

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 [5]. Weakly simple polygons are required as input for several geometric algorithms. As such, recognizing simple or weakly simple polygons is a fundamental problem.

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 (Fig. 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.

Several alternative definitions have been proposed for weakly simple polygons, formalizing the intuition that such polygons do not self-cross. Some of these definitions were unnecessarily restrictive or incorrect; see [5] for a detailed discussion and five equivalent definitions for weak simplicity of a polygon. Among others, a result by Ribó Mor [14][Theorem 3.1] implies an equivalent definition in terms of Fréchet distance, in which a polygon is perturbed into a simple closed curve (see Section 2). This definition is particularly useful for recognizing weakly simple polygons, since it

*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.

[‡]Department of Computer Science, University of Illinois, Urbana-Champaign, IL

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

21 allows transforming edges into polylines (by subdividing the edges with Steiner points which may
 22 be perturbed). With suitable Steiner points, the perturbation of a vertex incurs only local changes
 23 (in other words, we do not need to worry about stretchability of the perturbed configuration).

24 We can decide whether an n -gon in the plane is simple in $O(n \log n)$ time by a sweepline
 25 algorithm [15]. Chazelle’s polygon triangulation algorithm also recognizes simple polygons (in
 26 $O(n)$ time), because it only produces a triangulation if the input is simple [6]. Recognizing weakly
 27 simple polygons, however, is more subtle. Skopenkov [16] gave a combinatorial characterization of
 28 the topological obstructions to weak simplicity in terms of line graphs. Cortese et al. [9] gave an
 29 $O(n^3)$ -time algorithm to recognize weakly simple n -gons. Chang et al. [5] improved the running
 30 time to $O(n^2 \log n)$ in general; and to $O(n \log n)$ in several special cases. They identified two
 31 features that are difficult to handle: A *spur* is a vertex whose incident edges overlap, and a *fork* is
 32 a vertex that lies in the interior of an edge (a vertex may be both a fork and a spur). They gave
 33 an easy algorithm for polygons with neither forks nor spurs, and a more involved one for polygons
 34 with spurs but no forks, both running in $O(n \log n)$ time. In the presence of both forks and spurs,
 35 they presented an $O(n^2 \log n)$ time algorithm that eliminates forks by subdividing all edges that
 36 contain vertices in their interiors, potentially creating a quadratic number of vertices.

37 We show how to manage both forks and spurs efficiently, while building on ideas from [5, 9] and
 38 from Arkin et al. [2], and obtain the following main results.

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

41 **Theorem 2.** *Given a weakly simple polygon P with n vertices and an $\varepsilon > 0$, a simple polygon with*
 42 *$2n$ vertices within Fréchet distance ε from P can be computed in $O(n \log n)$ time.*

43 Our decision algorithm is detailed in Sections 3–5. It consists of three phases, simplifying the
 44 input polygon by a sequence of reduction steps. First, the *preprocessing* phase applies known
 45 methods such as *crimp reductions* and *node expansions* (Section 3). Second, the *bar simplification*
 46 phase successively eliminates all forks (Section 4). Third, the *spur elimination* phase eliminates all
 47 spurs (Section 5). When neither forks nor spurs are present, we can decide weak simplicity in $O(n)$
 48 time [9]. Finally, by reversing the sequence of operations, we can also perturb any weakly simple
 49 polygon into a simple polygon in $O(n \log n)$ time (Section 6).

50 2 Preliminaries

51 In this section, we review previously established definitions and known methods from [5] and [9].

52 **Polygons and weak simplicity.** An *arc* in \mathbb{R}^2 is a continuous function $\gamma : [0, 1] \rightarrow \mathbb{R}^2$. A *closed*
 53 *curve* is a continuous function (map) $\gamma : \mathbb{S}^1 \rightarrow \mathbb{R}^2$. A closed curve γ is *simple* (also known as a
 54 *Jordan curve*) if it is injective. A (*simple*) *polygon* is the image of a piecewise linear (*simple*) closed
 55 curve. Thus a polygon P can be represented by a cyclic sequence of points (p_0, \dots, p_{n-1}) , called
 56 *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 cyclic
 57 order. Similarly, a *polygonal chain* (alternatively, *path*) is the image of a piecewise linear arc, and
 58 can be represented by a sequence of points $[p_0, \dots, p_{n-1}]$.

59 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
 60 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
 61 to work with because a small perturbation of a vertex modifies the two incident edges, which may

62 be long, and the effect of a perturbation is not localized. Combining earlier results from [8], [9], and
63 [14][Theorem 3.1], an equivalent definition was formulated by Chang et al. [5] in terms of Fréchet
64 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
65 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
66 between two closed curves. The curve γ' can approximate an edge of the polygon by a polyline,
67 and any perturbation of a vertex can be restricted to a small neighborhood. With this definition,
68 recognizing weakly simple polygons becomes a combinatorial problem, as explained below. Note
69 that in topology, the broader question of *isotopic embeddability* has been considered [13, 16]: Given
70 a continuous map $f : A \rightarrow \mathbb{R}^d$ for a simplicial complex A , is it isotopic to some *injective* continuous
71 map (i.e., (*embedding*) $g : A \rightarrow \mathbb{R}^d$?

72 **Bar decomposition and image graph.** Two edges of a polygon P *cross* if their interiors intersect
73 at precisely one point; we call this an *edge crossing*. Weakly simple polygons cannot have edge
74 crossings. In the following, we assume that such crossings have been ruled out. Two edges of P
75 *overlap* if their intersection is a (nondegenerate) line segment. The transitive closure of the overlap
76 relation is an equivalence relation on the edges of P ; see Fig. 2(a) where equivalence classes are
77 represented by purple regions. The union of all edges in an equivalence class is called a *bar*.¹ All
78 bars of a polygon can be computed in $O(n \log n)$ time [5]. The bars are line segments that are
79 pairwise noncrossing and nonoverlapping. There are at most n bars, since each bar is the union of
80 some edges of P .

81 The vertices and bars of P define a planar straight-line graph G , called the *image graph* of P .
82 We call the vertices and edges of G *nodes* and *segments*¹ to distinguish them from the vertices and
83 edges of P . Every node that is not in the interior of a bar is called *sober*¹. The set of nodes in G
84 is $\{p_0, \dots, p_{n-1}\}$ (note that P may have repeated vertices that correspond to the same node); two
85 nodes are connected by a segment in G if they are consecutive nodes along a bar; see Fig. 2(b).
86 Hence G has $O(n)$ nodes and segments, and it can be computed in $O(n \log n)$ time [5]. Note,
87 however, that up to $O(n)$ edges of P may pass through a node of G , and there may be $O(n^2)$
88 edge-node pairs such that an edge of P passes through a node of G . An $O(n \log n)$ -time algorithm
89 cannot afford to compute these pairs explicitly.




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 .

90 **Operations.** We use certain elementary operations that successively modify a polygon and ul-
91 timately eliminate forks and spurs. An operation that produces a weakly simple polygon iff it is

¹We adopt terminology from [5].

92 performed on a weakly simple polygon is called *ws-equivalent*. Several such operations are already
 93 known (e.g., crimp reduction, node expansion, bar expansion). We shall use these and introduce
 94 several new operations in Sections 3.3–5.

95 **Combinatorial characterization of weak simplicity.** To show that an operation is ws-
 96 equivalent, it suffices to provide suitable simple ε -perturbations for all $\varepsilon > 0$. We use a combi-
 97 natorial representation of an ε -perturbation (independent of ε or any specific embedding). When a
 98 weakly simple polygon P is perturbed into simple polygon, overlapping edges in P are perturbed
 99 into interior-disjoint near-parallel edges, which define an ordering. It turns out that these orderings
 100 over all segments of the image graph are sufficient to encode an ε -perturbation and to (re)construct
 101 an ε -perturbation.

102 We rely on the notion of “strip system” introduced in [5][Appendix B]. Similar concepts have
 103 previously been used in [8, 9, 10, 13, 16]. Let P be a polygon and G its image graph. Without loss
 104 of generality, we assume that no bar is vertical. For every $\varepsilon > 0$, the ε -strip-system of P consists
 105 of the following regions:

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

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

113 We say that a polygon is *in the ε -strip-system* of P if its edges alternate between an edge that
 114 connects the boundary of two disks D_u and D_v and whose interior is contained in N_{uv} ; and an
 115 edge between two points on the boundary of a disk. See Fig. 2(c) for an example, where the edges
 116 within the disk D_u are drawn with circular arcs for clarity. Let $\Phi(P)$ be the set of simple polygons
 117 in the ε_0 -strip-system of P that crosses the disks and corridors in the same order as P traverses the
 118 corresponding nodes and segments of G . It is clear that every $Q \in \Phi(P)$ is within Fréchet distance
 119 ε_0 from P . By [5][Theorem B.2], P is weakly simple iff $\Phi(P) \neq \emptyset$.

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

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

133 **Remark 1.** The construction above has two immediate consequences: (1) To prove weak simplicity,
 134 it is enough to find a signature that defines a simple perturbation; that is the signature witnesses

135 weak simplicity (independent of the value of ε). (2) Weak simplicity of a polygon depends only on
 136 the *combinatorial embedding* of the image graph G (i.e., the counterclockwise order of edges incident
 137 to each vertex), as long as G is a plane graph. Consequently, when an operation modifies the image
 138 graph, it is enough to maintain the combinatorial embedding of G (the precise coordinates of the
 139 nodes do not matter).

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

144 **The number of combinatorially different perturbations.** In the absence of spurs, a polygon
 145 P determines a unique noncrossing perfect matching in each disk D_u , hence a unique noncrossing
 146 2-regular graph in the ε_0 -strip-system of P [5][Section 3.3]. Consequently, to decide whether P is
 147 weakly simple it is enough to check whether this graph is connected. The uniqueness no longer
 148 holds in the presence of spurs. In fact, it is not difficult to construct weakly simple n -gons that
 149 admit $2^{\Theta(n)}$ perturbations into simple polygons that are combinatorially different (i.e., have different
 150 bar-signatures); see Fig. 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.

151 3 Preprocessing

152 By a standard line sweep [15], we can detect and halt if any two edges properly cross. We then
 153 simplify the polygon, using some known steps from [2, 5], and some new ones. All of this takes
 154 $O(n \log n)$ time.

155 3.1 Crimp reduction

156 Arkin et al. [2] gave an $O(n)$ -time algorithm for recognizing weakly simple n -gons in the special case
 157 where all edges are collinear (in the context of flat foldability of a polygonal linkage). They define
 158 the ws-equivalent *crimp-reduction* operation. A *crimp* is a chain of three consecutive collinear edges
 159 $[a, b, c, d]$ such that both the first edge $[a, b]$ and the last edge $[c, d]$ contain the middle edge $[b, c]$
 160 (the containment need not be strict). The operation $\text{crimp-reduction}(a, b, c, d)$ replaces the crimp
 161 $[a, b, c, d]$ with edge $[a, d]$; see Fig. 4.




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

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

163 *Proof.* Let P_1 and P_2 be two polygons such that P_2 is obtained from P_1 by the operation **crimp-**
 164 **reduction**(a, b, c, d). Without loss of generality, assume that ad is horizontal with a on the left and
 165 d on the right.

166 First assume that P_1 is weakly simple. Then there exists a simple polygon $Q_1 \in \Phi(P_1)$. We
 167 modify Q_1 to obtain a simple polygon $Q_2 \in \Phi(P_2)$. Without loss of generality, assume that edge
 168 $[a, b]$ is above $[b, c]$ (consequently, $[c, d]$ is below $[b, c]$) in Q_1 . The modification involves the perfect
 169 matchings at the disks D_b and D_c , and all disks and corridors along the line segment bc . Denote by
 170 W_{top} the set of maximal paths that lie in the convex hull of $D_b \cup D_c$, below $[a, b]$ and above $[b, c]$;
 171 similarly, let W_{bot} be the set of maximal paths that lie in the convex hull of $D_b \cup D_c$, below $[b, c]$
 172 and above $[c, d]$. We proceed in two steps; refer to Fig. 5. First, replace the path $[a, b, c, d]$ with
 173 the path $[a, c, b, d]$ such that the new edge $[a, c]$ replaces old $[a, b]$ in the edge ordering of segment
 174 ac , new $[c, b]$ replaces old $[b, c]$ in segments contained in bc , and finally new $[b, d]$ replaces $[c, d]$ in
 175 bd . Second, exchange W_{top} and W_{bot} such that the above-below order within each set of paths
 176 remains the same. Since the above-below order within W_{top} and W_{bot} is preserved, and the paths
 177 in W_{top} (resp., W_{bot}) lie below (resp., above) the new path $[a, c, b, d]$, no edge crossings have been
 178 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 $[a, b, c, d]$.

179 Next assume that P_2 is weakly simple. Then, there exists a simple polygon $Q_2 \in \Phi(P_2)$. We
 180 modify Q_2 to obtain a simple polygon $Q_1 \in \Phi(P_1)$; refer to Fig. 6. Replace edge $[a, d]$ by $[a, b, c, d]$
 181 also replacing $[a, d]$ in the ordering of the affected segments by $[c, d]$, $[b, c]$, $[a, b]$, in this order. The
 182 new ordering produces a polygon Q_1 in the strip system of P . Because Q_2 is simple, by construction
 183 the new matchings do not interact with the preexisting edges in the disks. Hence, $Q_1 \in \Phi(P_1)$,
 184 which shows that P_1 is weakly simple. \square

185 Given a chain of two edges $[a, b, c]$ such that $[a, b]$ and $[b, c]$ are collinear but do not overlap, the
 186 **merge** operation replaces $[a, b, c]$ with a single edge $[a, c]$. The merge operation (as well as its inverse,
 187 **subdivision**) is ws-equivalent by the definition of weak simplicity in terms of Fréchet distance [5]. If
 188 we greedily apply **crimp-reductions** and **merge** operations, in linear time we obtain a polygon with
 189 the following two properties:

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

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

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

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

202 3.2 Node expansion

203 Compute the bar decomposition of P and its image graph G (defined in Section 2, see Fig. 2).
 204 For every sober node of the image graph, we perform the ws-equivalent **node-expansion** operation,
 205 described in [5][Section 3] (Cortese et al. [9] call this a *cluster expansion*). Let u be a sober node of
 206 the image graph. Let D_u be the disk centered at u with radius $\delta > 0$ sufficiently small so that D_u
 207 intersects only the segments incident to u . For each segment ux incident to u , create a new node
 208 u^x at the intersection point $ux \cap \partial D_u$. Then modify P by replacing each subpath $[x, u, y]$ passing
 209 through u by $[x, u^x, u^y, y]$; see Fig. 7. If a node expansion produces an edge crossing, report that
 P is not weakly simple.




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.

210

211 3.3 Bar expansion

212 Chang et al. [5][Section 4] define a bar expansion operation. In this paper, we will refer to it as
 213 **old-bar-expansion**. For a bar b of the image graph, draw a long and narrow ellipse D_b around the
 214 interior nodes of b , create subdivision vertices at the intersection of ∂D_b with the edges, and replace
 215 each maximal path in D_b by a straight-line edge. If b contains no spurs, **old-bar-expansion** is known
 216 to be ws-equivalent [5]. Otherwise, it can produce false positives, hence it is not ws-equivalent; see
 217 Fig. 8 for an example.

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

223 Similar to **old-bar-expansion**, the operation **new-bar-expansion** introduces subdivision vertices on
 224 ∂D_b , however we keep all interior vertices of a bar at their original positions. In Section 4, we apply
 225 a sequence of new operations to eliminate all vertices on b sequentially while creating new nodes in




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

226 the vicinity of D_b . Our bar expansion operation can be considered as a preprocessing step for this
 227 subroutine.

228 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
 229 intersection point $ux \cap \partial D_b$ and subdivide every edge $[u, x]$ to a path $[u, u^x, x]$. For each endpoint
 230 v of b , create two new nodes, v' and v'' , as follows. Node v is adjacent to a unique segment $vw \subset b$,
 231 where $w \in b \cap D_b$. Create a new node $v' \in \partial D_b$ sufficiently close to the intersection point $vw \cap \partial D_b$,
 232 but strictly above b ; and create a new node v'' in the interior of segment $vw \cap D_b$. Subdivide every
 233 edge $[v, y]$, where $y \in b$, into a path $[v, v', v'', y]$. Since the new-bar-expansion operation consists of
 234 only subdivisions (and slight perturbations of the edges passing through the end-segments of the
 235 bars), it is ws-equivalent.




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

236 **Crossing paths.** Apart from node-expansion and old-bar-expansion, none of our operations will
 237 create edge crossings. In some cases, our bar simplification algorithm (Section 4) will detect whether
 238 two subpaths cross. Crossings between overlapping paths is not easy to identify (see [5][Section 2]
 239 for a discussion). We rely on the following simple condition to detect some (but not all) crossings.

240 **Lemma 3.** *Let P be a weakly simple polygon parameterized by a curve $\gamma_1 : \mathbb{S}^1 \rightarrow \mathbb{R}^2$; and let
 241 $\gamma_2 : \mathbb{S}^1 \rightarrow \mathbb{R}^2$ be a closed Jordan curve that does not pass through any vertices of P and intersects
 242 every edge of P transversely. If $\gamma_1(\mathbb{S}^1)$ and $\gamma_2(\mathbb{S}^1)$ intersect in four distinct points, then the cyclic
 243 order of the intersection points along γ_1 and γ_2 is either the same or reverse.*

244 *Proof.* If P is weakly simple, then γ_1 can be perturbed to a closed Jordan curve γ'_1 with the same
 245 properties. Denote the intersection points by p_i , $i = 1, \dots, 4$, in counterclockwise order along
 246 γ_2 . By the Jordan curve theorem, $\mathbb{R}^2 \setminus \gamma_2(\mathbb{S})$ has two connected components, the interior and the
 247 exterior of γ_2 . Since γ'_1 and γ_2 intersect transversely, the intersection points partition γ'_1 into four
 248 Jordan arcs that lie alternately in the interior and exterior of γ_2 . The two Jordan arcs of γ'_1 lying
 249 in the interior (resp., exterior) of γ_2 cannot cross. Without loss of generality, the two arcs of γ'_1 in
 250 the interior of γ_2 connect the pairs $\{p_1, p_2\}$ and $\{p_3, p_4\}$. Then the two arcs of γ'_1 in the exterior of
 251 γ_2 must connect the pairs $\{p_2, p_3\}$ and $\{p_4, p_1\}$. That is, the cyclic order of the intersection points
 252 along γ'_1 (hence, along γ_1) is also (p_1, p_2, p_3, p_4) , as claimed. \square

Specifically, a weakly simple polygon cannot contain certain configurations.

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

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 this cyclic order about u_2 (node crossing; see Fig. 10(a)).
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 segments vu_2 and wu_3 lie in a closed halfplane bounded by this line (Fig. 10(b)).
3. $[u_1, u_2]$ and $[v_1, v_2, \dots, v_{k-1}, v_k]$ where $v_2, \dots, v_{k-1} \in \text{int}(u_1u_2)$, and segments v_1v_2 and $v_{k-1}v_k$ lie in two different halfplanes bounded by line containing u_1u_2 (Fig. 10(c)).

Proof. In all four cases, Lemma 3 with a suitable Jordan curve γ_2 completes the proof. In case 1, 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 case 3, let γ_2 be a small neighborhood of the convex hull of $\{v_2, \dots, v_{k-1}\}$. \square




Figure 10: Three pairs of incompatible paths. Nodes u_1 , u_2 , u_3 , and u_4 are collinear.

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

- *cross chain* if u_1^x and u_k^y are top and bottom nodes respectively;
- *top chain* (resp., *bottom chain*) if both u_1^x and u_k^y are top nodes (resp., bottom nodes);
- *pin* if $p = [u_1^x, u_1, u_1^x]$ (note that every pin is a top or a bottom chain);
- *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.

Let $\mathcal{Pin} \subset \mathcal{P}$ be the set of pins, and $\mathcal{V} \subset \mathcal{P}$ the set of V-chains. Let M_{cr} be the set of longest edges of *cross chains* in \mathcal{P} (by Lemma 2, each cross chain contributes one or two edges).

By Corollary 1, every weakly simple polygon has the following properties.

- (A3) Polygon P has no node-crossings (cf. Fig. 10(a)).
- (A4) No edge in M_{cr} lies in the interior of any other edge of P (cf. Fig. 10(c)).

We can test properties (A3)–(A4) in $O(n \log n)$ time at preprocessing: For each bar, sort all edges by their endpoints, and compute M_{cr} ; and for each node, temporarily compute a node-expansion. If property (A3) or (A4) fails, we report that P is not weakly simple. The operations introduced in Section 4 maintain properties (A1)–(A4) for all maximal paths inside an elliptical disk D_b .

281 **3.4 Clusters**

282 As a preprocessing step for spur elimination (Section 5), we group all nodes that do not lie inside
 283 a bar into *clusters*. After **node-expansion** and **new-bar-expansion**, all such nodes lie on a boundary
 284 of a disk (circular or elliptical). For every sober node u , we create $\deg(u)$ clusters as follows. Refer
 285 to Fig. 11. The node expansion has replaced u with new nodes on ∂D_u . Subdivide each segment
 286 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
 287 all adjacent (subdivision) nodes inside D_u . For each node u on the boundary of an elliptical disk
 288 D_b , subdivide the unique edge outside D_b incident to u with a node u^* . Form a cluster $C(u^*)$
 containing u and u^* . Every cluster will maintain the following invariants.




Figure 11: 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.

289

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

- 291 (I1) $C(u)$ induces a tree $T[u]$ in the image graph rooted at u .
 292 (I2) Every maximal path of P in $C(u)$ is of one of the following two types:
 293 (a) both endpoints are at the root of $T[u]$ and the path contains a single spur;
 294 (b) one endpoint is at the root, the other is at a leaf, and the path contains no spurs.

295 Additionally, each leaf node ℓ satisfies the following:

- 296 (I3) ℓ has degree one or two in the image graph of P ;
 297 (I4) there is no spur at ℓ ;
 298 (I5) no edge passes through ℓ (i.e., there is no edge $[a, b]$ such that $\ell \in ab$ but $\ell \notin \{a, b\}$).

299 Initially, every cluster trivially satisfies (I1) and (I2)a and every leaf node satisfies (I3)–(I5)
 300 since it was created by a subdivision.

301 **Dummy vertices.** Although the operations described in Sections 4 and 5 introduce new nodes in
 302 the clusters, the image graph will always have $O(n)$ nodes and segments. A vertex at a cluster node
 303 is called a *benchmark* if it is a spur or if it is at a leaf node; otherwise it is called a *dummy vertex*.
 304 Paths traversing clusters may jointly contain $\Theta(n^2)$ dummy vertices in the worst case, however we
 305 do not store these explicitly. By (I1), (I2) and (I4) a maximal path in a cluster can be uniquely
 306 encoded by one benchmark vertex: if it goes from a root to a spur at an interior node s and back,
 307 we record only $[s]$; and if it traverses $T[u]$ from the root to a leaf ℓ , we record only $[\ell]$.

308 **4 Bar simplification**

309 In this section we introduce three new ws-equivalent operations and show that they can eliminate
 310 all vertices from each bar independently (thus eliminating all forks). The bar decomposition is

311 pre-computed, and the bars remain fixed during this phase (even though all edges along each bar
 312 are eliminated).

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

316 4.1 Overview

317 After preprocessing in Section 3, we may assume that P has no edge crossings, no node crossings,
 318 and satisfies (A1)–(A4). We summarize the overall effect of the bar simplification subroutine for a
 319 given expanded bar.




Figure 12: The changes in the image graph caused by a bar simplification.

320 **Changes in the image graph G .** Refer to Fig. 12. All nodes in the interior of the ellipse D_b are
 321 eliminated. Some spurs on b are moved to new nodes in the clusters along ∂D_b . Segments inside
 322 D_b connect two leaves of trees induced by clusters.

323 **Changes in the polygon P .** Refer to Fig. 13. Consider a maximal path p in P that lies in D_b .
 324 The bar simplification will replace $p = [u, \dots, v]$ with a new path p' . By (I4)–(I5), only nodes u and
 325 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 vertices
 326 of p_1 in reverse order), then p' will be a spur in the cluster containing u (Fig. 13 (a)). If p has no
 327 such decomposition, but its two endpoints are at the same node, $u = v$, then p' will be a single
 328 edge connecting two leaves in the cluster containing u (Fig. 13 (b)). If the endpoints of p are at
 329 two different nodes, p' is an edge between two leaves of the clusters containing u and v respectively
 330 (Fig. 13 (c), (d)).

331 4.2 Primitives

332 The operations in Section 4.3 rely on two basic steps, **spur-reduction** and **node-split** (see Fig. 14).
 333 Together with merge and subdivision, these operations are called *primitives*.

334 **spur-reduction(u, v).** Assume² that every vertex at node u has at least one incident edge
 335 $[u, v]$. Replace any path $[u, v, u]$, with a single-vertex path $[u]$. (See Fig. 14, left.)

336 **node-split(u, v, w).** Assume segments uv and vw are consecutive in radial order around
 337 v , and not collinear with any adjacent segment; and P contains no spurs of the form
 338 $[u, v, u]$ or $[w, v, w]$. Create node v^* in the interior of the wedge $\angle(u, v, w)$ sufficiently
 339 close to v ; and replace every path $[u, v, w]$ with $[u, v^*, w]$. (See Fig. 14, right.)

340 The following two lemmas are generalizations of the results in [5][Section 5].

²In the terminology of [5], a node u satisfying this assumption was called a *base* of segment uv .

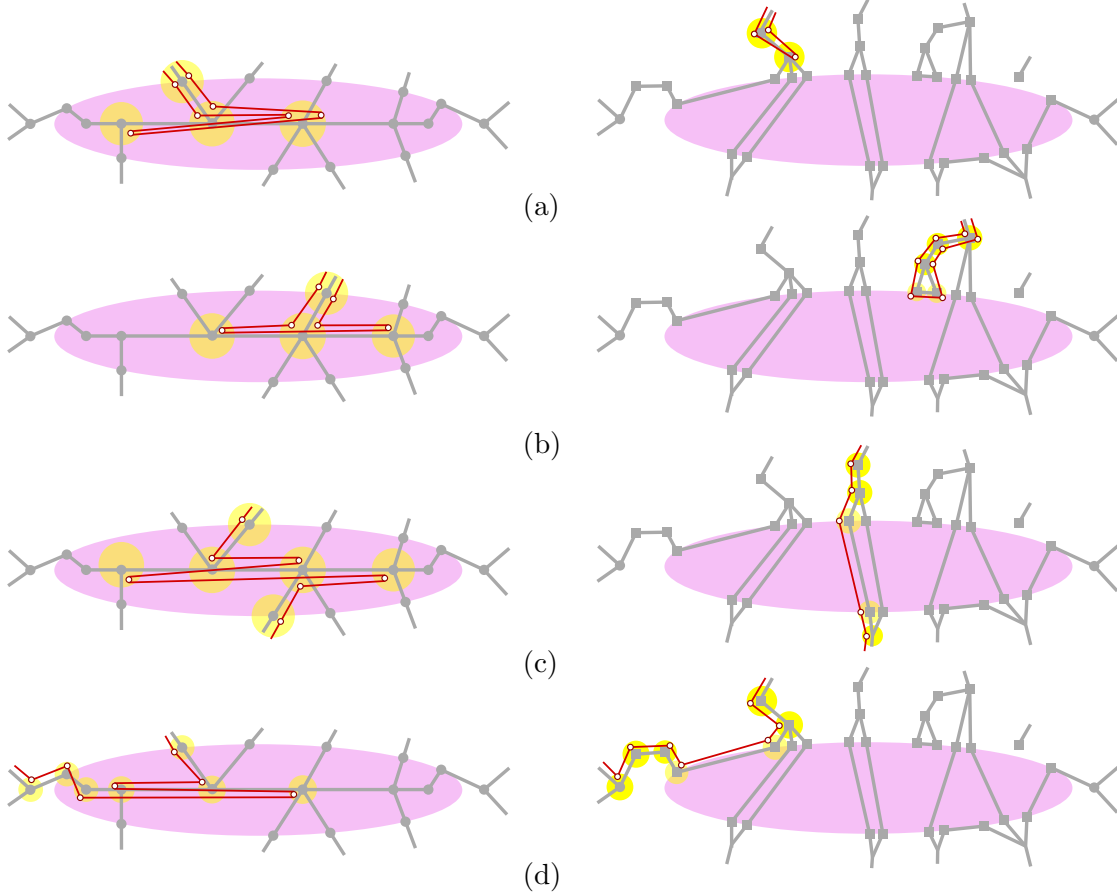


Figure 13: The changes in the polygon caused by a bar simplification.

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

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

350 Now suppose P' is weakly simple. Then, there exists a simple polygon $Q' \in \Phi(P')$ represented
 351 by its signature. Consider the clockwise order of edges in D_u . List all edges starting after the last
 352 edge $[u, v]$. Suppose that the i -th edge, $[p, u]$, is the first edge that is not adjacent to an edge $[u, v]$.
 353 Replace $[p, u]$ by $[p, u, v, u]$, and modify the signature by inserting two new edges into the total
 354 order of the edges along uv immediately after the i -th position. The resulting polygon P'' and the
 355 new signature define a polygon Q'' in the strip system of P'' . By construction, the new edge in
 356 D_v connects two consecutive endpoints in counterclockwise order around v , thus the new matching




Figure 14: Left: Spur-reduction(u, v). Right: Node-split(u, v, w).

357 in D_v has no crossing. All edges that appear before the i -th edge in the clockwise order of edges
 358 around u are also incident to an edge $[u, v]$ located below the new edges, thus, none of them cross
 359 the new edges in D_u . The operation deletes one edge and inserts two new edges in the perfect
 360 matching in D_u . Since Q' is simple, all endpoints at ∂D_u between the endpoints of the deleted
 361 edge can only be adjacent to other such points. Then, the new matching in D_u has no crossing and
 362 $Q'' \in \Phi(P'')$. By repeating this procedure until all vertices at u are adjacent to an edge $[u, v]$, we
 363 obtain a simple polygon $Q \in \Phi(P)$, hence P is weakly simple. \square

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

365 *Proof.* Let P' be obtained from P via node-split(u, v, w). First assume that P is weakly simple.
 366 Then there is a simple polygon $Q \in \Phi(P)$. Consider the clockwise order of edges around v . Since
 367 Q is simple, the order of the edges $[u, v]$ of paths $[u, v, w]$ must be the reverse order of its adjacent
 368 edges $[v, w]$ (the paths must be nested as shown in Fig. 14(right)). Because P has no spurs $[u, v, u]$
 369 or $[w, v, w]$, every edge between a pair of adjacent edges $[u, v]$ and $[v, w]$ is also part of a path
 370 $[u, v, w]$. Replace the paths $[u, v, w]$ by $[u, v^*, w]$ and set the order of edges at segments wv^* and
 371 v^*w to be the same order of the removed edges at wv and vw . That defines a polygon $Q' \in \Phi(P')$
 372 which is simple because the circular order of endpoints around D_u and D_w remains unchanged and
 373 the matching in D_{v^*} is a subset of the matching in D_v .

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

379 4.3 Operations

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

384 **pin-extraction(u, v).** Assumes that P satisfies (I1)–(I5) and contains a pin $[v, u, v] \in \text{Pin}$.
 385 By (I3), node v is adjacent to a unique node w outside of D_b . Perform the following three
 386 primitives: (1) subdivision of every path $[v, w]$ into $[v, w^*, w]$; (2) spur-reduction(v, u).
 387 (3) spur-reduction(w^*, v). See Fig. 15 for an example.

388 **V-shortcut(v_1, u, v_2).** Assumes that P satisfies (I1)–(I5) and $[v_1, u, v_2] \in \mathcal{V}$. Furthermore,
 389 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

390
391
392
393
394

uq is in the interior of the wedge $\angle(v_1, u, v_2)$. By (I3), nodes v_1 and v_2 are each adjacent to unique nodes w_1 and w_2 outside of D_b , respectively.

The operation executes the following primitives sequentially: (1) **node-split** (v_1, u, v_2) , which creates u^* ; (2) **node-split** (u^*, v_1, w_1) and **node-split** (u^*, v_2, w_2) ; which create $v_1^*, v_2^* \in \partial D_b$, respectively; (3) **merge** every path $[v_1^*, u^*, v_2^*]$ to $[v_1^*, v_2^*]$. See Fig. 16 for an example.




Figure 15: pin-extraction. Changes in the image graph (top), changes in the polygon (bottom).




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

395
396
397
398
399
400
401
402
403

Lemma 6. *pin-extraction and V-shortcut are ws-equivalent and maintain (I1)–(I5).*

Proof. **pin-extraction.** By construction, the operation maintains (I1)–(I5). Also, (I3)–(I5) ensure that **spur-reduction** (v, u) in step (2) satisfies its preconditions. Consequently, all three primitives are ws-equivalent.

V-shortcut. By construction, the operation maintains (I1)–(I5). The first two primitives are ws-equivalent by Lemma 5. The third step is ws-equivalent because triangle $\Delta(u^*v_1^*v_2^*)$ is empty of nodes and segments, by assumption. \square

L-shortcut operation. The purpose of this operation is to eliminate a vertex of a path that has an edge along a given bar. Before describing the operation, we introduce some notation; refer to

404 Fig. 17. For a node $v \in \partial D_b$, let L_v be the set of paths $[v, u_1, u_2]$ in P such that $u_1, u_2 \in \text{int}(D_b)$.
 405 Each path in \mathcal{P} is either in $\mathcal{P}in$, in \mathcal{V} or has two subpaths in some L_v . Recall that M_{cr} is the set
 406 of 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 $[u_1, u_2]$
 407 is *not* in M_{cr} .




Figure 17: Paths in $\mathcal{P}in$, \mathcal{V} , L_v^{TR} , L_v^{TL} , L_v^{BR} , and L_v^{BL} .

408 We partition L_v into four subsets (refer to Fig. 17): a path $[v, u_1, u_2] \in L_v$ is in

- 409 1. L_v^{TR} (*top-right*) if v is a *top* vertex and $x(u_1) < x(u_2)$;
- 410 2. L_v^{TL} (*top-left*) if v is a *top* vertex and $x(u_1) > x(u_2)$;
- 411 3. L_v^{BR} (*bottom-right*) if v is a *bottom* vertex and $x(u_1) < x(u_2)$;
- 412 4. L_v^{BL} (*bottom-left*) if v is a *bottom* vertex and $x(u_1) > x(u_2)$.

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

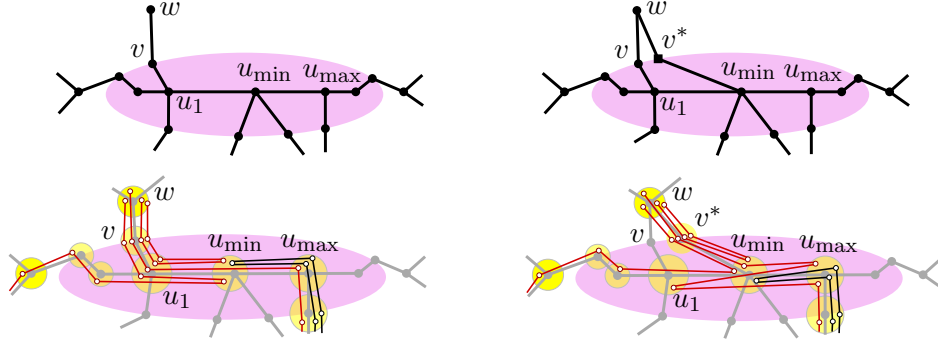


Figure 18: L-shortcut. Changes in the image graph (top), changes in the polygon (bottom).

415 **L-shortcut(v, TR)**. Assume that P satisfies (I1)–(I5), $v \in \partial D_b$ and $L_v^{TR} \neq \emptyset$. By (I3), v
 416 is adjacent to a unique node $u_1 \in b$ and to a unique node $w \notin D_b$. Let U denote the set
 417 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
 418 and rightmost node in U , respectively. Further assume that P satisfies:

- 419 (B1) no pins of the form $[v, u_1, v]$;
- 420 (B2) no edge $[p, u_1]$ such that segment pu_1 is in the interior of the wedge $\angle(v, u_1, u_2)$;
- 421 (B3) no edge $[p, q]$ such that $p \in \partial D_b$ is a top vertex and $q \in b$, $x(u_1) < x(q) < x(u_{\max})$.

422 Do the following (see Fig. 18 for an example).

- 423 (0) Create a new node $v^* \in \partial D_b$ to the right of v sufficiently close to v .

- 424 (1) For every path $[v, u_1, u_2] \in L_v^{TR}$ where u_1u_2 is the *only* longest edge of a cross
425 chain, create a crimp by replacing $[u_1, u_2]$ with $[u_1, u_2, u_1, u_2]$.
426 (2) Replace every path $[w, v, u_1, u_{\min}]$ by $[w, v^*, u_{\min}]$.
427 (3) Replace every path $[w, v, u_1, u_2]$, where $u_2 \in U$, by $[w, v^*, u_{\min}, u_2]$.

428 See Fig. 19 for an explanation of why L-shortcut requires conditions (B2)–(B3) and phase (1)
of the operation. If we omit any of these conditions, L-shortcut would not be ws-equivalent.




Figure 19: Cases in which L-shortcut is not ws-equivalent. Top left: P has a pin $[v, u_1, v]$ not satisfying (B2). Top right: P does not satisfy (B3). Bottom: the operation skips phase (1).

429

430 **Lemma 7.** *L-shortcut is ws-equivalent and maintains (I1)–(I5).*

431 *Proof.* Let P_1 be the polygon obtained from P after phase (1) of $L\text{-shortcut}(v, TR)$. Similarly, let
432 P_2 and P_3 be the polygons obtained after phase (2) and (3), respectively. Note that phase (1) of
433 the operation only creates crimps, and it is ws-equivalent by [2]. Let H be the set of edges $[u_1, u_2]$
434 of paths $[v, u_1, u_2] \in L_v^{TR}$. Phases (2)–(3) are equivalent to the concatenation of the primitives:
435 subdivision, node-split, and merge. Specifically, phase (2) is equivalent to subdividing every edge in
436 H into $[u_1, u_{\min}, u_2]$, where $u_2 \neq u_{\min}$, and applying $\text{node-split}(v, u_1, u_{\min})$ (which creates u_1^*) to P_2 .
437 Phase (3) consists of $\text{node-split}(w, v, u_1^*)$ (which creates v^*), and merging every path $[v^*, u_1^*, u_{\min}]$
438 to $[v^*, u_{\min}]$. The only primitive that may not satisfy its preconditions is $\text{node-split}(v, u_1, u_{\min})$:
439 segment u_1u_{\min} may be collinear with several segments of b , and P_2 may contain spurs that overlap
440 with u_1u_{\min} . We show that the spurs that may overlap with u_1u_{\min} do not pose a problem, and
441 we can essentially repeat the proof of Lemma 5. It remains to show that P_1 is weakly simple iff P_2
442 is weakly simple.

443 Assume that P_1 is weakly simple and consider a polygon $Q_1 \in \Phi(P_1)$ and its signature. Let Z
444 be the set of edges that are neither in H nor adjacent to an edge in H but lie above some edge in
445 H in the segments between u_1 and u_{\max} . Due to (B2) and (B3), the edges in Z must form paths
446 of the form $[z_1, z_2, z_3]$ where $x(u_1) \leq x(z_1) < x(u_{\max}) \leq x(z_2)$ and $x(u_{\max}) \leq x(z_3)$, otherwise Q_1
447 would not be simple. Let $e = [u_3, u_4]$, $x(u_3) < x(u_4)$, be the lowest edge in b that is adjacent to
448 some edge in H (shown in blue in Fig. 20). Due to (A1), the right endpoint u_4 is incident to an
449 edge in H , and we have $u_4 \in U$, thus $x(u_{\min}) \leq x(u_4) \leq x(u_{\max})$. We have $x(u_3) \leq x(u_1)$ either by
450 property (A2) or by a crimp introduced in phase (1). Subdivide every edge $[z_1, z_2] \in Z$ into a path

451 $[z_1, u_{\max}, z_2]$ if $z_2 \neq u_{\max}$ leaving $[u_{\max}, z_2]$ in the original position of $[z_1, z_2]$ in the ordering; place
452 the new edges $[z_1, u_{\max}]$ immediately below e in the ordering of the segments between u_1 and u_{\max}
453 while maintaining their relative order. We obtain a polygon Q'_1 in the strip system of P_1 . Notice
454 that a new edge $[z_1, u_{\max}]$ must be part of a path $[u_{\max}, z_1, u_{\max}]$. Since the relative order of the
455 new edges is maintained and every new edge is contained by e , no crossing can occur at a disk to
456 the left of u_{\max} . At $D_{u_{\max}}$, the circular order of endpoints is the same as before apart from the $|Z|$
457 edges below e . By construction there is no edge in the matching of $D_{u_{\max}}$ connecting the left and
458 right parts of $\partial D_{u_{\max}}$ above a subdivided path $[z_1, u_{\max}, z_2]$, or this edge would be in Z . Therefore,
459 Q'_1 is a simple polygon and $Q'_1 \in \Phi(P_1)$. We can proceed as in the proof of Lemma 5 using Q'_1 to
460 show that P_2 is weakly simple iff P_1 is weakly simple, that is, phase (2) is ws-equivalent.

461 Note that the intermediate polygons, P_1 and P_2 , may violate condition (A2), since phase (1)
462 introduces crimps. However, after phase (3), conditions (A1) and (A2) are restored, and operation
463 L-shortcut maintains (A1)–(A4) in the ellipse D_b .

464 □




Figure 20: If P_1 is weakly simple, we can change the signature maintaining weak simplicity.

465 4.4 Bar simplification algorithm

466 In this section, we show that the three operations (pin-extraction, V-shortcut, and L-shortcut) can
467 successively remove all spurs of the polygon P from a bar b .

468 Algorithm `bar-simplification(P, b)`.

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

- 470 (i) If $\mathcal{P}in \neq \emptyset$, pick an arbitrary pin $[v, u, v]$ and perform `pin-extraction(u, v)`.
- 471 (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, and perform
472 `V-shortcut(v_1, u, v_2)`.
- 473 (iii) Else if there exists $v \in \partial D_b$ such that $\widehat{L}_v^{TR} \neq \emptyset$, do:
 - 474 (a) Let v be the rightmost node where $L_v^{TR} \neq \emptyset$.
 - 475 (b) If $L_{v'}^{TL} = \emptyset$ for all $v' \in \partial D_b$, $x(v) < x(v')$ and $x(u'_1) < x(u_{\max})$, where u'_1 is the unique
476 neighbor of v' on b , do `L-shortcut(v, TR)`.
 - 477 (c) Else let v' be the leftmost node such that $x(v) < x(v')$ and $L_{v'}^{TL} \neq \emptyset$. If $L_{v'}^{TL}$ satisfies
478 (B3) do `L-shortcut(v', TL)`, otherwise halt and report that P is not weakly simple.
- 479 (iv) Else if there exists $v \in \partial D_b$ such that $L_v^{TL} \neq \emptyset$, repeat steps (iii)a–(iii)c with left–right and
480 TR – TL interchanged. (Note the use of L_v instead of \widehat{L}_v . The same applies to (vi)).
- 481 (v) Else if there exists $v \in \partial D_b$ such that $\widehat{L}_v^{BL} \neq \emptyset$, repeat steps (iii)a–(iii)c using BL and BR
482 in place of TR and TL , respectively, and left–right interchanged.

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

485 After the loop ends, perform old-bar-expansion (cf. Section 3.3) in the ellipse D_b ;
 486 Return P (end of algorithm).

487
 488 Informally, bar-simplification “unwinds” each polygonal chain in the bar, while extracting pins
 489 and V-chains as they appear, by alternating between steps (iii) to (vi) (see Fig. 21). Note that
 490 step (iii) uses \widehat{L}_v^{TR} (instead of L_v^{TR}) to avoid an infinite loop.




Figure 21: Life cycle of a cross chain in the while loop of bar-simplification. The steps applied, from left to right, are: (iv), (iii), (iv), and (vi).

491 **Lemma 8.** *The operations performed by bar-simplification(P, b) are ws-equivalent, and maintain*
 492 *properties (A1)–(A4) and (I1)–(I5) inside D_b . The algorithm either removes all nodes from the*
 493 *ellipse D_b , or reports that P is not weakly simple. The L-shortcut operations performed by the*
 494 *algorithm create at most two crimps in each cross-chain in \mathcal{P} .*

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

497 **Steps (i)–(ii).** Since every pin can be extracted from a polygon satisfying (I1)–(I5), we may assume
 498 that $\mathcal{Pin} = \emptyset$. 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
 499 minimal. Since $\mathcal{Pin} = \emptyset$, the only obstacle for the precondition is an edge $[u, q]$ such that segment
 500 uq is in the interior of the wedge $\angle(v_1, u, v_2)$ (or else the image graph would have a crossing). This
 501 edge is part of a path $[p, u, q]$. Node q must be on ∂D_b between v_1 and v_2 , otherwise there would
 502 be a node-crossing at u , contrary to (A3). However, $p \neq q$, otherwise $[p, u, q]$ would be a pin.
 503 Consequently, $[p, u, q]$ is a V-chain where $|x(p) - x(q)| < |x(v_1) - x(v_2)|$, contrary to the choice
 504 of $[v_1, u, v_2] \in \mathcal{V}$. This confirms that $\mathbf{V}\text{-shortcut}(v_1, u, v_2)$ satisfies all preconditions. Henceforth,
 505 assume that $\mathcal{Pin} = \emptyset$ and $\mathcal{V} = \emptyset$.

506 **Step (iii)–(iv).** By symmetry, we consider only step (iii). We distinguish between two cases.

507 **Case 1: the conditions of (iii)b are satisfied.** We need to show that $\mathbf{L}\text{-shortcut}(v, TR)$ satisfies
 508 (B1)–(B3). Since $\mathcal{Pin} = \emptyset$, condition (B1) is met. Suppose there is an edge $[p, u_1]$ such that segment
 509 pu_1 is in the interior of the wedge $\angle(v, u_1, u_{\min})$. Clearly, $p \in \partial D_b$ is a top node. Then edge $[p, u_1]$
 510 is part of a path $[p, u_1, q]$. Recall that q cannot be a top vertex on ∂D_b since $\mathcal{Pin} = \mathcal{V} = \emptyset$. If
 511 q is a bottom vertex or a vertex on b such that $x(q) < x(u_{\min})$ there is a node crossing at u_{\min}
 512 contradicting (A3). Then q must be on b and $x(q) > x(u_{\min})$ which contradicts the choice of node v .
 513 This confirms (B2). We argue similarly for (B3). Suppose there is an edge $[p, q]$ such that $p \in \partial D_b$
 514 is a top vertex and $q \in b$, $x(u_1) < x(q) < x(u_{\max})$. This edge is part of a path $[p, q, r]$. Node r must
 515 be on or above b , otherwise there would be a node-crossing at q . It cannot be a top vertex, since

516 $\text{Pin} = \mathcal{V} = \emptyset$. It cannot be to the left of q , otherwise the conditions of (iii)b are satisfied; and it
517 cannot be to the right of q , otherwise $L_p^{TR} \neq \emptyset$ with $x(v) < x(p)$, contrary to the choice of v . This
518 confirms that $\text{L-shortcut}(v, TR)$ satisfies (B2)-(B3) and can be performed.

519 **Case 2: the conditions of (iii)b are not satisfied.** Let the path $[v', u'_1, u'_{\min}] \in L_{v'}^{TL}$ be selected
520 in $\text{L-shortcut}(v', TL)$ by the algorithm. Conditions (B1)-(B2) are satisfied similar to Case 1. If (B3)
521 fails, there is an edge $[p, q]$ such that $p \in \partial D_b$ is a top vertex and $q \in b$, $x(u'_{\max}) < x(q) < x(u_{\max})$
522 (Recall that left and right are interchanged in L^{TL}). Edge $[p, q]$ is part of a path $[p, q, r]$, where $r \in b$,
523 similar to Case 1. This implies $[p, q, r] \in L_p^{TR} \cup L_p^{TL}$. If $x(v) < x(p) < x(v')$, then either $L_p^{TR} \neq \emptyset$,
524 which contradicts the choice of v , or $L_p^{TL} \neq \emptyset$, which contradicts the choice of v' . Consequently,
525 $x(p) \leq x(v)$. This implies $x(u'_{\max}) < x(p) \leq x(v)$, so the paths $[v, u_1, u_{\max}]$ and $[v', u'_1, u'_{\max}]$ form
526 a forbidden configuration described in Corollary 1(2); see also Fig. 10(b). Therefore the algorithm
527 correctly finds that P is not weakly simple.

528 **Steps (v)-(vi).** If steps (i)-(iv) do not apply, then $\widehat{L}_v^{TR} \cup L_v^{TL} = \emptyset$. That is, for every path
529 $[v, u_1, u_2] \in L^{TR}$, we have $[u_1, u_2] \in M_{cr}$. In particular, there are no top chains. The operations
530 in (v)-(vi) do not change these properties. Consequently, once steps (v)-(vi) are executed for
531 the first time, steps (iii)-(iv) are never executed again. By a symmetric argument, steps (v)-(vi)
532 eliminate all paths in $\widehat{L}_v^{BL} \cup L_v^{BR}$. If the while loop terminates, every edge in b is necessarily in
533 M_{cr} and $L_v^{TL} \cup L_v^{BR} = \emptyset$. Consequently, by Lemma 2, b contains no spurs and *old-bar-expansion* is
534 ws-equivalent. This operation eliminates all nodes in the interior of D_b .

535 **Termination.** Each pin-extraction and V-shortcut operation reduces the number of vertices of
536 P within D_b . Operation $\text{L-shortcut}(v, X)$, $X \in \{TR, TL, BR, BL\}$, either reduces the number
537 of interior vertices, or produces a crimp if edge $[u_1, u_2]$ is a longest edge of a cross-chain. For
538 termination, it is enough to show that, for each cross-chain $c \in \mathcal{P}$, the algorithm introduces a crimp
539 at most once in steps (iii)-(iv), and at most once in steps (v)-(vi). Without loss of generality,
540 consider step (iii). We apply an L-shortcut in two possible cases. We show that it does not
541 introduce crimps in Case 2. In step (iii)c, we only perform $\text{L-shortcut}(v', TL)$ if (B3) is satisfied and
542 $x(u'_1) < x(u_{\max})$. So for all $[v', u'_1, u'_2] \in L_{v'}^{TL}$, we have $x(u_1) < x(u'_2)$. Suppose, for contradiction,
543 that $[u'_1, u'_2]$ is the only longest edge of some cross chain (and hence L-shortcut would introduce a
544 crimp). Then, $[u'_1, u'_2] \in M_{cr}$ is contained in $[u_1, u_{\max}]$, contradicting (A4).

545 Consider Case 1. Notice that $\text{L-shortcut}(v, TR)$ is executed only if there exists a top node p with
546 $x(p) \leq x(v)$ such that $\widehat{L}_p^{TR} \neq \emptyset$. Suppose that $\text{L-shortcut}(v, TR)$ introduces a crimp in the path
547 $[v, u_1, u_2] \in L_v^{TR}$. This operation removes this subpath of a cross chain from L_v^{TR} , but introduces
548 $[v^*, u_2, u_1]$ into $L_{v^*}^{TL}$. By the time the algorithm executes $\text{L-shortcut}(v^*, TL)$, we know that for every
549 top vertex p with $x(p) < x(v^*)$, $\widehat{L}_p^{TR} = L_p^{TL} = \emptyset$. This implies that, after $\text{L-shortcut}(v^*, TL)$ is
550 performed, although a path $[v^{**}, u_1, u_2]$ is introduced in $L_{v^{**}}^{TR}$, operation $\text{L-shortcut}(v^{**}, TR)$ can
551 never be performed, because for all top vertex p with $x(p) \leq x(v^{**})$, $\widehat{L}_p^{TR} = \emptyset$ or else (A4) would
552 be violated. The same arguments apply to steps (v)-(vi). \square

553 **Lemma 9.** *Algorithm bar-simplification(P, b) takes $O(m \log m)$ time, where m is the number of*
554 *vertices in b .*

555 *Proof.* Operations pin-extraction, V-shortcut, and L-shortcut each make $O(1)$ changes in the image
556 graph. Operations pin-extraction and V-shortcut decrease the number of vertices inside D_b . Each

557 L-shortcut does as well, but they may jointly create $2|\mathcal{P}| = O(m)$ crimps, by Lemma 7. So the total
 558 number of operations is $O(m)$.

559 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
 560 $[u_1]$ shifts to $[u_2]$, but no vertex is eliminated. In the worst case, one L-shortcut modifies $\Theta(m)$
 561 paths, so in $\Theta(m)$ operations the total number of vertex shifts is $\Theta(m^2)$.

562 Our implementation does not maintain the paths in \mathcal{P} explicitly. Instead, we use set operations.
 563 We maintain the sets $\mathcal{P}in$, \mathcal{V} , and L_v^X , with $v \in \partial D_b$ and $X \in \{TR, TL, BR, BL\}$, in sorted lists.
 564 The pins $[v, u, v] \in \mathcal{P}in$ are sorted by $x(v)$; the wedges $[v_1, u, v_2] \in \mathcal{V}$ are sorted by $|x(v_1) - x(v_2)|$.
 565 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 (I4), and so it
 566 is enough to store vertex $[u_2]$; these vertices are stored in a list sorted by $x(u_2)$. We also maintain
 567 binary variables to indicate for each path $[v, u_1, u_2] \in L_v^X$ whether it is part of a cross chain, and
 568 whether $[u_1, u_2]$ is the only longest edge of that chain.

569 Steps (i)-(ii) remove pins and V-chains, taking linear time in the number of removed vertices,
 570 without introducing any path in any set. Consider L-shortcut(v, TR), executed in one of steps (iii)–
 571 (iv) which can be generalized to other occurrences of the L-shortcut operation. The elements
 572 $[v, u_1, u_{\min}] \in L_v^{TR}$ are simplified to $[v^*, u_{\min}]$. For each of these paths, say that the next edge along
 573 P is $[u_{\min}, u_3]$. Then, the paths $[v^*, u_{\min}, u_3]$ are inserted into either $\mathcal{P}in \cup \mathcal{V}$ if $u_3 \in \partial D_b$ is a top
 574 vertex, or $L_{v^*}^{TL}$ if $u_3 \in b$. We can find each chain $[v, u_1, u_{\min}] \in L_v^{TR}$ in $O(1)$ time since L_v^{TR} is
 575 sorted by $x(u_2)$. Finally, all other paths $[v, u_1, u_2] \in L_v^{TR}$, where $u_2 \neq u_{\min}$, become $[v^*, u_{\min}, u_2]$
 576 and they form the new set $L_{v^*}^{TR}$. Since we store only the last vertex $[u_2]$, which is unchanged, we
 577 create $L_{v^*}^{TR}$ at no cost.

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

582 5 Spur elimination algorithm

583 After bar-simplification (Section 4), we obtain a polygon that has no forks and every spur is at
 584 an interior node of some cluster (formed on the boundary of some ellipse D_b). In the absence
 585 of forks, we can decide weak simplicity using [5][Theorem 5.1], but a naïve implementation runs
 586 in $O(n^2 \log n)$ time: successive applications of spur-reduction would perform an operation at each
 587 dummy vertex. In this section, we show how to eliminate spurs in $O(n \log n)$ time.

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

594 In the course of our algorithm, an analogue of the pin-extraction operation will extract a spur
 595 from one group into an “adjacent” group. This requires a well-defined adjacency relation between
 596 groups. By construction, if a segment uv connects nodes in different clusters, both u and v are
 597 leaves or both are root nodes. For every pair of clusters, $C(u)$ and $C(v)$, with adjacent roots, u
 598 and v , create a *group* $G_{uv} = C(u) \cup C(v)$; see Fig 22. By construction, the groups are pairwise
 599 disjoint. Two groups are called *adjacent* if they have two adjacent leaves in the image graph.




Figure 22: The formation of a group G_{uv} , containing clusters $C(u)$ and $C(v)$. Leaf nodes are shown as black dots.

600 Recall that a maximal path in each cluster is represented by benchmark vertices (leaves and
 601 spurs). We denote by $[u_1; \dots; u_k]$ (using semicolons) a maximal path inside a group defined by the
 602 benchmark vertices u_1, \dots, u_k . For a given group G_{uv} , let \mathcal{P} denote the set of maximal paths with
 603 vertices in G_{uv} ; and let \mathcal{B} be the set of subpaths in \mathcal{P} between consecutive benchmark vertices.

604 **Remark 2.** *By invariants (I1), (I2) and (I4), a path in \mathcal{P} of a group G_{uv} has alternating benchmark*
 605 *vertices between $C(u)$ and $C(v)$. Consequently, every path in \mathcal{B} has one endpoint in $C(u)$ and one*
 606 *in $C(v)$, and each spur in G_{uv} is incident to two paths in \mathcal{B} .*

607 **Spur-elimination algorithm.** Assume that \mathcal{G} is a partition of the nodes of the image graph into
 608 groups satisfying (I1)–(I5). We consider one group at a time, and eliminate all spurs from one
 609 cluster of that group. When we process one group, we may split it into two groups, create a new
 610 group, or create a new spur in an adjacent group (similar to pin-extraction in Section 4). The latter
 611 operation implies that we may need to process a group several times. Termination is established
 612 by showing that each operation reduces a weighted sum of the number of benchmark vertices (i.e.,
 613 spurs and boundary vertices). Initially, the number of benchmarks is $O(n)$.

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

615 While P contains a spur, do:

- 616 1. Choose a group $G_{uv} \in \mathcal{G}$ that contains a spur, w.l.o.g. contained in cluster $C(u)$
 617 and compute its supporting data structures.
- 618 2. While $T[u]$ contains an interior node, do:
 - 619 (a) If u contains no spurs and is incident to only two edges uv and uw , eliminate
 620 u with a merge operation. Rename node w to u which becomes the new root
 621 of the tree $T[u]$.
 - 622 (b) If u contains spurs, eliminate them as described in Section 5.2.
 - 623 (c) If u contains no spurs, split G_{uv} into two groups along a chain of segments
 624 that contains uv as described in Section 5.3. Rename a largest resulting group
 625 to G_{uv} .

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

628 **5.1 Data structures**

629 Consider a group G_{uv} composed of two trees $T[u]$ and $T[v]$ rooted at u and v , respectively. In the
 630 algorithm *spur-elimination*, we dynamically maintain the image trees $T[u] \cup T[v]$, and the set of paths
 631 \mathcal{B} . In each group G_{uv} , we maintain only $O(|\mathcal{B}|)$ nodes that contain benchmark vertices or have
 632 degree higher than 2. Dummy nodes of degree-2 that contain no benchmark vertices are redundant
 633 for the combinatorial representation, and will be eliminated with *merge* operations. However, a
 634 polyline formed by a chain of dummy nodes of degree-2 cannot always be replaced by a straight-
 635 line segment (this might introduce artificial crossings). By Remark 1, it suffices to maintain the
 636 combinatorial embeddings of the trees $T[u]$ and $T[v]$ (ie., the counterclockwise order of incident the
 637 segments around each node).

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




Figure 23: The geometry of crossing benchmark-to-benchmark paths.

645 For the image graph of G_{uv} , we maintain the following data structures.

- 646 • We store trees $T[u]$ and $T[v]$ each using the dynamic data structure of [7], which supports
 647 $O(1)$ -time insertion and deletion of leaves, merging interior nodes of degree 2, subdivision of
 648 edges, and lowest common ancestor queries.
- 649 • Imagine that G_{uv} is inside an axis-aligned rectangle with the leaves of $T[u]$ along the top
 650 edge and leaves of $T[v]$ along the bottom edge (see Fig. 23). For each tree, we maintain
 651 a left-to-right Euler tour in an order-maintenance data structure [17, 3], which supports
 652 insertions immediately before or after an existing item, deletions, and precedence queries,

653 each in $O(1)$ amortized time. For any node w , let w^b and w^\sharp respectively denote the first and
654 last occurrences of w in the Euler tour. We refer to the elements of the Euler tour as *tokens*.
655 We write $x < y$ to denote that some token x occurs before (“to the left of”) another token y
656 in their common Euler tour.

- 657 • We also maintain the cyclic list of all leaves of the tree $T[u] \cup T[v]$ (in the order determined
658 by the Euler tour above; note that we have $w^\sharp = w^b$ for a leaf w).

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

- 663 • The set of benchmark-to-benchmark paths $[s; t] \in \mathcal{B}$ is stored in four lists, sorted by s^b , s^\sharp , t^b ,
664 and t^\sharp , respectively. The sorted lists can be computed in $O(|\mathcal{B}|)$ time by an Eulerian traversal
665 of the tree.
- 666 • 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,
667 sorted by t^b and t^\sharp , respectively.
- 668 • An *interval tree* for all intervals $I[s_1; t; s_2]$ that can report, for a query node q , all intervals
669 containing q in $O(\log n)$ time.

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

671 **Crossing paths.** The data structure described above can determine in $O(1)$ time whether two
672 paths in \mathcal{B} cross. Straightforward case analysis implies the following characterization of path
673 crossings (refer to Fig. 23).

674 **Lemma 10.** *Let s and u be arbitrary nodes in tree T_1 , and let t and v be arbitrary nodes in T_2 .
675 Paths $[s; t]$ and $[u; v]$ cross if and only if either (1) $s^\sharp < u^b$ and $t^b > v^\sharp$, or (2) $s^b > u^\sharp$ and $t^\sharp > v^b$.*

676 5.2 Eliminating spurs from a root

677 We describe step 2b of Algorithm *spur-elimination*. Suppose that the root node u contains a
678 spur. The following operation eliminates all spurs from u , but the resulting cluster $C(v)$ need not
679 satisfy (I2) and (I4), and we need to perform other operations to restore these properties. Refer to
680 Fig. 24(a)–(b) for an example.

681 **spur-shortcut(u).** Assume that G_{uv} satisfies invariants (I1)–(I5), and u contains a spur.
682 Replace every path $[t_1; u; t_2]$ by $[t_1; t_2]$. Let \mathcal{S} be the set of all such modified paths.

683 **Lemma 11.** *spur-shortcut is ws -equivalent and maintains properties (I1), (I3) and (I5).*

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

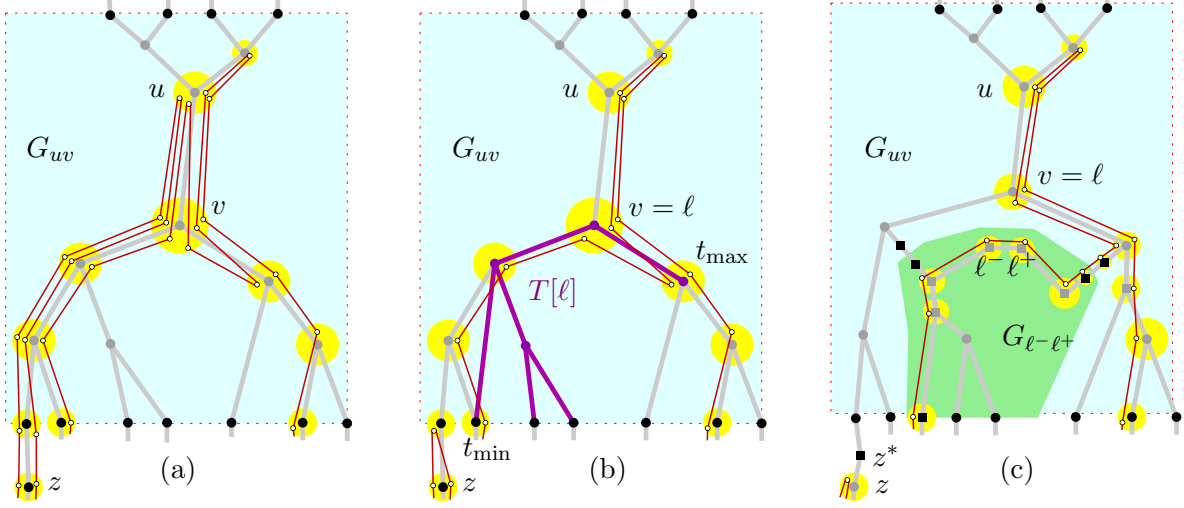


Figure 24: (a) u contain spurs. (b) After eliminating spurs, $T[v]$ does not satisfy (I2). (c) The analogues of pin-extraction and V-shortcut.

689 Note that for every path $[t_1; u; t_2]$, both t_1 and t_2 are in $T[v]$ (cf. Remark 2) and the path $[t_1; t_2]$
 690 is uniquely defined by (I1). However, the paths $[t_1; t_2] \in \mathcal{S}$ violate (I2) or (I4). We proceed with
 691 a sequence of “repair” steps to restore them, after which the total number of benchmark vertices
 692 decrease by at least $|\mathcal{S}|$. The following applies for when t_1 and t_2 are in ancestor-descendent relation,
 693 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 root.

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

- 695 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 $[\min(t_1, t_2)]$.
- 696 2. If $t_1 = t_2$, and t_1 is not a leaf of $T[v]$, then replace $[t_1; t_2]$ with $[t_1]$.
- 697 3. If $t_1 = t_2$, and t_1 is a leaf of $T[v]$, then do: by (I3), node t_1 is adjacent to a unique
 698 node $z \notin G_{uv}$. Subdivide segment t_1z with a new node z^* (added to the cluster
 699 containing z), subdivide every edge $[t_1, z]$ into $[t_1, z^*, z]$, and then replace every
 700 path $[z, z^*, t_1, z^*, z]$ with $[z]$. See Fig. 24(b)–(c) for an example.

701 These steps restore (I2) and (I4) for the affected paths $[t_1; t_2] \in \mathcal{S}$, and are ws-equivalent: Steps 1–
 702 2 do not modify the polygon (they change only the benchmarks); and step 3 is analogous to
 703 $\text{pin-extraction}(v)$.

704 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,
 705 we perform an elaborate version of the V-shortcut operation, that creates a new group. For every
 706 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
 707 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
 708 (refer to Fig. 24).

709 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
 710 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^\sharp , and
 711 let t_{\min} be the minimum node; and similarly sort the nodes $t_2 \in N^+$ by t_2^\flat , and let t_{\max} be the
 712 maximum node. Since all nodes in N are descendants of ℓ in $T[v]$, we have $t_{\min}^\sharp \leq t_1^\sharp < \ell^\sharp$ for all
 713 $t_1 \in N^-$ and $\ell^\flat < t_2^\flat \leq t_{\max}^\flat$ for all $t_2 \in N^+$.

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

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

720 We can find the set N' of nodes t such that $t_{\min}^\# < t^\#$ and $t^b < t_{\max}^b$ in $O(|N'|)$ time, by lowest
 721 common ancestor queries from the leaves between t_{\min} and t_{\max} . The algorithm halts and reports
 722 that the input polygon is not weakly simple if some node in N' has a path satisfying Lemma 12.
 723 We can now assume that $N' \subset N$. The nodes in N induce a binary tree, denoted $T[\ell]$, of size at
 724 most $2|N|$: its nodes are all nodes in N and the lowest common ancestors of consecutive nodes in
 725 N^- and N^+ respectively. Note that a segment of $T[\ell]$ might not correspond to a segment $T[v]$ (see
 726 Fig. 24(b)). Denote by C^* the path between t_{\min} and t_{\max} in $T[\ell]$.

727 We now define the changes in the image graph. Every node $t \in N \setminus C^*$ is deleted from G_{uv} ,
 728 and added to the new group. Create two nodes, ℓ^- and ℓ^+ , in $G_{\ell-\ell^+}$ sufficiently close to ℓ in the
 729 wedge between the two children of ℓ , and connect them by a segment $\ell^-\ell^+$. Duplicate each node
 730 $t \in C^* \setminus \{\ell\}$, by creating a node t' (added to $G_{\ell-\ell^+}$) sufficiently close to t , and add a segment tt' .
 731 Subdivide every segment tt' with two new boundary nodes, t_{leaf} (added to $T[v]$) and t'_{leaf} (added
 732 to $G_{\ell-\ell^+}$). The nodes t or t' might now have degree 4. Adjust the image graph so that the group
 733 trees are binary. Finally partition the nodes in $G_{\ell-\ell^+}$ into two trees, $T[\ell^-]$ and $T[\ell^+]$, rooted at ℓ^-
 734 and ℓ^+ , respectively.

735 We now define the changes in the polygon. Replace every path $[t; t_1] \in \mathcal{S}_\ell$, where $t \in C^*$, by
 736 $[t'; t_1]$ if it is adjacent to a path $[t; t_2] \in \mathcal{S}_\ell$; and by $[t_{leaf}, t'_{leaf}, t_1]$ otherwise. Now we can build \mathcal{B}'
 737 as the set of benchmark-to-benchmark paths $[t'_1; t'_2]$ where $t'_1, t'_2 \in G_{\ell-\ell^+}$ in $O(|\mathcal{B}'|)$ time.

738 To prove ws-equivalence, we consider the changes in the polygon. These amount to a sequence
 739 of ws-equivalent primitives: a **node-split** at ℓ , a sequence of **node-splits** along the chain C from ℓ
 740 to t_{\min} and t_{\max} , respectively, and **subdivision** operations that create the new leaf nodes and **merge**
 741 **operations** at degree-2 nodes that no longer contain spurs. The creation of new groups takes $O(|\mathcal{S}_\ell|)$
 742 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
 743 updated in $O(|\mathcal{S}_\ell| \log n)$ time. Overall, operation **spur-shortcut**(u) and the repair steps that follow
 744 take $O(|\mathcal{S}| \log n)$ time.

745 5.3 Splitting a group in two

746 In this section we describe step 2c of Algorithm **spur-elimination**(P, \mathcal{G}). Assume that G_{uv} satisfies
 747 invariants (I1)–(I5) and u contains no spur.

748 Denote the left and right child of u by u^- and u^+ , respectively. Let $\mathcal{B}^-, \mathcal{B}^+ \subset \mathcal{B}$, resp., be
 749 the set of benchmark-to-benchmark paths that contain u^- and u^+ . We split G_{uv} into two groups
 750 induced by \mathcal{B}^- and \mathcal{B}^+ , respectively. Refer to Fig. 25.

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




Figure 25: Splitting group G_{uv} . (a) Changes in the image graph. (b) Changes in the polygon.

757 First, we find \mathcal{B}^- and \mathcal{B}^+ . Compute \mathcal{B}^- using the list of paths $[s; t] \in \mathcal{B}$ sorted by s^\sharp or s^\flat .
 758 Since both lists naturally split into corresponding lists for \mathcal{B}^- and \mathcal{B}^+ , we can split these lists in
 759 $O(\min(|\mathcal{B}^-|, |\mathcal{B}^+|)) = O(|\mathcal{B}^-|)$ time. To construct the list of \mathcal{B}^+ sorted by t^\sharp and t^\flat , we start with
 760 the corresponding lists for \mathcal{B} , and delete all elements of \mathcal{B}^- in $O(|\mathcal{B}^-|)$ time. To compute the lists
 761 sorted by t^\sharp and t^\flat for \mathcal{B}^- , we shall first compute the subtree $T[v^-]$ induced by \mathcal{B}^- . However, we
 762 can already find the maximum t^\sharp of a path $[s; t] \in \mathcal{B}^-$ in $O(|\mathcal{B}^-|)$ time.

763 Next, we test for crossings between the paths in \mathcal{B}^- and the paths in \mathcal{B}^+ . Let t_-^\sharp be the
 764 maximum t^\sharp 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,
 765 there is such a crossing iff $t_+^\flat < t_-^\sharp$, which can be determined in $O(1)$ time based on our sorted lists.
 766 If a crossing is detected, the algorithm halts and reports that P is not weakly simple.

767 The trees $T[u^-]$ and $T[u^+]$ are simple subtrees of $T[u]$; but splitting $T[v]$ is nontrivial. We use
 768 the Eulerian cycle of all leaves to find the rightmost leaf ℓ_0 in $T[v]$ for which $\mathcal{B}_{\ell_0} \cap \mathcal{B}^- \neq \emptyset$. Let
 769 $C_0 = [\ell_0; u]$. We do not compute C_0 explicitly, as it may contain more than $O(|\mathcal{B}^-|)$ nodes, but we
 770 can test whether a query node t of $T[v]$ is in C_0 in $O(1)$ time by checking whether $\text{lca}(\ell_0, t) = t$.
 771 Since the 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 ,
 772 and all common nodes of $T[v^-]$ and $T[v^+]$ are in C_0 . The image graph of $T[v^-]$ can be computed
 773 from scratch using \mathcal{B}^- in $O(|\mathcal{B}^-|)$ time. Replace each node t of $T[v^-]$ that is in C_0 by a duplicate
 774 copy t^- located sufficiently close to t , to the right of t . The tree $T[v^+]$ is computed from $T[v]$ by
 775 node deletion and merge operations as follows. First delete all nodes that are in $T[v^-]$ but not in
 776 C_0 . For every node t of $T[v^-]$ that lies in C_0 , if t has degree-3 in $T[v^-]$ and $\mathcal{B}_t^+ = \emptyset$, then it is a
 777 degree-2 node in $T[v]$ with no spurs, and so we can delete t by merging its two incident segments.
 778 As a result, $T[v]$ becomes a tree induced by \mathcal{B}^+ . It remains to resolve the connections between
 779 trees.

780 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
 781 on all chains $[s_1; t; s_2] \in \mathcal{V}^0$ will be eliminated (they will become adjacent leaves in the two resulting
 782 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
 783 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
 784 a segment between t and t^- , and subdivide the segment t^-t with two new nodes t_{leaf}^- and t_{leaf} in
 785 $T[v^-]$ and $T[v]$, respectively. The degree of nodes t or t^- might increase to 4; and so we adjust the
 786 image graphs so that both trees are binary. The image graph is now split into groups $G_{u^-v^-}$ and
 787 G_{uv} .

788 We next define the changes in the polygon. Replace every chain $[s_1; t; s_2] \in \mathcal{V}^0$ with a new chain

789 $[s_1; t_{leaf}^-; t_{leaf}; s_2]$, replacing also the corresponding paths in the lists \mathcal{B}^- and \mathcal{B}^+ in $O(|\mathcal{V}^0|)$ time.
790 In the sorted lists for \mathcal{B}^- and \mathcal{B}^+ , this is done by deletions and reinsertions. Note that all leaves
791 t_{leaf}^- (resp., t_{leaf}) are at the end (resp., beginning) of the Euler tour of $T[v^-]$ (resp., $T[v]$), so the
792 deletions can be performed in $O(|\mathcal{V}^0|)$ time; and the insertions take $O(|\mathcal{V}^0| \log n)$ time.

793 The changes in the polygon are equivalent to a sequence of ws-equivalent primitives: a node-
794 split operation at u , followed by a sequence of node-splits along the chain C_0 from ℓ_0 to u , and
795 subdivision operations that create the new leaf nodes between the two groups. The interval tree
796 does not need to be updated since its query time remains $O(\log n)$. Consequently, we can split G_{uv}
797 in $O(\min(|\mathcal{B}^-|, |\mathcal{B}^+|) + |\mathcal{V}^0| \log n + \log n)$ time.

798 5.4 Analysis of the spur-elimination algorithm

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

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

805 Step 2a removes an interior node of degree 2 in $O(1)$ time. Interior nodes are created only when
806 they contain a spur, so at most $O(m)$ interior nodes are ever created, and steps 2a altogether take
807 $O(m)$ time. Step 2b eliminate $|\mathcal{S}|$ spurs in $O(|\mathcal{S}| \log m)$ time. Eventually, all spurs are eliminated,
808 thus steps 2b altogether take $O(m \log m)$ time. Step 2c takes $O(\min(|\mathcal{B}^-|, |\mathcal{B}^+|) + |\mathcal{V}^0| \log m + \log m)$
809 time. By a standard heavy-path decomposition argument, the terms $\min(|\mathcal{B}^-|, |\mathcal{B}^+|)$ contribute
810 $O(m \log m)$ time. Every chain in \mathcal{V}^0 corresponds to a spur that is destroyed in a step 2c (and
811 no new spurs are created in step 2c), therefore the terms $O(|\mathcal{V}^0| \log m)$ sum to $O(m \log m)$ over
812 the course of the algorithm. Since every execution of step 2c increases the number of groups by
813 one, and this step is repeated $O(m)$ times, and the $\log m$ terms sum to $O(m \log m)$ in the entire
814 algorithm. \square

815 Algorithm $\text{spur-elimination}(P, \mathcal{G})$ returns a polygon P' , a set \mathcal{G}' of groups, and a set \mathcal{B}' of
816 benchmark-to-benchmark paths, each of which connects two leaves in two different clusters of a
817 group. We can now decide whether P' is weakly simple in $O(n)$ time similarly to [5][Section 3.3]. If
818 P' is weakly simple, then the benchmark-to-benchmark paths in each group G_{uv} can be perturbed
819 into disjoint paths, and they traverse the corridor N_{uv} from D_u to D_v , in a specific order. The
820 Euler tours of the leaves in the clusters $C(u)$ and $C(v)$ determine this ordering uniquely, up to
821 permutations of the paths between the same pair of leaves, and such an ordering can be computed
822 in $O(n)$ time. By concatenating the benchmark-to-benchmark each each benchmark node according
823 to the these ordering, we obtain a unique 2-regular graph Q' , given as one or more cyclic sequences
824 of benchmarks. Polygon P' is weakly simple if and only if Q' is connected and visits the benchmarks
825 in the same cyclic order as P' . These properties can be verified by a simple traversal of P' and Q'
826 in $O(n)$ time. This completes the proof of Theorem 1.

6 Perturbing weakly simple polygons into simple polygons

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

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

Let $P = (p_0, \dots, p_{n-1})$, $n \geq 3$, be a weakly simple polygon with an image graph G . Assume that the nodes of G are partitioned into a set \mathcal{C} of disjoint clusters satisfying invariants (I1)–(I5) such that every bar is either entirely in a cluster or outside of all clusters. Let $Q = (p'_0, \dots, p'_{n-1})$ 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 the image graph G has no vertical segments. In each segment uv of G outside of clusters, the above-below relationship yields a total ordering over the edges of Q that contain uv . For each bar b outside of clusters, the total orders of the segments along b are consistent (since the above-below relationship between two edges is the same in every corridor). Consequently, the transitive closure of these total orders is a partial order over all edges in b . Consider a linear extension of such a partial order. The collection of these total orders is a *bar-signature* of Q . Since the linear extensions need not be unique, a polygon $Q \in \Phi(P)$ may have several bar-signatures.

Given a bar-signature of a perturbation of P , we can (re)construct an approximate simple polygon Q' as follows; refer to Fig. 26. For every bar $b = uv$ of G outside of clusters, let the *volume* $\text{vol}(uv)$ be the number of edges of P that lie on b . Place $\text{vol}(uv)$ parallel line segments, called *lanes*, between ∂D_u and ∂D_v in the region U_ε , ordered from bottom to top (the lanes will contain the edges of Q'). For the i -th edge pq in the total order of b , let the corresponding edge in Q' be the shortest edge connecting ∂D_p and ∂D_q in the i -th lane. For each cluster $C(u)$, denote by $R(u)$ the union of all disks D_v , $v \in C(u)$, and all corridors between nodes in $C(u)$. If $C(u)$ contains only the node u , then $R(u) = D_u$, but $R(u)$ is always simply connected since $C(u)$ induces a tree $T[u]$. For each cluster $C(u)$, construct a noncrossing polyline matching between the endpoints of the edges in $\partial R(u)$ that connects the endpoints corresponding to a maximal subpath in $T[u]$. The edges in the lanes and the perfect matchings in the regions $R(u)$ produce a polygon Q' . If the Euclidean diameter of each region $R(u)$ is at most δ , then the Fréchet distance between P and Q' is at most $\varepsilon + \delta$. Denote by $\Psi(P)$ the set of all simple polygons that can be constructed in this manner from a bar-signature for some $0 < \varepsilon < \varepsilon_0$.




Figure 26: Bar uv of a simple polygon obtained from an order compatible with the polygon shown in Fig. 2(c).

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

877 We show that, by reversing the steps of Algorithm `spur-elimination`(P', \mathcal{G}'), we can compute a bar-
 878 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),
 879 we can construct an ordering of the benchmark-to-benchmark paths of G_{uv} by concatenate the
 880 orders of \mathcal{B}^- and \mathcal{B}^+ (the sets of benchmark-to-benchmark paths of the resulting two groups).

881 If G_{uv} had spurs eliminated from u in some step 2b (cf. Section 5.2), we reverse each of the steps
 882 in the following manner. If a new group $G_{\ell-\ell+}$ was created, recall that every path $[t'_1; t'_2] \in \mathcal{B}'$ was
 883 created from two paths $[t_1; u]$ and $[t_2; u]$. Use the ordering of \mathcal{B}' to order the paths that originated
 884 \mathcal{B}' so that they form nested spurs, i.e., if $[t'_1; t'_2]$ is the topmost edge in \mathcal{B}' , $[t_1; u]$ (resp. $[t_2; u]$)
 885 should be the leftmost (resp., rightmost) path (without loss of generality, we use the orientation of
 886 Fig. 24). Identify the leftmost path in the segment that connects ℓ and its right child and place all
 887 nested paths that created $G_{\ell-\ell+}$ immediately to its left.

888 If one or more spurs were created at a node z in an adjacent group, we can find the position
 889 of the edges incident to each spur in the ordering of the adjacent group. Using this order, identify
 890 the first path in G_{uv} to the right of the edges incident to $[z]$. Then, immediately to the left of
 891 such a path, we can place the paths $[t_1; u; t_1]$ that generated the spurs at z . The relative order of
 892 these paths is the same as the one obtained by reversing a `spur-reduction`, described in the proof of
 893 Lemma 4, and therefore produces a simple polygon. If a path $[t_1; u; t_2]$ is simplified to $[t_1]$ (Step 1
 894 with $\min(t_1, t_2) = t_1$ without loss of generality; or Step 2), we can proceed analogously to the
 895 reversal of a `crimp-reduction` (cf. Lemma 1) from a path $[t_1; u; t_2; s]$ to $[t_1; s]$. Identify the path
 896 $[t_1; s]$ 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.

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

901 Updating the bar-signature when we reverse a `L-shortcut` operation is a bit more challenging.
 902 Determining the edge order in segments vw and vu_1 can be trivially done by just concatenating the
 903 order of merging segments. But phase (1) introduces a crimp in some cross-chains, and the reverse
 904 operation, `crimp-reduction`, may require nontrivial reordering in the bar-signature. Suppose that
 905 P' is obtained from P after a `crimp-reduction`. The proof of Lemma 1 shows a straightforward way
 906 to obtain a bar-signature of a polygon in $\Psi(P)$ given a polygon in $\Psi(P')$. However, obtaining a

907 bar-signature of $Q' \in \Psi(P')$ given $Q \in \Psi(P)$ requires identifying W_{top} and W_{bot} , which takes $O(n)$
 908 time.

909 In order to handle the reversal of phase (1) in $O(1)$ time, we divide the signature of each bar
 910 into pieces. Recall that the **bar-simplification** algorithm does not eliminate any cross chains from
 911 D_b , and when **bar-simplification** terminates, only the cross chains remain in the interior of D_b . Let
 912 K denote the set of cross chains of D_b . The bar-signature yields a linear order (from left to right)
 913 of K ; and the cross chains subdivide D_b into $|K| + 1$ regions. We maintain a linear order for the
 914 edges along the bar in each such region (including the boundary of the region), and denote the set
 915 of these edges in b by $E_1, \dots, E_{|K|+1}$.

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

923 When all operations in the bar simplification algorithm have been reversed, we have to combine
 924 the linear orders of $E_1, \dots, E_{|K|+1}$ into a total order, a common linear extension of these orders.
 925 The above-below relationship between the edges of each cross chain is uniquely determined by
 926 Lemma 2. Since the ordering of each subset guarantees that its paths can be realized without
 927 crossing, any linear extension of these ordering produces a bar-signature of a simple polygon.

928 **Pre processing.** The cluster formation and **new-bar-expansion** consist of subdivision operations
 929 that do not influence in the order of edges that define the bar-signature. If an edge $[a, c]$ in a bar
 930 b is subdivided into $[a, b, c]$, where $[b, c]$ is in D_b , we can assign $[a, c]$ to the same position of $[b, c]$
 931 in the ordering of edges in b . The **crimp-reduction** operations can be reversed by making the three
 932 edges that form a crimp consecutive in the ordering, as in the proof of Lemma 1. This completes
 933 the proof of Theorem 2.

934 7 Conclusion

935 We presented an $O(n \log n)$ -time algorithm for deciding whether a polygon with n vertices is
 936 weakly simple. There is a natural generalization for planar graphs [5][Appendix D]. We can
 937 define the *weak embedding* for graphs in terms of Fréchet distance. A graph $H = (V, E)$ can
 938 be considered a 1-dimensional simplicial complex. A *drawing* of H is a continuous maps of
 939 H to \mathbb{R}^2 . The Fréchet distance between two drawings P, Q of H is defined as $\text{dist}_F(P, Q) =$
 940 $\inf_{\phi: H \rightarrow H} \max_{x \in H} \text{dist}(P(\phi(x)), Q(x))$, where ϕ is an automorphism of H (a homeomorphism from
 941 H to itself). It is an open problem to decide efficiently whether a drawing of a graph H is weakly
 942 simple, i.e., whether a straight-line drawing P of H is within ε Fréchet distance from some embed-
 943 ding Q of H , for all $\varepsilon > 0$. An efficient algorithm has been announced [11] in the special case that
 944 the desired embedding is restricted to a given homotopy class (i.e., given combinatorial embedding).

945 We can also generalize the problem to higher dimensions. A polyhedron can be described as a
 946 map $\gamma : \mathbb{S}^2 \rightarrow \mathbb{R}^3$. A simple polyhedron is an injective function. A polyhedron P is weakly simple if
 947 there exists a simple polyhedron within ε Fréchet distance from P for all $\varepsilon > 0$. This problem can
 948 be reduced to origami flat foldability. The results of [4] imply that, given a convex polygon P and
 949 a piecewise isometric function $f : P \rightarrow \mathbb{R}^2$ (called *crease pattern*), it is NP-hard to decide if there

950 exists an injective embedding of P in three dimensions $\lambda : P \rightarrow \mathbb{R}^3$ within ε Fréchet distance from
 951 f for all $\varepsilon > 0$, i.e., if f is *flat foldable*. Given P and f , we can construct a continuous function
 952 $g : \mathbb{S}^2 \rightarrow P$ mapping each hemisphere of \mathbb{S}^2 to P ($g^{-1}(x)$, for a point $x \in P$, maps to two points
 953 in different hemispheres of \mathbb{S}^2). Then, the polyhedron $\gamma = g \circ f$ is weakly simple if and only if f is
 954 flat foldable. Therefore, it is also NP-hard to decide whether a polyhedron is weakly simple.

955 Finally it is an open problem to find a linear-time algorithm for recognizing weakly simple
 956 polygons. Chang et al. [5] conjectured that this is possible in the absence of spurs and forks.

957 **Acknowledgements.** Research by Akitaya, Aloupis, and Tóth was supported in part by the NSF
 958 awards CCF-1422311 and CCF-1423615. Research by Erickson was supported in part by the NSF
 959 award CCF-1408763. We thank Anika Rounds and Diane Souvaine for many helpful conversations
 960 that contributed to the completion of this project.

961 References

- 962 [1] Zachary Abel, Erik D. Demaine, Martin L. Demaine, David Eppstein, Anna Lubiw, and Ryuhei
 963 Uehara, Flat foldings of plane graphs with prescribed angles and edge lengths, in *Proc. 22nd*
 964 *Graph Drawing*, LNCS 8871, Springer, 2014, pp. 272–283.
- 965 [2] Esther M. Arkin, Michael A. Bender, Erik D. Demaine, Martin L. Demaine, Joseph S.B.
 966 Mitchell, Saurabh Sethia, and Steven S. Skiena, When can you fold a map?, *Computational*
 967 *Geometry: Theory and Applications* **29** (2004), 23–46.
- 968 [3] Michael A. Bender, Richard Cole, Erik D. Demaine, Martin Farach-Colton, and Jack Zito. Two
 969 simplified algorithms for maintaining order in a list, *Proc. 10th Annual European Symposium*
 970 *on Algorithms*, LNCS 2461, Springer, 2002, pp. 152–164.
- 971 [4] Marshall Bern and Barry Hayes, The complexity of flat origami, in *Proc. 7th ACM-SIAM*
 972 *Symposium on Discrete Algorithms*, SIAM, 1996, pp. 175–183.
- 973 [5] Hsien-Chih Chang, Jeff Erickson, and Chao Xu, Detecting weakly simple polygons, in *Proc.*
 974 *26th Symposium on Discrete Algorithm*, SIAM, 2015, pp. 1655–1670.
- 975 [6] Bernard Chazelle, Triangulating a simple polygon in linear time, *Discrete & Computational*
 976 *Geometry* **6** (1991), 485–524.
- 977 [7] Richard Cole and Ramesh Hariharan, Dynamic LCA queries on trees, *SIAM J. Comput.* **34** (4)
 978 (2005), 894–923.
- 979 [8] Robert Connelly, Erik D. Demaine, and Günter Rote, Infinitesimally locked self-touching link-
 980 ages with applications to locked trees, in *Physical Knots: Knotting, Linking, and Folding of*
 981 *Geometric Objects in \mathbb{R}^3* , American Mathematical Society, Providence, RI, 2002, pages 287–311.
- 982 [9] Pier Francesco Cortese, Giuseppe Di Battista, Maurizio Patrignani, and Maurizio Pizzonia, On
 983 embedding a cycle in a plane graph, *Discrete Mathematics* **309** (7) (2009), 1856–1869.
- 984 [10] Andrea Francke and Csaba D. Tóth, A census of plane graphs with polyline edges, in *Proc.*
 985 *30th Symposium on Computational Geometry*, ACM Press, 2014, pp. 242–250.
- 986 [11] Radoslav Fulek, Embedding graphs into embedded graphs, manuscript, [arXiv:1608.02087](https://arxiv.org/abs/1608.02087),
 987 2016.
- 988 [12] Branko Grünbaum, Polygons: Meister was right and Poinot was wrong but prevailed, *Beiträge*
 989 *zur Algebra und Geometrie* **53(1)** (2012), 57–71.

- 990 [13] Piotr Minc, Embedding of simplicial arcs into the plane, *Topology Proceedings*, **22** (1997),
991 305–340.
- 992 [14] Ares Ribó Mor, Realization and Counting Problems for Planar Structures: Trees and Linkages,
993 Polytopes and Polyominoes, Ph.D. thesis, Freie Universität Berlin, 2006.
- 994 [15] Michael Ian Shamos and Dan Hoey, Geometric intersection problems, in *Proc. 17th Symposium*
995 *on Foundations of Computer Science*, IEEE, 1976, pp. 208–215.
- 996 [16] Mikhail Skopenkov, On approximability by embeddings of cycles in the plane. *Topology and*
997 *its Applications* **134** (2003), 1–22.
- 998 [17] Daniel D. Sleator and Paul F. Dietz, Two algorithms for maintaining order in a list, in *Proc.*
999 *19th ACM Symposium on Theory of Computing*, ACM, 1987, pp. 365–372. Full version in Tech.
1000 Rep. CMU-CS-88-113, Carnegie Mellon University, 1988.