of v sufficiently close to v .
- 441 (1) For every path $[v, u_1, u_2] \in L_v^{TR}$ in which u_1u_2 is the *only* longest edge of a cross-
 442 chain, create a crimp by replacing $[u_1, u_2]$ with $[u_1, u_2, u_1, u_2]$.
- 443 (2) Replace every path $[w, v, u_1, u_{\min}]$ by $[w, v^*, u_{\min}]$.
- 444 (3) Replace every path $[w, v, u_1, u_2]$, where $u_2 \in U$ and $u_2 \neq u_{\min}$, by $[w, v^*, u_{\min}, u_2]$.
- 445 (4) Update the image graph.

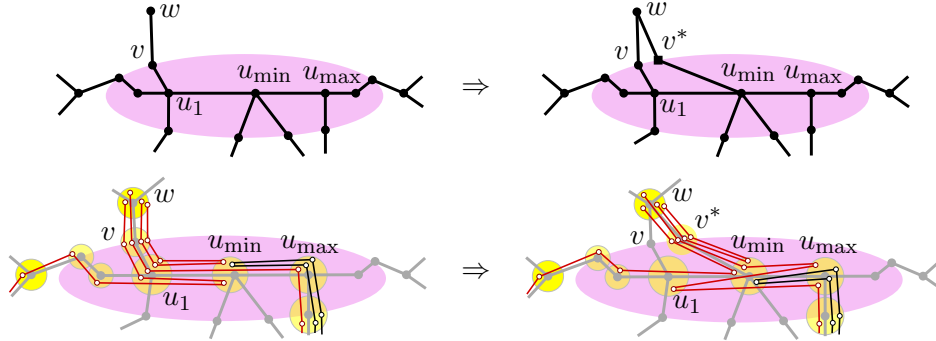


Figure 19: **L-shortcut**. Changes in the image graph (top), changes in the polygon (bottom).

446 See Figure 20 for an explanation of why **L-shortcut** requires conditions (B2)–(B3) and phase (1)
 447 of the operation. If we omit any of these conditions, **L-shortcut** would not be ws-equivalent.

448 **Lemma 7.** **L-shortcut** is ws-equivalent and maintains (A1)–(A2) in D_b and (I1)–(I4) in adjacent
 449 clusters.

450 *Proof.* Let P_1 be the polygon obtained from P after phase (1) of **L-shortcut**(v, TR) and P_2 be the
 451 polygon obtained after phase (3). Note that phase (1) of the operation only creates crimps, and it
 452 is ws-equivalent by Lemma 1. Let H be the set of edges $[u_1, u_2]$ of paths $[v, u_1, u_2] \in L_v^{TR}$. Phases
 453 (2)–(3) are equivalent to the concatenation of the primitives: **subdivision**, **node-split**, and **merge**.
 454 Specifically, they are equivalent to subdividing every edge in H into $[u_1, u_{\min}, u_2]$ whenever $u_2 \neq$
 455 u_{\min} , and applying **node-split**(v, u_1, u_{\min}) (which creates u_1^*) to P_2 followed by **node-split**(w, v, u_1^*)




Figure 20: Cases in which L-shortcut is not ws-equivalent. Top left: P does not satisfy (B2). Top right: P does not satisfy (B3). Bottom: the operation skips phase (1).

456 (which creates v^*), and merging every path $[v^*, u_1^*, u_{\min}]$ to $[v^*, u_{\min}]$. The only primitive that may
 457 not satisfy its preconditions is `node-split`(v, u_1, u_{\min}): segment $u_1 u_{\min}$ may be collinear with several
 458 segments of b , and P_2 may contain spurs that overlap with $u_1 u_{\min}$. In the next paragraph, we show
 459 that the spurs that may overlap with $u_1 u_{\min}$ do not pose a problem, and we can essentially repeat
 460 the proof of Lemma 5.

461 Assume that P_1 is weakly simple and consider a polygon $Q_1 \in \Phi(P_1)$. Due to (A1)–(A2) and
 462 phase (1), every path in L_v^{TR} is a sub-path of some path $[v, u_1, u_2, u_3]$ where $x(u_3) \leq x(u_2)$. We
 463 show that P_1 has a perturbation in $\Phi(P_1)$ with the following property:

464 (\star) Every edge $[u_1, u_2] \in H$ lies above all overlapping edges $e \notin H$.

465 Let $Q_1 \in \Phi(P_1)$ be a perturbation of P_1 into a simple polygon that has the minimum number of
 466 edges $[u_1, u_2] \in H$ that violate (\star) . We claim that Q_1 satisfies (\star) . Suppose the contrary, that Q_1
 467 does not satisfy (\star) . For a contradiction, we modify $Q_1 \in \Phi(P_1)$ and obtain another perturbation
 468 $Q'_1 \in \Phi(P_1)$ that has strictly fewer edges that violate (\star) as shown in Figure 21. Recall that
 469 Q_1 yields a total order of edges in each segment of b based on the above-below relationship. Let
 470 $[u_1, u'_2] \in H$ be the highest edge that violates (\star) , and assume that this edge is part of a path
 471 $[v, u_1, u'_2, u'_3]$. Let Z be the set of edges that are above $[u_1, u'_2]$ within the corridors between u_1 and
 472 u'_2 , and are not in H . By (B2)–(B3) and Lemma 2, every edge $[z_1, z_2] \in Z$ must be part of a path
 473 $[z_1, z_2, z_3]$ where $x(u_1) \leq x(z_2) < x(u'_2) \leq x(z_1)$ and $x(u'_2) \leq x(z_3)$, otherwise Q_1 would not be
 474 simple. We modify $\sigma(Q_1)$ by moving the edges in Z , maintaining their relative order, immediately
 475 below edge $[u'_2, u'_3]$ in all segments between u_1 and u'_2 . This results in a simple polygon $Q'_1 \in \Phi(P_1)$
 476 such that $[u_1, u'_2]$ and all edges in H above $[u_1, u'_2]$ satisfy (\star) , contradicting the choice of Q_1 .

477 We can proceed as in the proof of Lemma 5, using a perturbation $Q_1 \in \Phi(P_1)$ that satisfies
 478 (\star) to show that P_2 is weakly simple if and only if P_1 is weakly simple, that is, phases (2)–(3) are
 479 ws-equivalent.

480 By construction, (I1)–(I4) are maintained. Note that the intermediate polygon P_1 may violate
 481 condition (A2), since phase (1) introduces crimps. However, after phase (3), conditions (A1) and
 482 (A2) are restored, and operation L-shortcut maintains (A1)–(A2) in the ellipse D_b . \square




Figure 21: (a) A perturbation Q_1 that violates property (\star) ; the highest edge $[u_1, u'_2] \in H$ that violates (\star) is red, and edges in Z are blue. (b) We can modify Q_1 to reduce the number of edges in H that violate (\star) . (c) There exists a perturbation Q_1 that satisfies (\star) .

483 4.4 Bar simplification algorithm

484 In this section, we describe an algorithm, called **bar-simplification**, that incrementally removes all
 485 spurs of the polygon P from a bar b using a sequence of pin-extraction, V-shortcut, and L-shortcut
 486 operations. Informally, our algorithm “unwinds” each polygonal chain in the bar. It extracts
 487 pins and V-chains whenever possible. Any other chain in D_b contains edges along bar b , and
 488 the sequence of these edge lengths is unimodal (cf. Lemma 2). Our algorithm “unwinds” these
 489 chains by a sequence of L-shortcut operations. Each operation eliminates or reduces one of the
 490 shortest edges along b (see Figure 22). The algorithm alternates between L-shortcut(v, TR) and
 491 L-shortcut(v, TL) to unwind the chains from their top endpoints to the longest edge in b ; and then
 492 uses L-shortcut(v, BR) and L-shortcut(v, BL) to resolve the bottom part.

493 When we unwind the chains in D_b starting from their top vertices using L-shortcut(v, TR) and
 494 L-shortcut(v, TL), we cannot hope to remove the longest edge of a cross-chain. We stop using the
 495 operations when every path in L_v^{TR} contains a longest edge of a cross-chain. This motivates the
 496 use of \widehat{L}_v^{TR} (instead of L_v^{TR}) in step (iii) below. We continue with the algorithm and its analysis.




Figure 22: Life cycle of a cross-chain in the while loop of bar-simplification. The steps applied, from left to right, are: (iii), (iv), (iii), (iv), and (vi).

497 Algorithm **bar-simplification**(P, b).

498 While P has an edge along b , perform one operation as follows.

- 499 (i) If $\mathcal{P}in \neq \emptyset$, pick an arbitrary pin $[v, u, v]$ and perform **pin-extraction**(u, v).
- 500 (ii) Else if $\mathcal{V} \neq \emptyset$, then let $[v_1, u, v_2] \in \mathcal{V}$ be a path where $|x(v_1) - x(v_2)|$ is minimal. If there is
 501 no segment uq in the wedge $\angle v_1uv_2$, perform **V-shortcut**(v_1, u, v_2), else report that P is not
 502 weakly simple and halt.
- 503 (iii) Else if there exists $v \in \partial D_b$ such that $\widehat{L}_v^{TR} \neq \emptyset$, do:
- 504 (a) Let v be the rightmost node where $L_v^{TR} \neq \emptyset$.
- 505 (b) If L_v^{TR} satisfies (B2)–(B3), do **L-shortcut**(v, TR).
- 506 (c) Else let v' be the leftmost node such that $x(v) < x(v')$ and $L_{v'}^{TL} \neq \emptyset$, or record that no
 507 such vertex v' exists.

508 (c.1) If v' does not exist, or $L_{v'}^{TL}$ does not satisfy (B2)–(B3), or any path in $L_{v'}^{TL}$ contains
 509 a longest edge of a cross-chain, then report that P is not weakly simple and halt.

510 (c.2) Else do **L-shortcut**(v', TL).

511 (iv) Else if there exists $v \in \partial D_b$ such that $L_v^{TL} \neq \emptyset$, perform steps (iii)a–(iii)c with left–right and
 512 TR – TL interchanged. (Note the use of L_v instead of \widehat{L}_v . The same applies to (vi) below).

513 (v) Else if there exists $v \in \partial D_b$ such that $\widehat{L}_v^{BL} \neq \emptyset$, perform steps (iii)a–(iii)c using BL and BR
 514 in place of TR and TL , respectively, and left–right interchanged.

515 (vi) Else if there exists $v \in \partial D_b$ such that $L_v^{BR} \neq \emptyset$, perform steps (iii)a–(iii)c using BR and BL
 516 in place of TR and TL , respectively.

517 (vii) Else invoke **old-bar-expansion**.

518 Return P (end of algorithm).

519

520 **Lemma 8.** *The operations performed by bar-simplification(P, b) are ws-equivalent, and maintain*
 521 *properties (A1)–(A2) in D_b and (I1)–(I4) in adjacent clusters. The algorithm either removes*
 522 *all nodes from the ellipse D_b , or reports that P is not weakly simple. The L-shortcut operations*
 523 *performed by the algorithm create at most two crimps in each cross-chain in \mathcal{P} .*

524 *Proof.* We show that the algorithm only uses operations that satisfy their preconditions, and reports
 525 that P is not weakly simple only when P contains a forbidden configuration.

526 **Steps (i)–(ii).** Since every pin can be extracted from a polygon satisfying (I1)–(I4), we may assume
 527 that $\mathcal{P}in = \emptyset$. Suppose that $\mathcal{V} \neq \emptyset$. Let $[v_1, u, v_2] \in \mathcal{V}$ be a V-chain such that $|x(v_1) - x(v_2)|$ is
 528 minimal. Since $\mathcal{P}in = \emptyset$, the only obstacle for the precondition of **V-shortcut** is an edge $[u, q]$
 529 such that segment uq is in the interior of the wedge $\angle v_1uv_2$ (or else the image graph would have a
 530 crossing). If such an edge exists, it is part of a path $[p, u, q]$. The node q is in ∂D_b between v_1 and v_2 .
 531 Note that $p \neq q$, otherwise $[p, u, q]$ would be a pin. Further, p cannot be a node in the interior of the
 532 wedge $\angle v_1uv_2$, otherwise $[p, u, q]$ would be a V-chain where $|x(p) - x(q)| < |x(v_1) - x(v_2)|$, contrary
 533 to the choice of $[v_1, u, v_2] \in \mathcal{V}$. Consequently, p must be in the exterior of the wedge $\angle v_1uv_2$. In
 534 this case, the paths $[v_1, u, v_2]$ and $[p, u, q]$ form the forbidden configuration in Corollary 1(1), and
 535 the algorithm correctly reports that P is not weakly simple. If no such edge $[u, q]$ exists, then
 536 **V-shortcut**(v_1, u, v_2) satisfies all preconditions and it is ws-equivalent by Lemma 6. Henceforth, we
 537 may assume that $\mathcal{P}in = \emptyset$ and $\mathcal{V} = \emptyset$.

538 **Step (iii)–(iv).** By symmetry, we consider only step (iii). Since $\mathcal{P}in = \emptyset$, condition (B1) is met.
 539 In step (iii)b, if (B2)–(B3) are also satisfied, then **L-shortcut**(v, TR) is ws-equivalent by Lemma 7.
 540 If condition (B2) or (B3) fails, we proceed with step (iii)c.

541 **Step (iii)(c.1).** We show that in these cases the algorithm correctly reports that P is not weakly
 542 simple. Assume first that v' does not exist. Since L_v^{TR} does not satisfy (B2) or (B3), there exists
 543 an edge $[p, q]$ such that $x(u_1) \leq x(q) < x(u_{\max})$ and $p \in \partial D_b$ is a top node. Edge $[p, q]$ is part of
 544 some path $[p, q, r]$. Note that r cannot be a top vertex of ∂D_b , since $\mathcal{P}in = \emptyset$ and $\mathcal{V} = \emptyset$. If r is
 545 on b and $x(q) < x(r)$, then $[p, q, r] \in L_p^{TR}$, which contradicts the choice of node v . If r is on b and
 546 $x(r) < x(q)$, then $[p, q, r] \in L_p^{TL}$ and v' exists. It follows that r is a bottom vertex, and then the
 547 paths $[v, u_1, u_{\max}]$ and $[p, q, r]$ form a forbidden configuration in Corollary 1(1) or (3).

548 Assume now that v' exists but $L_{v'}^{TL}$ does not satisfy (B2) or (B3). Let $[v', u'_1, u'_{\max}]$ be the path
 549 in $L_{v'}^{TL}$ with the longest edge on b . By the definitions of (B2)–(B3), $x(u_1) \leq x(u'_1) < x(u_{\max})$.

550 If $x(u'_{\max}) < x(u_1)$, then $[v, u_1, u_{\max}]$ and $[v', u'_1, u'_{\max}]$ form the forbidden configuration in Corol-
551 lary 1(2). Else, we have $x(u_1) \leq x(u'_{\max}) < x(u'_1) < x(u_{\max})$. This implies that any edge $[p, q]$ that
552 violates (B2) or (B3) for L_v^{TL} must also violate (B2) or (B3) for L_v^{TR} . However, this contradicts
553 the choice of v (rightmost where $L_v^{TR} \neq \emptyset$) and v' (leftmost, $x(v) < x(v')$, where $L_{v'}^{TL} \neq \emptyset$).

554 Next assume that there is a path $[v', u'_1, u'_2] \in L_{v'}^{TL}$ such that $[u'_1, u'_2]$ is the longest edge of a
555 cross-chain. Then this cross-chain is of the form $[v', u'_1, u'_2, \dots, p']$, where all interior vertices lie
556 on the line segment $u'_1 u'_2$, and p' is a bottom vertex. Now $[v, u_1, u_{\max}]$ and this cross-chain form
557 the forbidden configuration in Corollary 1(3). In all three cases in step (iii)(c.1), the algorithm
558 correctly reports that P is not weakly simple.

559 **Step (iii)(c.2).** Let the path $[v', u'_1, u'_{\max}] \in L_{v'}^{TL}$ be selected in $\text{L-shortcut}(v', TL)$ by the algorithm.
560 Since conditions (B1)–(B3) are satisfied, $\text{L-shortcut}(v', TL)$ is ws-equivalent by Lemma 7.

561 **Steps (v)–(vii).** If steps (i)–(iv) do not apply, then $\widehat{L}_v^{TR} \cup L_v^{TL} = \emptyset$. That is, for every path
562 $[v, u_1, u_2] \in L^{TR}$, we have $[u_1, u_2] \in M_{cr}$. In particular, there are no top chains. The operations in
563 (v)–(vi) do not change these properties. Consequently, once steps (v)–(vi) are executed for the first
564 time, steps (iii)–(iv) are never executed again. By a symmetric argument, steps (v)–(vi) eliminate
565 all paths in $\widehat{L}_v^{BL} \cup L_v^{BR}$. When the algorithm reaches step (vii), every edge in b is necessarily in
566 M_{cr} and $L_v^{TL} \cup L_v^{BR} = \emptyset$. Consequently, by Lemma 2, b contains no spurs and old-bar-expansion is
567 ws-equivalent. This operation eliminates all nodes in the interior of D_b .

568 **Termination.** Each pin-extraction and V-shortcut operation reduces the number of vertices of
569 P within D_b . Operation $\text{L-shortcut}(v, X)$, $X \in \{TR, TL, BR, BL\}$, either reduces the number
570 of interior vertices, or produces a crimp if edge $[u_1, u_2]$ is a longest edge of a cross-chain. For
571 termination, it is enough to show that, for each cross-chain $c \in \mathcal{P}$, the algorithm introduces a crimp
572 at most once in steps (iii)–(iv), and at most once in steps (v)–(vi). Without loss of generality,
573 consider step (iii).

574 Note that step (iii) may apply an L-shortcut operation in two possible cases: (iii)b and (iii)c.
575 However, an L-shortcut operation in (iii)c does not create crimps: L-shortcut is performed when all
576 three conditions in (iii)(c.1) fail. In this case, $L_{v'}^{TR}$ does not contain any edge in M_{cr} , and L-shortcut
577 does not create crimps. We may assume that step (iii) creates crimps in case (iii)b only.

578 Every cross-chain remains a cross-chain in algorithm bar-simplification: operations pin-extraction
579 and V-shortcut do not modify cross-chains; and operations L-shortcut and old-bar-expansion modify
580 only the first or last few edges of a cross-chain. A longest edge of a cross-chain c always connects
581 the same two nodes in b until step (vii) (old-bar-expansion), although the *number* of longest edges
582 in c may change. When $\text{L-shortcut}(v, X)$ modifies a cross-chain, it moves its endpoint from $v \in \partial D_b$
583 to a nearby new node $v^* \in \partial D$. Consequently, if L_v^X , $X \in \{TR, TL\}$ contains the first two edges
584 of two chains in \mathcal{P} , then they have been modified by the same sequence of previous L-shortcut
585 operations.

586 Suppose, for contradiction, that two invocations of step (iii)b create crimps in a cross-chain c ,
587 say, in operations $\text{L-shortcut}(v_0, TR)$ and $\text{L-shortcut}(v_2, TR)$ (see Figure 23). The first invocation
588 replaces $[v_0, u_1, u_2]$ with $[v_0^*, u_{\min}, u_2, u_1, u_2]$ (where the edge $[u_{\min}, u_2]$ may vanish if $u_{\min} = u_2$).
589 The resulting cross-chain has two maximal longest edges, $[u_2, u_1]$ and $[u_1, u_2]$. Since L-shortcut
590 creates crimps only if the longest edge is unique, there must be an intermediate operation L-
591 shortcut(v_1, TL) that removes or shortens the edge $[u_2, u_1]$, so that $[u_1, u_2]$ becomes the unique
592 longest edge again. When $\text{L-shortcut}(v_1, TL)$ is performed in a step (iv), we have $\widehat{L}_v^{TR} = \emptyset$ for all
593 top nodes v , and $L_{v'}^{TL} = \emptyset$ for all top nodes v' , $x(v') < x(v_1)$. The steps between $\text{L-shortcut}(v_1, TL)$
594 and $\text{L-shortcut}(v_2, TR)$ modify only cross-chains whose top node is at or to the right of the top node

595 of c (L-shortcut operations move the top vertex of c to the left, from v_1 to v_2 in one or more steps).
 596 Consequently, when $\text{L-shortcut}(v_2, TR)$ is performed in a step (iii), we still have $\widehat{L}_{v'}^{TR} = L_{v'}^{TL} = \emptyset$
 597 for all top nodes v' , $x(v') < x(v_2)$.




Figure 23: At most one crimp can be created in a cross-chain by steps (iii)b.

598 When $\text{L-shortcut}(v_2, TR)$ is performed, we have $[v_2, u_1, u_2] \in L_{v_2}^{TR}$ but $[v_2, u_1, u_2] \notin \widehat{L}_{v_2}^{TR}$ (since
 599 $u_1 u_2$ is the longest edge of c). Step (iii) is performed only if $\widehat{L}_p^{TR} \neq \emptyset$ for some top vertex p .
 600 Since the rightmost top vertex where $L_v^{TR} \neq \emptyset$ is $v = v_2$, we have $x(p) \leq x(v_2)$. This implies
 601 $p = v_2$. Consequently there exists a chain $c' \in \mathcal{P}$ that contains a subpath $[v_2, u_1, u_3] \in L_{v_2}^{TR}$, such
 602 that $[u_1, u_3]$ is not the longest edge of c' . Since $L_{v_2}^{TR}$ contains the first two edges of both c and
 603 c' , they have been modified by the same sequence of L-shortcut operations. Therefore c' contained
 604 $[v, u_1, u_2, u_1]$ initially. By Lemma 2, only the longest edge can repeat, hence $[u_1, u_2]$ is the longest
 605 edge of c' . This implies that $u_3 = u_2$ and $\widehat{L}_{v_2}^{TR} = \emptyset$, contradicting the condition in Step (iii).

606 We conclude that $\text{bar-simplification}(P, b)$ introduces a crimp at most once in steps (iii)–(iv), and
 607 at most once in steps (v)–(vi) in each cross-chain. Since all other steps decrease the number of
 608 vertices in D_b , the algorithm terminates, as claimed. \square

609 **Lemma 9.** *Algorithm $\text{bar-simplification}(P, b)$ takes $O(m \log m)$ time using suitable data structures,*
 610 *where m is the number of vertices in b .*

611 *Proof.* Operations pin-extraction, V-shortcut, and L-shortcut each make $O(1)$ changes in the image
 612 graph. Operations pin-extraction and V-shortcut decrease the number of vertices inside D_b . Each
 613 L-shortcut does as well, except for the steps that create crimps. By Lemma 7, L-shortcut operations
 614 may create at most $2|\mathcal{P}| = O(m)$ crimps. So the total number of operations is $O(m)$.

615 When $[v, u_1, u_2] \in L_v^{TR}$ and $u_2 \neq u_{\min}$, L-shortcut replaces $[v, u_1, u_2]$ by $[v^*, u_{\min}, u_2]$: vertex
 616 $[u_1]$ shifts to $[u_2]$, but no vertex is eliminated. In the worst case, one L-shortcut modifies $\Theta(m)$
 617 paths, so in $\Theta(m)$ operations the total number of vertex shifts is $\Theta(m^2)$.

618 **Data structures.** We maintain a cyclic list of nodes in ∂D_b given by the combinatorial embedding
 619 of the image graph. Since each operation adds a constant number of nodes to ∂D_b at positions
 620 adjacent to the nodes to which the operation was applied, such a list can be maintained using $O(1)$
 621 time per operation. Our implementation does not maintain the paths in \mathcal{P} explicitly. Instead, we use
 622 set operations. We maintain the sets \mathcal{Pin} , \mathcal{V} , and L_v^X , with $v \in \partial D_b$ and $X \in \{TR, TL, BR, BL\}$,
 623 in sorted lists. The pins $[v, u, v] \in \mathcal{Pin}$ are sorted by $x(v)$; the wedges $[v_1, u, v_2] \in \mathcal{V}$ are sorted by
 624 $|x(v_1) - x(v_2)|$. In every set L_v^X , the first two nodes in the paths $[v, u_1, u_2] \in L_v^X$ are the same by
 625 (I3)b, and so it is enough to store vertex $[u_2]$; these vertices are stored in a list sorted by $x(u_2)$.
 626 We also maintain binary variables to indicate for each path $[v, u_1, u_2] \in L_v^X$ whether it is part of a
 627 cross-chain, and whether $[u_1, u_2]$ is the only longest edge of that chain.

628 **Running time analysis.** The condition in step (ii) can be tested in $O(1)$ time by checking whether
 629 uv_1 and uv_2 are consecutive segments in the rotation of node u in the image graph. Steps (i)-(ii)

630 remove pins and V-chains, taking linear time in the number of removed vertices, without introducing
 631 any path in any set.

632 Consider $L\text{-shortcut}(v, TR)$, executed in step (iii), which can be generalized to other occurrences
 633 of the $L\text{-shortcut}$ operation performed in one of steps (iii)–(vi). Recall that $\mathcal{P}in = \mathcal{V} = \emptyset$. Let p'
 634 be the leftmost top vertex in ∂D_b to the right of v , which can be found in $O(1)$ time using the
 635 cyclic list of nodes in ∂D_b . By (I3)b, every path $[p', q', r'] \in L_{p'}^X$, $X \in \{TR, TL\}$, must contain the
 636 edge $[p', q']$. If $x(q') < x(u_{\max})$, then (B2) or (B3) are not satisfied. Assume that $x(q') \geq x(u_{\max})$
 637 and (B2) (resp., (B3)) is not satisfied. Then there must exist an edge $[p, u_1]$ (resp., $[p, q]$ where
 638 $x(q) < x(u_{\max})$) such that p is to the right of p' . Then, segments $p'q'$ and pu_1 (resp., pq) properly
 639 cross. This is a contradiction since no operation introduces crossings in the image graph. Hence
 640 (B2)–(B3) are satisfied if and only if either p' does not exist (i.e., v is the rightmost top vertex), or
 641 $x(u_{\max}) \leq x(q')$; this can be tested in $O(1)$ time. The elements $[v, u_1, u_{\min}] \in L_v^{TR}$ are simplified to
 642 $[v^*, u_{\min}]$. Consider one of these paths, and assume that the next edge along P is $[u_{\min}, u_3]$. Then,
 643 the path $[v^*, u_{\min}, u_3]$ is inserted into either $\mathcal{P}in \cup \mathcal{V}$ if $u_3 \in \partial D_b$ is a top vertex, or $L_{v^*}^{TL}$ if $u_3 \in b$.
 644 We can find each chain $[v, u_1, u_{\min}] \in L_v^{TR}$ in $O(1)$ time since L_v^{TR} is sorted by $x(u_2)$. Finally, all
 645 other paths of the form $[v, u_1, u_2] \in L_v^{TR}$, where $u_2 \neq u_{\min}$, become $[v^*, u_{\min}, u_2]$ and they form the
 646 new set $L_{v^*}^{TR}$. Since we store only the last vertex $[u_2]$, which is unchanged, we create $L_{v^*}^{TR}$ at no
 647 cost.

648 This representation allows the manipulation of $O(m)$ vertices with one set operation. The
 649 number of insert and delete operations in the sorted lists is proportional to the number of vertices
 650 that are removed from the interior of D_b , which is $O(m)$. Each insertion and deletion takes $O(\log m)$
 651 time, and the overall time complexity is $O(m \log m)$. \square

652 5 Spur elimination algorithm

653 After bar-simplification (Section 4), we obtain a polygon that has no forks and every spur is at
 654 an interior node of some cluster (formed on the boundary of some ellipse D_b). In the absence
 655 of forks, we can decide weak simplicity using [6, Theorem 5.1], but a naïve implementation runs
 656 in $O(n^2 \log n)$ time: successive applications of `spur-reduction` would perform an operation at each
 657 dummy vertex. In this section, we show how to eliminate spurs in $O(n \log n)$ time.

658 **Formation of Groups.** We create *groups* by gluing pairs of clusters with adjacent roots together.
 659 Recall that by (I1) each cluster induces a tree. We modify the image graph, transforming each
 660 tree in a cluster into a binary tree using `ws-equivalent` primitives. For each node s with more than
 661 two children, let s_1 and s_2 be the first two children in counterclockwise order. Create new nodes
 662 s'_1 and s'_2 by `subdivision` in ss_1 and ss_2 , respectively, and create a segment $s'_1s'_2$. Use the inverse of
 663 `node-split` to merge nodes s'_1 and s'_2 into a node s' , reducing the number of children of s by one.

664 In the course of our algorithm, an analogue of the `pin-extraction` operation extracts a spur from
 665 one group into an “adjacent” group. This requires a well-defined adjacency relation between groups.
 666 By construction, if a segment uv connects nodes in different clusters, both u and v are leaves or
 667 both are root nodes. For every pair of clusters, $C(u)$ and $C(v)$, with adjacent roots, u and v , create
 668 a *group* $G_{uv} = C(u) \cup C(v)$; see Figure 24. By construction, the groups are pairwise disjoint. Two
 669 groups are called *adjacent* if they have two adjacent leaves in the image graph.

670 Recall that a maximal path in each cluster is represented by benchmark vertices (leaves and
 671 spurs). We denote by $[u_1; \dots; u_k]$ (using semicolons) a maximal path inside a group defined by the




Figure 24: The formation of a group G_{uv} , containing clusters $C(u)$ and $C(v)$. Leaf nodes are shown as black dots.

672 benchmark vertices u_1, \dots, u_k . For a given group G_{uv} , let \mathcal{P} denote the set of maximal paths with
 673 vertices in G_{uv} ; and let \mathcal{B} be the set of subpaths in \mathcal{P} between consecutive benchmark vertices.

674 **Remark 2.** *By invariants (I1)–(I3), a path in \mathcal{P} of a group G_{uv} has alternating benchmark vertices*
 675 *between $C(u)$ and $C(v)$. Consequently, every path in \mathcal{B} has one endpoint in $C(u)$ and one in $C(v)$,*
 676 *and each spur in G_{uv} is incident to two paths in \mathcal{B} .*

677 **Spur-elimination algorithm.** Assume that \mathcal{G} is a partition of the nodes of the image graph into
 678 groups satisfying (I1)–(I4). We consider one group at a time, and eliminate all spurs from one
 679 cluster of that group. When we process one group, we may split it into two groups, create a new
 680 group, or create a new spur in an adjacent group (similar to pin-extraction in Section 4). The latter
 681 operation implies that we may need to process a group several times. Termination is established
 682 by showing that each operation reduces a weighted sum of the number of benchmark vertices (i.e.,
 683 spurs and boundary vertices). Initially, the number of benchmarks is $O(n)$.

684 Algorithm spur-elimination(P, \mathcal{G}).

685 While P contains a spur, do:

- 686 1. Choose a group $G_{uv} \in \mathcal{G}$ that contains a spur, w.l.o.g. contained in cluster $C(u)$,
 687 and create its supporting data structures (described in Section 5.1 below).
- 688 2. While $T[u]$ contains an interior node, do:
 - 689 (a) If u contains no spurs and is incident to only two edges uv and uw , eliminate
 690 u with a merge operation. Rename node w to u which becomes the new root
 691 of the tree $T[u]$.
 - 692 (b) If u contains spurs, eliminate them as described in Section 5.2.
 - 693 (c) If u contains no spurs, split G_{uv} into two groups along a chain of segments
 694 that contains uv as described in Section 5.3. Rename a largest resulting group
 695 to G_{uv} .

696 The detailed description of steps 2b and 2c are in Sections 5.2 and 5.3, respectively. We first
 697 present supporting data structures in Section 5.1, and then analyze the algorithm in Section 5.4.

698 5.1 Data structures

699 In this section, we describe the data structures that we maintain for a group G_{uv} . We start
 700 with reviewing and introducing some notation. Consider a group G_{uv} composed of two binary

701 trees $T[u]$ and $T[v]$ rooted at u and v , respectively. Recall that \mathcal{B} denotes the set of benchmark-
702 to-benchmark paths, each with one benchmark in $T[u]$ and one in $T[v]$. In the algorithm spur-
703 elimination, we dynamically maintain the image trees $T[u] \cup T[v]$, and the set of paths \mathcal{B} . In each
704 group G_{uv} , we maintain only $O(|\mathcal{B}|)$ nodes that contain benchmark vertices or have degree higher
705 than 2. Dummy nodes of degree two that contain no benchmark vertices are redundant for the
706 combinatorial representation, and will be eliminated with merge operations. However, a polyline
707 formed by a chain of dummy nodes of degree two cannot always be replaced by a straight-line
708 segment (this might introduce unnecessary crossings). By Remark 1, it suffices to maintain the
709 combinatorial embeddings of the trees $T[u]$ and $T[v]$ (i.e., the counterclockwise order of the incident
710 segments around each node).

711 The partition of a group into two groups is driven by the partition of the paths in \mathcal{B} . For a
712 set $\mathcal{B}' \subset \mathcal{B}$ of benchmark-to-benchmark paths, we define a subtree $T(\mathcal{B}')$ induced by \mathcal{B}' as follows.
713 Let $N = N(\mathcal{B}')$ be the set of nodes that contain endpoints of some path in \mathcal{B}' . The tree $T(\mathcal{B}')$ is
714 obtained in two steps: take the minimum subtree of $T[u] \cup T[v]$ that contains all nodes in N , and
715 then merge all nodes of degree two that are not in N . In particular, the nodes of $T(\mathcal{B}')$ include
716 N and the lowest common ancestor of any two nodes in $N \cap C(u)$ and in $N \cap C(v)$, respectively.
717 Denote by $\text{lca}(r, s)$ the *lowest common ancestor* of nodes r and s in $T[u]$ (resp., $T[v]$).

718 **Description of data structures.** For the image graph of G_{uv} , we maintain the following data
719 structures.

- 720 • We store trees $T[u]$ and $T[v]$ each using the dynamic data structure of [8], which supports
721 $O(1)$ -time insertion and deletion of leaves, merging interior nodes of degree 2, subdivision of
722 edges, and lowest common ancestor queries.
- 723 • Imagine that G_{uv} is inside an axis-aligned rectangle with the leaves of $T[u]$ along the top
724 edge and leaves of $T[v]$ along the bottom edge (see Figure 25(a)). For each tree, we maintain
725 a left-to-right Euler tour in an order-maintenance data structure [3, 19], which supports
726 insertions immediately before or after an existing item, deletions, and precedence queries,
727 each in $O(1)$ amortized time. For any node w , let w^b and $w^\#$ respectively denote the first and
728 last occurrences of w in the Euler tour. Note that we have $w^\# = w^b$ for a leaf w . We refer to
729 the elements of the Euler tour as *tokens*. We write $x < y$ to denote that some token x occurs
730 before (“to the left of”) another token y in their common Euler tour.
- 731 • We also maintain the cyclic list of all leaves of the tree $T[u] \cup T[v]$ (in the order determined
732 by the Euler tour above).

733 We now describe data structures for \mathcal{P} and \mathcal{B} . For every benchmark-to-benchmark path $[s; t] \in$
734 \mathcal{B} , we assume that s is in $T[u]$ and t is in $T[v]$. A path $[s; t]$ is associated with the intervals $[s^b, s^\#]$
735 and $[t^b, t^\#]$. For two consecutive benchmark-to-benchmark paths $[s_1; t; s_2]$, where t is in $T[v]$, we
736 define the interval $I[s_1; t; s_2] = [s_1^b, s_2^b]$.

- 737 • The set of benchmark-to-benchmark paths $[s; t] \in \mathcal{B}$ is stored in four lists, sorted by s^b , $s^\#$, t^b ,
738 and $t^\#$, respectively, with ties broken arbitrarily. The sorted lists can be computed in $O(|\mathcal{B}|)$
739 time by an Eulerian traversal of the tree.
- 740 • For each node s of $T[u]$, let \mathcal{B}_s denote the set of paths $[s; t] \in \mathcal{B}$. We store \mathcal{B}_s in two lists,
741 sorted by t^b and $t^\#$, respectively.




Figure 25: The geometry of crossing benchmark-to-benchmark paths. (a) Paths $[s_1; t_1]$ and $[s_2; t_2]$ cross. (b) If $t_{\min}^\# < t^b \leq t^\# < t_{\max}^\#$, then any benchmark-to-benchmark path $[s; t]$ crosses path $[t_{\min}; t_{\max}]$.

- 742 • We use a centered *interval tree* [4] for all $O(n)$ intervals $I[s_1; t; s_2]$ that can report, for a query
 743 node q , all intervals containing q in output-sensitive $O(\log n + k)$ time, where k is the number
 744 of intervals that contain q . Since the interval endpoints s^b are already sorted, the interval tree
 745 can be constructed in $O(|\mathcal{B}|)$ time. The interval tree can handle the deletion of an interval in
 746 $O(1)$ time (without re-balancing, hence maintaining the $O(\log n + k)$ query time).

747 All data structures described in this section can be constructed in $O(|\mathcal{B}|)$ preprocessing time.

748 **Crossing paths.** The data structure described above can determine in $O(1)$ time whether two
 749 paths in \mathcal{B} cross. Straightforward case analysis implies the following characterization of path
 750 crossings (refer to Figure 25(a)).

751 **Lemma 10.** *Let s_1 and s_2 be arbitrary nodes in tree $T[u]$, and let t_1 and t_2 be arbitrary nodes in*
 752 *$T[v]$. Paths $[s_1; t_1]$ and $[s_2; t_2]$ cross if and only if either (1) $s_1^\# < s_2^\#$ and $t_1^b > t_2^\#$, or (2) $s_2^\# < s_1^\#$*
 753 *and $t_2^b > t_1^\#$.*

754 5.2 Eliminating spurs from a root

755 We describe step 2b of Algorithm *spur-elimination*. Suppose that the root node u contains a
 756 spur. The following operation eliminates all spurs from u , but the resulting cluster $C(v)$ need not
 757 satisfy (I2) and (I3), and we need to perform other operations to restore these properties. Refer to
 758 Figure 26(a)–(b) for an example.

759 **spur-shortcut(u).** Assume that G_{uv} satisfies invariants (I1)–(I4), and u contains a spur.
 760 Replace every path $[t_1; u; t_2]$ by $[t_1; t_2]$. Let \mathcal{S} be the set of all such modified paths.

761 **Lemma 11.** *spur-shortcut is ws-equivalent and maintains properties (I1) and (I4).*

762 *Proof.* The operation is equivalent to a sequence of spur-reduction operations: First perform spur-
 763 reduction(v, u). In a BFS traversal of all nodes x of $T[v]$, except for the root, perform spur-
 764 reduction($x, \text{parent}(x)$). All these operations satisfy spur-reduction’s constraints. Initially, every
 765 path through the node x has an edge in the segment $x \text{parent}(x)$, by (I2). The BFS traversal
 766 ensures that this property still holds when the algorithm performs spur-reduction($x, \text{parent}(x)$). \square

767 Note that for every path $[t_1; u; t_2]$, both t_1 and t_2 are in $T[v]$ (cf. Remark 2) and path $[t_1; t_2]$ is
 768 uniquely defined by (I1). However, a maximal path in $C(v)$ that contains $[t_1; t_2]$ violates (I2), and if
 769 $t_1 = t_2$ is a leaf in $C(v)$, then it forms a spur that may violate (I3). We proceed with a sequence of
 770 “repair” steps to restore them, after which the total number of benchmark vertices decreases by at
 771 least $|\mathcal{S}|$. The following three steps restore (I2) and (I3) when t_1 and t_2 are in ancestor-descendent
 772 relation, that is, $\text{lca}(t_1, t_2) \in \{t_1, t_2\}$. Let $\min(t_1, t_2)$ denote the node in $\{t_1, t_2\}$ farther from the
 773 root.

774 For every path $[t_1; t_2] \in \mathcal{S}$, do

- 775 1. If $\text{lca}(t_1, t_2) \in \{t_1, t_2\}$ and $t_1 \neq t_2$, then replace $[t_1; t_2]$ with $[\min(t_1, t_2)]$.
- 776 2. If $t_1 = t_2$ and t_1 is not a leaf of $T[v]$ that has degree two in the image graph, then
 777 replace $[t_1; t_2]$ with $[t_1]$.
- 778 3. If $t_1 = t_2$ and t_1 is a leaf of $T[v]$ that has degree two in the image graph, then do:
 779 by (I3), node t_1 is adjacent to a unique node $z \notin G_{uv}$ and z is incident to a single
 780 segment in the cluster containing z . Subdivide such segment creating a new node
 781 z^* (added to the cluster containing z), and replace every path $[z, t_1, z]$ with $[z^*]$.
 782 See Figure 26(b)–(c) for an example.




Figure 26: (a) Node u contains spurs. (b) After eliminating spurs, $T[v]$ does not satisfy (I2). (c) The analogues of pin-extraction and V-shortcut.

783 These steps restore (I3) at all leaves, and (I2) for the affected paths $[t_1; t_2] \in \mathcal{S}$. Note that these
784 steps are ws-equivalent: Steps 1–2 do not modify the polygon (they change only the benchmarks);
785 and step 3 is analogous to `pin-extraction`(t_1, z).

786 We are left with paths $[t_1; t_2] \in \mathcal{S}$ where t_1 and t_2 are in different branches of $T[v]$. In this case,
787 we perform an elaborate version of the V-shortcut operation, that creates a new group. For every
788 node ℓ of $T[v]$, let \mathcal{S}_ℓ be the set of paths $[t_1; t_2] \in \mathcal{S}$ such that $\text{lca}(t_1, t_2) = \ell$. Consider every node
789 ℓ of $T[v]$ where $\mathcal{S}_\ell \neq \emptyset$ in a bottom-up traversal of $T[v]$; and create a new group $G_{\ell-\ell^+}$ as follows
790 (refer to Figure 26).

791 Let N^- (resp., N^+) be the set of nodes t_1 (resp., t_2) such that there is a path $[t_1; t_2] \in \mathcal{S}_\ell$, and
792 t_1 is in the left (resp., right) subtree of ℓ . Let $N = N^- \cup N^+$. Sort the nodes $t_1 \in N^-$ by $t_1^\#$, and
793 let t_{\min} be the minimum node; and similarly sort the nodes $t_2 \in N^+$ by $t_2^\#$, and let t_{\max} be the
794 maximum node. The following lemma shows that interior nodes of the path from t_{\min} to ℓ in $T[v]$
795 have no right branches, and the interior nodes of the path from t_{\max} to ℓ have no left branches.

796 **Lemma 12.** *If there is a path $[s; t] \in \mathcal{B} \setminus \mathcal{S}_\ell$ such that $t_{\min}^\# < t^b \leq t^\# < t_{\max}^b$, then it crosses some*
797 *path in \mathcal{S}_ℓ , hence P is not weakly simple.*

798 *Proof.* Let C be the path between t_{\min} and t_{\max} in $T[v]$. Refer to Figure 25(b). By the choice of ℓ
799 (in a bottom-up traversal of $T[v]$), we have $\mathcal{S}_{\ell'} = \emptyset$ for all descendants of ℓ . Path $[s; t]$ reaches C at
800 some interior node $t^* \in C$, and then continues to ℓ , and farther to $\text{parent}(\ell)$. If t^* is in a left (resp.,
801 right) subtree of ℓ , then $[s; t]$ crosses every path in \mathcal{S}_ℓ that starts at t_{\min} (resp., ends at t_{\max}). \square

802 We can find the set N' of nodes t such that $t_{\min}^\# < t^b \leq t^\# < t_{\max}^b$, in $O(|N'| + \log n)$ time,
803 by binary search in the list of leaves to find all the leaves between t_{\min} and t_{\max} , and by lowest
804 common ancestor queries to find nodes in N' . The algorithm reports that the input polygon is not
805 weakly simple and halts if some node in N' has a path satisfying the condition in Lemma 12. We
806 can now assume that $N' \subset N$. The nodes in N induce a binary tree, denoted $T[\ell]$, of size at most
807 $2|N|$: its nodes are all nodes in N and the lowest common ancestors of consecutive nodes in N^-
808 and N^+ respectively. Note that a segment of $T[\ell]$ might not correspond to a segment of $T[v]$ (see
809 Figure 26(b)). Denote by C^* the path between t_{\min} and t_{\max} in $T[\ell]$.

810 We now define the changes in the image graph. Every node $t \in N \setminus C^*$ is deleted from G_{uv} ,
811 and added to the new group. Create two nodes, ℓ^- and ℓ^+ , in $G_{\ell-\ell^+}$ sufficiently close to ℓ in the
812 wedge between the two children of ℓ , and connect them by a segment $\ell^-\ell^+$. Duplicate each node
813 $t \in C^* \setminus \{\ell\}$, by creating a node t' (added to $G_{\ell-\ell^+}$) sufficiently close to t , and add a segment tt' .
814 Subdivide every segment tt' with two new boundary nodes, t_{leaf} (added to $T[v]$) and t'_{leaf} (added
815 to $G_{\ell-\ell^+}$). The nodes t or t' might now have degree 4. Adjust the image graph so that the group
816 trees are binary. Finally partition the nodes in $G_{\ell-\ell^+}$ into two trees, $T[\ell^-]$ and $T[\ell^+]$, rooted at ℓ^-
817 and ℓ^+ , respectively.

818 We now define the changes in the polygon. Replace every path $[t; t_1] \in \mathcal{S}_\ell$, where $t \in C^*$, by
819 $[t'; t_1]$ if it is adjacent to a path $[t; t_2] \in \mathcal{S}_\ell$, i.e., replacing the path $[t_1; t; t_2]$ by $[t_1; t'; t_2]$. Otherwise,
820 replace $[t; t_1]$ by $[t_{\text{leaf}}; t'_{\text{leaf}}; t_1]$. Now we can build \mathcal{B}' as the set of benchmark-to-benchmark paths
821 $[t'_1; t'_2]$ where $t'_1, t'_2 \in G_{\ell-\ell^+}$ in $O(|\mathcal{B}'|)$ time.

822 To prove ws-equivalence, we consider the changes in the polygon. These amount to a sequence of
823 ws-equivalent primitives: a `node-split` at ℓ , a sequence of `node-splits` along the chain C from ℓ to t_{\min}
824 and t_{\max} , respectively, `subdivision` operations that create the new leaf nodes, and `merge` operations
825 at degree two nodes that no longer contain spurs. The creation of new groups takes $O(|\mathcal{S}_\ell| + \log n)$
826 time and $O(|\mathcal{S}_\ell|)$ paths in \mathcal{B} are removed or modified in G_{uv} . Thus the data structures for G_{uv} are

827 updated in $O(|\mathcal{S}_\ell| \log n)$ time. Overall, operation `spur-shortcut(u)` and the repair steps that follow
 828 take $O(|\mathcal{S}| \log n)$ time.

829 5.3 Splitting a group in two

830 In this section we describe step 2c of Algorithm `spur-elimination(P, \mathcal{G})`. Assume that G_{uv} satisfies
 831 invariants (I1)–(I4) and u contains no spur.

832 Denote the left and right child of u by u^- and u^+ , respectively. Let $\mathcal{B}^-, \mathcal{B}^+ \subset \mathcal{B}$, resp., be
 833 the set of benchmark-to-benchmark paths that contain u^- and u^+ . We split G_{uv} into two groups
 834 induced by \mathcal{B}^- and \mathcal{B}^+ , respectively. Refer to Figure 27.

835 It would be easy to compute the groups induced by \mathcal{B}^- and \mathcal{B}^+ in $O(|\mathcal{B}|)$ time. However, for
 836 an overall $O(n \log n)$ -time algorithm, we can afford $O(\min(|\mathcal{B}^-|, |\mathcal{B}^+|))$ time for the split operation,
 837 and an additional $O(\log n)$ time for each eliminated spur and each node that we split into two
 838 nonempty nodes. Without loss of generality, we may assume $|\mathcal{B}^-| \leq |\mathcal{B}^+|$. The group induced
 839 by $|\mathcal{B}^-|$ can be computed from scratch in $O(|\mathcal{B}^-|)$ time, and we construct the group for \mathcal{B}^+ by
 840 modifying G_{uv} , and updating the corresponding data structures.




Figure 27: Splitting group G_{uv} . (a) Changes in the image graph. (b) Changes in the polygon.

841 First, we find \mathcal{B}^- and \mathcal{B}^+ . Compute \mathcal{B}^- using the list of paths $[s; t] \in \mathcal{B}$ sorted by $s^\#$ or s^\flat .
 842 Since both lists naturally split into corresponding lists for \mathcal{B}^- and \mathcal{B}^+ , we can split these lists in
 843 $O(\min(|\mathcal{B}^-|, |\mathcal{B}^+|)) = O(|\mathcal{B}^-|)$ time. To construct the list of \mathcal{B}^+ sorted by $t^\#$ and t^\flat , we start with
 844 the corresponding lists for \mathcal{B} , and delete all elements of \mathcal{B}^- in $O(|\mathcal{B}^-|)$ time. To compute the lists
 845 sorted by $t^\#$ and t^\flat for \mathcal{B}^- , we shall first compute the subtree $T[v^-]$ induced by \mathcal{B}^- . However, we
 846 can already find the maximum $t^\#$ of a path $[s; t] \in \mathcal{B}^-$ in $O(|\mathcal{B}^-|)$ time.

847 Next, we test for crossings between the paths in \mathcal{B}^- and the paths in \mathcal{B}^+ . Let $t_-^\#$ be the
 848 maximum $t^\#$ of a path $[s; t] \in \mathcal{B}^-$, and t_+^\flat the minimum t^\flat of a path $[s; t] \in \mathcal{B}^+$. By Lemma 10,
 849 there is such a crossing if and only if $t_+^\flat < t_-^\#$, which can be determined in $O(1)$ time using our
 850 order-maintenance structures. If a crossing is detected, the algorithm reports that P is not weakly
 851 simple and halts.

852 Trees $T[u^-]$ and $T[u^+]$ are simple subtrees of $T[u]$; but splitting $T[v]$ is nontrivial. We use
 853 binary search in the Eulerian cycle of all leaves to find the rightmost leaf ℓ_0 in $T[v]$ such that
 854 $\mathcal{B}_{\ell_0} \cap \mathcal{B}^- \neq \emptyset$, if such a leaf exists, otherwise the leftmost leaf ℓ_0 in $T[v]$. Let $C_0 = [\ell_0; u]$. We do
 855 not compute the path C_0 explicitly, as it may contain more than $O(|\mathcal{B}^-|)$ nodes, but we can test
 856 whether a query node t of $T[v]$ is in C_0 in $O(1)$ time by checking whether $\text{lca}(\ell_0, t) = t$. Since the

857 paths in \mathcal{B}^- and \mathcal{B}^+ do not cross, all nodes of $T[v^-]$ are in or to the left of the chain C_0 , and all
858 common nodes of $T[v^-]$ and $T[v^+]$ are in C_0 . The image graph of $T[v^-]$ can be computed from
859 scratch using \mathcal{B}^- in $O(\mathcal{B}^-)$ time. Replace each node t of $T[v^-]$ that is in C_0 by a duplicate copy
860 t^- located sufficiently close to t , to the right of t . The tree $T[v^+]$ is computed from $T[v]$ by node
861 deletion and merge operations as follows. First delete all nodes that are in $T[v^-]$ but not in C_0 .
862 For every node t of $T[v^-]$ that lies in C_0 , if t has degree two in $T[v^-]$ and $\mathcal{B}_t^+ = \emptyset$, then it would
863 be a degree two node in $T[v^+]$ with no spurs, and so we can delete t by merging its two incident
864 segments. Let v^+ be the node not in $T[u^+]$ adjacent to u^+ . The resulting $T[v^+]$ becomes a tree
865 induced by \mathcal{B}^+ . It remains to resolve the connections between trees.

866 Let \mathcal{V}^0 denote the set of chains $[s_1; t; s_2]$ such that $[s_1; t] \in \mathcal{B}^-$ and $[t; s_2] \in \mathcal{B}^+$. The spurs at t
867 on all chains $[s_1; t; s_2] \in \mathcal{V}^0$ will be eliminated (they will become adjacent leaves in the two resulting
868 groups). \mathcal{V}^0 can be found with a query for u in the interval tree. Let N^0 be the set of all nodes t
869 such that $[s_1; t; s_2] \in \mathcal{V}^0$. Each node $t \in N^0$ is in C_0 and, therefore, has a copy t^- in $T[v^-]$. Create
870 a segment between t and t^- , and subdivide the segment t^-t with two new nodes t_{leaf}^- and t_{leaf} in
871 $T[v^-]$ and $T[v^+]$, respectively. The degree of nodes t or t^- might increase to 4; and so we adjust
872 the image graphs so that both trees are binary. The image graph is now split into groups $G_{u^-v^-}$
873 and $G_{u^+v^+}$.

874 We next define the changes in the polygon. Replace every chain $[s_1; t; s_2] \in \mathcal{V}^0$ with a new chain
875 $[s_1; t_{\text{leaf}}^-; t_{\text{leaf}}; s_2]$, while also replacing the corresponding paths in the lists \mathcal{B}^- and \mathcal{B}^+ in $O(|\mathcal{V}^0|)$
876 time. In the sorted lists for \mathcal{B}^- and \mathcal{B}^+ , this is done by deletions and reinsertions. Note that all
877 leaves t_{leaf}^- (resp., t_{leaf}) are at the end (resp., beginning) of the Euler tour of $T[v^-]$ (resp., $T[v^+]$),
878 so deletions can be performed in $O(|\mathcal{V}^0|)$ time; and insertions take $O(|\mathcal{V}^0| \log n)$ time.

879 The changes in the polygon are equivalent to a sequence of ws-equivalent primitives: a **node-split**
880 operation at u , followed by a sequence of **node-splits** along the chain C_0 from ℓ_0 to u , and subdivision
881 operations that create the new leaf nodes between the two groups. The interval tree is updated
882 by deleting the intervals that contain u , and the query time remains the same output-sensitive
883 $O(\log n + k)$. Consequently, we can split G_{uv} in $O(\min(|\mathcal{B}^-|, |\mathcal{B}^+|) + |\mathcal{V}^0| \log n + \log n)$ time.

884 5.4 Analysis of the spur-elimination algorithm

885 **Lemma 13.** *Given m benchmark vertices, spur-elimination(P, \mathcal{G}) takes $O(m \log m)$ time.*

886 *Proof.* Let σ be the number of spurs, β the number of benchmark vertices at the leaves of clusters,
887 and let $\phi = 2\sigma + \beta$. Initially, $\phi = O(m)$ by (I1). All operations in **spur-elimination** monotonically
888 decrease both σ and ϕ . Step 2b decreases ϕ by the number of spurs at u , and steps 2a and 2c both
889 maintain ϕ . In particular, Step 2c converts some spurs into pairs of adjacent benchmark vertices at
890 leaves. Consequently, the number of benchmark vertices remains $O(m)$ throughout the algorithm.

891 Step 1 creates data structures for new groups: For a group containing m benchmarks, all
892 supporting data structures can be computed in $O(m)$ time, that is, in $O(1)$ time per benchmark.
893 A new benchmark v appears in a group when (i) a benchmark is extracted into an adjacent group,
894 or (ii) a group of size m is split and v is part of the smaller group of size at most $m/2$. Extraction
895 strictly decreases ϕ , so it occurs $O(m)$ times. The total number of benchmarks that are either
896 present initially or created by extraction is $O(m)$. Each of these benchmarks can move into a group
897 of half-size $O(\log m)$ times. Consequently, there are $O(m \log m)$ new benchmarks overall, and the
898 time spent on all instances of Steps 1 is $O(m \log m)$.

899 Step 2a removes an interior node of degree two; the update of supporting data structures takes

900 $O(\log m)$ time. Interior nodes are created only when they contain a spur, so at most $O(m)$ interior
901 nodes are ever created, and all instances of Step 2a take $O(m \log m)$ time. Step 2b eliminates
902 $|\mathcal{S}|$ spurs in $O(|\mathcal{S}| \log m)$ time. Eventually, all spurs are eliminated, thus all instances of Step 2b
903 take $O(m \log m)$ time. Step 2c takes $O(\min(|\mathcal{B}^-|, |\mathcal{B}^+|) + |\mathcal{V}^0| \log m + \log m)$ time. By a standard
904 heavy-path decomposition argument, the terms $\min(|\mathcal{B}^-|, |\mathcal{B}^+|)$ contribute $O(m \log m)$ time. Every
905 chain in \mathcal{V}^0 corresponds to a spur that is destroyed in a step 2c (and no new spurs are created in
906 step 2c), therefore the terms $O(|\mathcal{V}^0| \log m)$ sum to $O(m \log m)$ over the course of the algorithm.
907 Since every execution of step 2c increases the number of groups by one, and this step is repeated
908 $O(m)$ times, the $\log m$ terms sum to $O(m \log m)$ in the entire algorithm. \square

909 Algorithm `spur-elimination`(P, \mathcal{G}) returns a polygon P' , a set \mathcal{G}' of groups, and a set \mathcal{B}' of
910 benchmark-to-benchmark paths, each of which connects two leaves in two different clusters of a
911 group. In each group G_{uv} , the trees $T[u]$ and $T[v]$ have no interior nodes, thus G_{uv} consists of two
912 single-node clusters $C(u) = \{u\}$ and $C(v) = \{v\}$, connected by a single edge uv . Consequently, the
913 image graph is 2-regular. We can now decide whether P' is weakly simple in $O(n)$ time similarly
914 to [6, Section 3.3]. The polygon P' is weakly simple if and only if the image graph is connected
915 and each group contains precisely one benchmark-to-benchmark path. These properties can be
916 verified by a simple traversal of the image graph and P' in $O(n)$ time. This completes the proof of
917 Theorem 1.

918 6 Perturbing weakly simple polygons into simple polygons

919 In Sections 3–5, we have presented an algorithm that decides, in $O(n \log n)$ time, whether a given
920 n -gon P is weakly simple. If P is weakly simple, then for every $\varepsilon > 0$ it can be perturbed into a
921 simple polygon by moving each vertex a distance at most ε . In this section we show how to find,
922 for any $\varepsilon > 0$, a simple polygon Q with $2n$ vertices such that $\text{dist}_F(P, Q) < \varepsilon$. Let P' and P'' be
923 the polygons obtained after the bar-simplification and spur-elimination phases of the algorithm,
924 respectively. P'' has $O(n)$ vertices, none of which is a fork or a spur. Using the results in [6,
925 Section 3], we can construct a simple polygon $Q'' \in \Phi(P'')$ in $O(n)$ time. In this section, we show
926 that we can reverse the sequence of operations in $O(n \log n)$ time and perturb P as well into a
927 simple polygon $Q \in \Phi(P)$.

928 **Combinatorial representation by bar-signatures.** A perturbation of a weakly simple polygon
929 has a combinatorial representation, called a signature, which consists of total orders of the overlap-
930 ping edges in all segments of the image graph (cf. Section 2). In the absence of forks, every edge lies
931 in a segment, and the size of such a signature is $O(n)$. However, the signature may have size $\Theta(n^2)$
932 in the presence of forks. When our algorithm eliminates forks from a polygon, it may create $\Theta(n^2)$
933 dummy vertices and edges, which would again lead to a signature of size $\Theta(n^2)$. For reversing the
934 operations of the algorithm in Sections 3–5, we introduce a new combinatorial representation of
935 size $O(n)$ that maintains the total order of the edges in each bar that are outside of clusters.

936 For $n \geq 3$, let $P = (p_0, \dots, p_{n-1})$ be a weakly simple polygon with image graph G . Assume that
937 the sober nodes of G are partitioned into a set \mathcal{C} of disjoint clusters satisfying invariants (I1)–(I4)
938 such that every bar is either entirely in a cluster or outside of all clusters. Let $Q = (p'_0, \dots, p'_{n-1})$
939 be a simple polygon such that $|p_i, p'_i| < \varepsilon_0 = \varepsilon_0(P)$ for all $i = 0, \dots, n-1$. We may assume that G
940 has no vertical segments (so that the above-below relationship is defined between disjoint segments
941 parallel to a bar). In each segment uv of G outside of clusters, the above-below relationship yields

942 a total ordering over the edges of Q that contain uv . For each bar b outside of clusters, the total
 943 orders of the segments along b are consistent (since the above-below relationship between two edges
 944 is the same in every corridor). Consequently, the transitive closure of these total orders is a partial
 945 order over all edges in b . Consider a linear extension of such a partial order. The collection of
 946 these total orders for all bars in P is a *bar-signature* of Q . Since the linear extensions need not be
 947 unique, a polygon $Q \in \Phi(P)$ may have several bar-signatures.

948 Given a bar-signature of a perturbation of P , we can (re)construct an approximate simple
 949 polygon Q' as follows; refer to Figure 28. For every bar $b = uv$ of G outside of clusters, let the
 950 *volume* $\text{vol}(uv)$ be the number of edges of P that lie on b . Place $\text{vol}(uv)$ parallel line segments,
 951 called *lanes*, between ∂D_u and ∂D_v in the region U_ε , ordered from bottom to top (the lanes contain
 952 the edges of Q'). For the i -th edge pq in the total order of b , let the corresponding edge in Q' be the
 953 shortest edge connecting ∂D_p and ∂D_q in the i -th lane. For each cluster $C(u)$, denote by $R(u)$ the
 954 union of all disks D_v , $v \in C(u)$, and all corridors between nodes in $C(u)$. If $C(u)$ contains only the
 955 node u , then $R(u) = D_u$, but $R(u)$ is always simply connected since $C(u)$ induces a tree $T[u]$. For
 956 each cluster $C(u)$, construct a noncrossing polyline matching, between the endpoints of the edges
 957 in $\partial R(u)$, that connects the endpoints corresponding to a maximal subpath in $T[u]$. The edges in
 958 the lanes and the perfect matchings in the regions $R(u)$ produce a polygon Q' . If the Euclidean
 959 diameter of each region $R(u)$ is at most δ , then the Fréchet distance between P and Q' is at most
 960 $\varepsilon + \delta$. Denote by $\Psi(P)$ the set of all simple polygons that can be constructed in this manner from
 961 a bar-signature for some ε , $0 < \varepsilon < \varepsilon_0$.




Figure 28: Construction of a simple polygon $Q' \in \Psi(P)$ from a bar-signature. Left: Bar uv of a simple polygon obtained from an order compatible with the polygon shown in Figure 2(c). Right: maximal paths of Q and Q' inside clusters.

962 **Spur elimination.** If a given n -gon is weakly simple, our decision algorithm computes a polygon
 963 P'' , which is ws-equivalent to P and represented implicitly by a cyclic sequence of benchmark nodes.
 964 Specifically, P'' is represented by an image graph G'' , a set \mathcal{G}'' of groups, a set \mathcal{B}'' of benchmark-
 965 to-benchmark paths, and for every group $G_{uv} \in \mathcal{G}''$, a linear order of the paths in \mathcal{B}'' that cross the
 966 corridor N_{uv} between D_u and D_v . Consequently, the decision algorithm provides a bar-signature
 967 for the weakly simple polygon P'' .

968 We show that, by reversing the steps of Algorithm `spur-elimination`(P', \mathcal{G}'), we can compute a
 969 bar-signature of P' in $O(n \log n)$ time. If a group G_{uv} has been split in some step 2c (cf. Section 5.3),
 970 we can construct an ordering of the benchmark-to-benchmark paths of G_{uv} by concatenating the
 971 orders of \mathcal{B}^- and \mathcal{B}^+ (the sets of benchmark-to-benchmark paths of the resulting two groups).

972 If G_{uv} had spurs eliminated from u in some step 2b (cf. Section 5.2), we reverse each of the steps
 973 in the following manner. Recall that if a new group $G_{\ell-\ell+}$ was created, then every path $[t'_1; t'_2] \in \mathcal{B}'$
 974 in that group was created from a concatenation of two paths $[t_1; u]$ and $[t_2; u]$. Use the ordering
 975 of the paths in \mathcal{B}' to insert the paths $[t_1; u]$ and $[t_2; u]$ into the ordering of \mathcal{B} so that they form
 976 nested spurs, i.e., if $[t'_1; t'_2]$ is the topmost edge in \mathcal{B}' , $[t_1; u]$ (resp. $[t_2; u]$) should be the leftmost
 977 (resp., rightmost) path (without loss of generality, we use the orientation of Figure 26). Identify
 978 the leftmost path in the segment that connects ℓ and its right child and place all nested paths that
 979 created $G_{\ell-\ell+}$ immediately to its left.

980 If one or more spurs were created at a node z in an adjacent group, we can find the position
 981 of the edges incident to each spur in the ordering of the adjacent group. Using this order, we can
 982 identify the first path in G_{uv} to the right of the edges incident to $[z]$. Then, immediately to the
 983 left of such a path, we can place the paths $[t_1; u; t_1]$ that generated the spurs at z . The relative
 984 order of these paths is the same as the one obtained by reversing a *spur-reduction*, described in the
 985 proof of Lemma 4, and therefore produces a simple polygon. If a path $[t_1; u; t_2]$ is simplified to $[t_1]$
 986 (Step 1 with $\min(t_1, t_2) = t_1$ without loss of generality; or Step 2), we can proceed analogously to
 987 the reversal of a *crimp-reduction* (cf. Lemma 1) from a path $[t_1; u; t_2; s]$ to $[t_1; s]$. Identify the path
 988 $[t_1; s]$ in the ordering of \mathcal{B} and replace it with the paths $[t_1; u]$, $[u; t_2]$, and $[t_2; s]$ in this order.

989 **Bar simplification.** The bar-signature determines all segments between adjacent clusters. Using
 990 these orders, we can reverse the operation *pin-extraction*(u, v) assigning the same order for the
 991 edges in uv as the order of its adjacent benchmark-to-benchmark paths. *V-shortcut* is also trivially
 992 reversible by concatenating the order of segments that get merged.

993 Updating the bar-signature when we reverse an *L-shortcut* operation is a bit more challenging.
 994 Determining the edge order in segments vw and vu_1 can be trivially done by just concatenating the
 995 order of merged segments. But phase (1) introduces a crimp in some cross-chains, and the reverse
 996 operation, *crimp-reduction*, may require nontrivial reordering in the bar-signature. Suppose that
 997 P' is obtained from P after a *crimp-reduction*. The proof of Lemma 1 shows a straightforward way
 998 to obtain a bar-signature of a polygon in $\Psi(P)$ given a polygon in $\Psi(P')$. However, obtaining a
 999 bar-signature of $Q' \in \Psi(P')$ given $Q \in \Psi(P)$ requires identifying W_{top} and W_{bot} , which takes $O(n)$
 1000 time.

1001 In order to handle the reversal of phase (1) in $O(1)$ time, we divide the signature of each bar
 1002 into pieces. Recall that the *bar-simplification* algorithm does not eliminate any cross-chains from
 1003 D_b , and when *bar-simplification* terminates, only one-edge cross-chains remain in the interior of D_b .
 1004 Let K denote the set of cross-chains of D_b . The segments of the image graph that cross the ellipse
 1005 D_b , and the bar-signatures of these edges yield a linear order (from left to right) of K ; and the
 1006 cross-chains subdivide D_b into $|K| + 1$ regions. We maintain a linear order for the edges along the
 1007 bar in each such region (including the boundary of the region), and denote the set of these edges
 1008 in b by $E_1, \dots, E_{|K|+1}$.

1009 We reverse phases (2) and (3) of *L-shortcut*(v, TR) as follows (applying reflections for other
 1010 *L-shortcut* operations if necessary). Assign the new edges $[u_1, u_2]$ the highest lanes in the ordering
 1011 of the appropriate E_i , maintaining the relative order of affected paths. To reverse phase (1), first
 1012 notice that the three edges in the crimp $[u_1, u_2, u_1, u_2]$ are part of a cross-chain, consequently they
 1013 appear in two consecutive subsets E_i and E_{i+1} . In the ordering of the left (resp., right) subset,
 1014 assign the new edge $[u_1, u_2]$ to the highest (resp., lowest) position among the positions of the three
 1015 edges $[u_1, u_2]$.

1016 When all operations in the bar simplification algorithm have been reversed, we have to combine

1017 the linear orders of $E_1, \dots, E_{|K|+1}$ into a total order, a common linear extension of these orders.
 1018 The intersection of two edge sets, $E_i \cap E_j$ with $i < j$, is either disjoint or contains the edges of the
 1019 i -th cross-chain. The above-below relationship between the edges of each cross-chain is uniquely
 1020 determined by Lemma 2, and must be the same in each total order. Therefore, the union of the
 1021 total orders is a partial order for all edges in the bar. Since the ordering of each subset guarantees
 1022 that its paths can be realized without crossing, any linear extension of this partial order produces
 1023 a bar-signature of a simple polygon.

1024 **Preprocessing.** The cluster formation and new-bar-expansion consist of subdivision operations that
 1025 do not influence the order of edges that define the bar-signature. If an edge $[v, w]$ in a bar b is
 1026 subdivided into $[v, v', w]$, where $[v', w]$ is in D_b , we can assign $[v, w]$ to the same lane of $[v', w]$ in
 1027 the ordering of edges in b . The crimp-reduction operations can be reversed by making the three
 1028 edges that form a new crimp consecutive in the ordering, as in the proof of Lemma 1.

1029 We have shown how to maintain bar-signatures while reversing the operations of our algorithms,
 1030 in time proportional to those operations. For every $\varepsilon > 0$, the bar-signatures yield a perturbation
 1031 of a weakly simple polygon P into a simple polygon $Q \in \Phi(P)$ with $2n$ vertices, where each vertex
 1032 $[u]$ of P corresponds to two vertices of Q on the circle ∂D_u . This completes the proof of Theorem 2.

1033 7 Conclusion

1034 We presented an $O(n \log n)$ -time algorithm for deciding whether a polygon with n vertices is weakly
 1035 simple. Weak simplicity of polygons has a natural generalization for planar graphs [6, Appendix
 1036 D]. We can define the *weak embedding* for graphs in terms of Fréchet distance. A graph $H = (V, E)$
 1037 can be considered a 1-dimensional simplicial complex. A *drawing* of H is a continuous map of H
 1038 to \mathbb{R}^2 . The Fréchet distance between two drawings, P and Q , of H is defined as $\text{dist}_F(P, Q) =$
 1039 $\inf_{\phi: H \rightarrow H} \max_{x \in H} \text{dist}(P(\phi(x)), Q(x))$, where ϕ is an automorphism of H (a homeomorphism from
 1040 H to itself). Very recently, Fulek and Kynčl [13] gave a polynomial-time algorithm for deciding
 1041 whether a given drawing of a graph H is weakly simple, i.e., whether a straight-line drawing P of H
 1042 is within ε Fréchet distance from some embedding Q of H , for all $\varepsilon > 0$. Earlier, efficient algorithms
 1043 were known only in special cases: when the embedding is restricted to a given isotopy class (i.e.,
 1044 given combinatorial embedding) [12]; and when all n vertices are collinear and the isotopy class is
 1045 given [1].

1046 We can also generalize the problem to higher dimensions. A polyhedron can be described as
 1047 a map $\gamma : M \rightarrow \mathbb{R}^3$, where M is a 2-manifold without boundary. A simple polyhedron is an
 1048 injective function. A polyhedron P is weakly simple if there exists a simple polyhedron within ε
 1049 Fréchet distance from P for all $\varepsilon > 0$. This problem can be reduced to origami flat foldability. The
 1050 results of [5] imply that, given a convex polygon P and a piecewise isometric function $f : P \rightarrow \mathbb{R}^2$
 1051 (called *crease pattern*), it is NP-hard to decide if there exists an injective embedding of P in three
 1052 dimensions $\lambda : P \rightarrow \mathbb{R}^3$ within ε Fréchet distance from f for all $\varepsilon > 0$, i.e., if f is *flat foldable*.
 1053 Given P and f , we can construct a continuous function $g : \mathbb{S}^2 \rightarrow P$ mapping each hemisphere of \mathbb{S}^2
 1054 to P (for a point $x \in P$, the inverse image $g^{-1}(x)$ is a set of two points in opposite hemispheres of
 1055 \mathbb{S}^2). Then, the polyhedron $\gamma = g \circ f$ is weakly simple if and only if f is flat foldable. Therefore, it
 1056 is also NP-hard to decide whether a polyhedron is weakly simple.

1057 Finally it is an open problem to find a linear-time algorithm for recognizing weakly simple
 1058 polygons. Chang et al. [6] conjectured that this is possible in the absence of spurs and forks.

1059 **Acknowledgements.** Research by Akitaya, Aloupis, and Tóth was supported in part by the
1060 NSF awards CCF-1422311 and CCF-1423615. Akitaya was supported by the Science Without
1061 Borders program. Research by Erickson was supported in part by the NSF award CCF-1408763.
1062 We thank Anika Rounds and Diane Souvaine for many helpful conversations that contributed to
1063 the completion of this project. We thank the anonymous referees for many useful comments and
1064 suggestions.

1065 References

- 1066 [1] Zachary Abel, Erik D. Demaine, Martin L. Demaine, David Eppstein, Anna Lubiw, and Ryuhei
1067 Uehara, Flat foldings of plane graphs with prescribed angles and edge lengths, in *Proc. 22nd*
1068 *Symposium on Graph Drawing*, LNCS 8871, Springer, 2014, pp. 272–283.
- 1069 [2] Esther M. Arkin, Michael A. Bender, Erik D. Demaine, Martin L. Demaine, Joseph S.B.
1070 Mitchell, Saurabh Sethia, and Steven S. Skiena, When can you fold a map?, *Computational*
1071 *Geometry: Theory and Applications* **29** (2004), 23–46.
- 1072 [3] Michael A. Bender, Richard Cole, Erik D. Demaine, Martin Farach-Colton, and Jack Zito. Two
1073 simplified algorithms for maintaining order in a list, *Proc. 10th Annual European Symposium*
1074 *on Algorithms*, LNCS 2461, Springer, 2002, pp. 152–164.
- 1075 [4] Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars, *Computational Geom-*
1076 *etry: Algorithms and Applications*, third edition, Springer, Berlin, 2008.
- 1077 [5] Marshall Bern and Barry Hayes, The complexity of flat origami, in *Proc. 7th ACM-SIAM*
1078 *Symposium on Discrete Algorithms*, SIAM, 1996, pp. 175–183.
- 1079 [6] Hsien-Chih Chang, Jeff Erickson, and Chao Xu, Detecting weakly simple polygons, in *Proc.*
1080 *26th ACM-SIAM Symposium on Discrete Algorithm*, SIAM, 2015, pp. 1655–1670.
- 1081 [7] Bernard Chazelle, Triangulating a simple polygon in linear time, *Discrete & Computational*
1082 *Geometry* **6** (1991), 485–524.
- 1083 [8] Richard Cole and Ramesh Hariharan, Dynamic LCA queries on trees, *SIAM J. Comput.* **34**(4)
1084 (2005), 894–923.
- 1085 [9] Robert Connelly, Erik D. Demaine, and Günter Rote, Infinitesimally locked self-touching link-
1086 ages with applications to locked trees, in *Physical Knots: Knotting, Linking, and Folding of*
1087 *Geometric Objects in \mathbb{R}^3* , American Mathematical Society, Providence, RI, 2002, pages 287–311.
- 1088 [10] Pier Francesco Cortese, Giuseppe Di Battista, Maurizio Patrignani, and Maurizio Pizzonia,
1089 On embedding a cycle in a plane graph, *Discrete Mathematics* **309**(7) (2009), 1856–1869.
- 1090 [11] Andrea Francke and Csaba D. Tóth, A census of plane graphs with polyline edges, *SIAM J.*
1091 *Discrete Math.* **31**(2) (2017), 1174–1195.
- 1092 [12] Radoslav Fulek, Embedding graphs into embedded graphs, preprint, [arXiv:1608.02087](https://arxiv.org/abs/1608.02087), 2016.
- 1093 [13] Radoslav Fulek and Jan Kynčl, Hanani-Tutte for approximating maps of graphs, preprint,
1094 [arXiv:1705.05243](https://arxiv.org/abs/1705.05243), 2017.
- 1095 [14] Branko Grünbaum, Polygons: Meister was right and Poincot was wrong but prevailed, *Beiträge*
1096 *zur Algebra und Geometrie* **53**(1) (2012), 57–71.
- 1097 [15] Piotr Minc, Embedding of simplicial arcs into the plane, *Topology Proceedings* **22** (1997),
1098 305–340.

- 1099 [16] Ares Ribó Mor, *Realization and Counting Problems for Planar Structures: Trees and Linkages,*
1100 *Polytopes and Polyominoes*, Ph.D. thesis, Freie Universität Berlin, 2006.
- 1101 [17] Michael Ian Shamos and Dan Hoey, Geometric intersection problems, in *Proc. 17th IEEE*
1102 *Symposium on Foundations of Computer Science*, 1976, pp. 208–215.
- 1103 [18] Mikhail Skopenkov, On approximability by embeddings of cycles in the plane. *Topology and*
1104 *its Applications* **134** (2003), 1–22.
- 1105 [19] Daniel D. Sleator and Paul F. Dietz, Two algorithms for maintaining order in a list, in *Proc.*
1106 *19th ACM Symposium on Theory of Computing*, 1987, pp. 365–372. Full version in Tech. Rep.
1107 CMU-CS-88-113, Carnegie Mellon University, Pittsburgh, PA, 1988.