

The Left-Right Planarity Test

Ulrik Brandes^{*}

Department of Computer & Information Science, University of Konstanz

Abstract

A graph is planar if and only if it can be embedded in the plane without crossings. I give a detailed exposition of simple and efficient, yet poorly known algorithms for planarity testing, embedding, and Kuratowski subgraph extraction based on the left-right characterization of planarity.

Key words: Graph algorithms, planarity, algorithm review

1 Introduction

Two things appear to constitute the folklore about graph planarity testing:

- (1) There are two main strands of linear-time algorithms, the *vertex-addition approach* pioneered by Lempel, Even, and Cederbaum (1967), and the *path-addition approach* pioneered by Hopcroft and Tarjan (1974).
- (2) Both are a real challenge to understand, implement, and teach.

This is not a review of the exciting history of planarity testing at large, however, but of the lesser known *left-right approach*, which is seemingly different and usually associated with de Fraysseix and Rosenstiehl (1982). Even though the developments from its origins in Wu (1955) to its latest version in de Fraysseix, Ossona de Mendez, and Rosenstiehl (2006) and de Fraysseix (2008) is interesting in itself, my main goal here is to meet the apparent demand for an accessible exposition.

The left-right approach is remarkably elementary and does not require tricky data structures (e.g., Booth and Lueker 1976), a complicated embedding phase

^{*} I would like to thank Sabine Cornelsen, Giuseppe Di Battista, Daniel Kaiser, Martin Mader, and Maurizio Patrignani for helpful comments, suggestions, and corrections.

25 (e.g., Mehlhorn and Mutzel 1996), or even special treatment of biconnected
26 components. Moreover, it was found to be extremely fast (Boyer, Cortese,
27 Patrignani, and Di Battista, 2004) and can be augmented easily to return a
28 Kuratowski subgraph if the input is not planar.

29 This work is motivated by the stark contrast between the elegance and sim-
30 plicity of the left-right approach and its minimal adoption. It yields, I am
31 convinced, the simplest linear-time planarity algorithms known to date, but
32 to the best of my knowledge, there is not a single exposition or implementation
33 independent from the original group of authors.

34 The absence of an easily readable, yet fully detailed description may be the
35 main cause for its lack of popularity. In an attempt to remedy this situation,
36 the original description of de Fraysseix, Ossona de Mendez, and Rosenstiehl
37 (2006) is simplified with minor corrections, and it is extended by a new mo-
38 tivation, implementation-level pseudo-code, and more straightforward Kura-
39 towski subgraph extraction. While the planarity test given here differs from
40 the original paper, similar improvements have been introduced independently
41 into the only previous implementation, available in PIGALE (de Fraysseix and
42 Ossona de Mendez, 2002).

43 From the present description it should be possible to teach the algorithm in no
44 more than two sessions of an advanced algorithms course. With a planar graph
45 data structure art hand, transforming the pseudo-code into an implementation
46 should be a matter of hours.

47 The remainder is organized such that readers solely interested in understand-
48 ing the left-right approach can stop reading after Section 5. Therefore, only
49 minimal background on graph planarity and the associated algorithmic prob-
50 lems is provided in Section 2. A new motivation for the left-right approach is
51 given in Section 3, and the planarity characterization on which it is based in
52 Section 4. The left-right algorithm for planarity testing and planar embedding
53 is given in Section 5, including detailed pseudo-code. Kuratowski subgraph
54 extraction for non-planar graphs is treated separately in Section 6, and the
55 relation to other planarity criteria and algorithms as well as some notes on
56 the history of the left-right approach are postponed until Section 7.

57 **2 Planarity**

58 We consider simple undirected graphs $G = (V, E)$, since directions, loops and
59 multiple edges have no effect on planarity, and denote $n = n(G) = |V|$ and
60 $m = m(G) = |E|$ throughout.

61 A *drawing* of a graph is a mapping of its vertices onto points in the plane, and
62 of its edges onto curves connecting their endpoints. Where possible without
63 confusion, we neglect the distinction between vertices, edges, etc., and their
64 drawings. A drawing of a graph is *planar*, if edges do not intersect except at
65 common endpoints. A graph is planar, if it admits a planar drawing.

66 A planar drawing divides the plane into connected regions, called *faces*. Each
67 bounded face is an *inner* face, and the single unbounded one is called the
68 *outer* face.

69 A (*combinatorial*) *embedding* consists of cyclic orderings of the incident edges
70 for every vertex. An embedding is *realized* by a drawing, if the clockwise order-
71 ing of the edges around each vertex in the drawing agrees with the embedding.
72 Note that an embedding represents an equivalence class of drawings that re-
73 alize it. An embedding is planar, if it can be realized in a planar drawing.

74 Given a graph G , there are four major algorithmic problems related to pla-
75 narity:

- 76 (1) Decide whether G is planar.
- 77 (2) If G is planar, find a planar embedding.
- 78 (3) If G is not planar, find a Kuratowski subgraph.
- 79 (4) Given a planar embedding of G , realize it in a planar drawing.

80 A Kuratowski subgraph is an inclusion-minimal subgraph certifying non-planarity.
81 Since it is of less general interest, the topic is deferred to Section 6.

82 Our focus will be on the first two problems and we only note that, given
83 a planar embedding, realizations may be subject to various criteria such as
84 integer coordinates, straight-line edges, small area, polygonal edges with few
85 bends and/or slopes, etc., and there are many algorithms for drawing planar
86 graphs according to such criteria (see, e.g., Nishizeki and Rahman 2004).

87 The linear-time testing and embedding algorithm described in Section 5 is
88 based on a rather intuitive criterion that is motivated and established in the
89 next two sections.

90 **3 Motivation**

91 Since planarity is about the absence of crossings, cycles are the root cause of
92 difficulties: cycles yield closed curves that disconnect regions of the plane, and
93 one has to be careful about where to place which part of the graph.

94 There are only two significantly different ways to draw a simple cycle planarly,

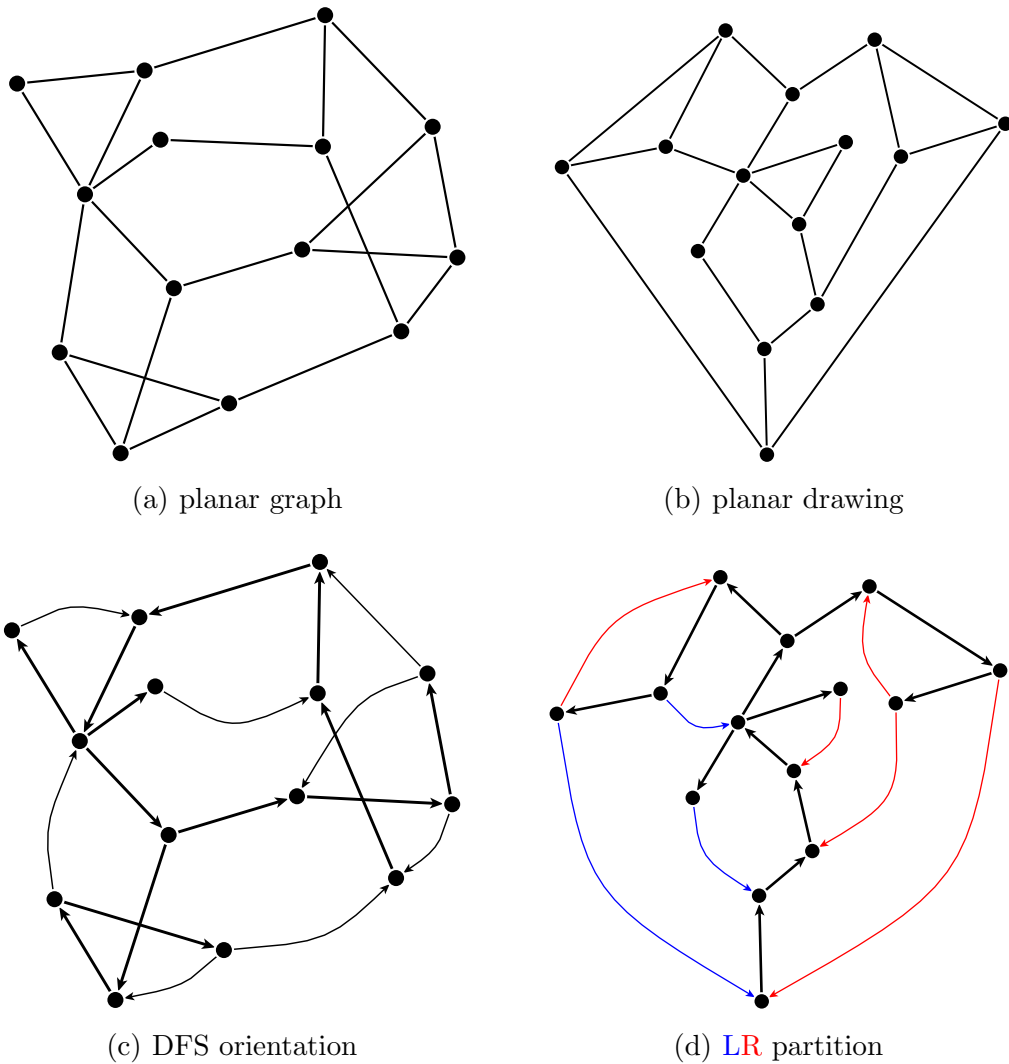


Fig. 1. Example of a planar graph (from Cai, Han, and Tarjan 1993). In both the planar and non-planar drawing, the same depth-first search (DFS) orientation is shown with thick tree edges and curved back edges. In any planar drawing the back edges can be partitioned into **left** and **right**, depending on whether their fundamental cycle is **counterclockwise** or **clockwise**. Note that the non-planar drawing contains self-intersecting fundamental cycles for both back edges entering the DFS root.

95 namely clockwise or counterclockwise. It turns out that fixing the orientation
 96 of some cycles may impose constraints on the choices for overlapping others
 97 via the ordering of edges around vertices. We will show that testing planarity
 98 amounts to deciding whether there is a consistent orientation of all cycles.
 99 Despite a potentially exponential number of cycles, this can be done efficiently,
 100 because constraints need not be resolved for all cycles, but only for a small
 101 set of cycles that represent the entire cycle structure.

102 Representative cycles are determined from a depth-first search as described
 103 next. This is followed by some apparent orientation constraints that also relate

104 cycle orientations to embeddings. In Section 4, more precise versions of these
 105 constraints are proven to characterize planarity. The proof is constructive and
 106 immediately yields a planar combinatorial embedding, if one exists.

107 3.1 Depth-first search

108 The left-right planarity criterion is inherently related to depth-first search
 109 (DFS). Important aspects of this relation are hinted at in this section, and
 110 DFS terminology is introduced along the way.

111 Recall that a depth-first search on a connected undirected graph $\bar{G} = (V, \bar{E})$
 112 yields a *DFS orientation* of \bar{G} , i.e. a directed graph $\vec{G} = (V, \vec{E})$ in which
 113 each undirected edge is oriented according to its traversal direction. Once
 114 the graph is oriented, we will only work with its directed version and hence
 115 neglect the distinction between \bar{E} and \vec{E} . In the oriented graph, we denote
 116 by $E^+(v) = \{(v, w) \in E\}$ the set of all outgoing edges of $v \in V$, so that
 117 $E = \bigcup_{v \in V} E^+(v)$.

118 In addition to an orientation, a DFS traversal yields a bipartition of the edges
 119 into $E = T \uplus B$, where those in T are called *tree edges* and induce a rooted
 120 spanning tree (the *DFS tree*), and the non-tree edges in B are called *back*
 121 *edges*. See Fig. 3. We write $u \rightarrow v$ and $v \hookleftarrow w$ for $(u, v) \in T$ and $(v, w) \in B$.
 122 Also, we use $\xrightarrow{+}$ for the transitive, and $\xrightarrow{*}$ for the reflexive and transitive closure
 123 of \rightarrow and call the unique sequence inducing $u \xrightarrow{*} v$ a *tree path*.

124 If $v \xrightarrow{*} w$ ($v \xrightarrow{+} w$), v is said to be (strictly) *lower* than w , and w (strictly)
 125 *higher* than v . A vertex is *lowest* (*highest*) in a set of vertices, if no other
 126 member of that set is lower (higher). The *height* of a vertex v is its distance
 127 from the root.

128 The characterizing property of DFS orientations is that the *target* w of every
 129 back edge $v \hookleftarrow w$ is a tree ancestor of (i.e., strictly below) its *source* v . Thus,
 130 each back edge $v \hookleftarrow w$ induces a *fundamental cycle* $C(v \hookleftarrow w) = w \xrightarrow{+} v \hookleftarrow$
 131 w , and these will be our primary objects of interest. Two cycles are called
 132 *overlapping*, if they share an edge, and it is the overlap of cycles that makes
 133 planarity testing challenging.

134 **Lemma 1** *Let $G = (V, T \uplus B)$ be a DFS-oriented graph.*

- 135 (1) *The fundamental cycles are exactly the simple directed cycles of G .*
 136 (2) *Two distinct fundamental cycles are either disjoint, or their intersection*
 137 *forms a tree path.*




Fig. 2. A fork with branching point v in the graph of Figure 1, and a sketched representation showing only those back edges that are return edges of $e = u \rightarrow v$. Note that edges to the lowpoint of e are dashed, and that e_2 is chordal but e_1 is not.

138 **PROOF.**

- 139 (1) All fundamental cycles are simple and, because of DFS, directed. Now
 140 consider any simple directed cycle and let $v \in V$ be lowest on that cycle.
 141 Since every cycle contains at least one back edge, let $x \hookrightarrow u$ be the first
 142 back edge after v . Vertex v is lowest, so that u must be in $v \xrightarrow{*} x$. Since
 143 the cycle is simple, $u = v$ and there are no more edges.
- 144 (2) Let $w \xrightarrow{*} v \hookrightarrow w$ and $u \xrightarrow{*} x \hookrightarrow u$ be two fundamental cycles. Since they
 145 are distinct, $v \hookrightarrow w \neq x \hookrightarrow u$. Since there is exactly one path between
 146 any pair of vertices in a tree, tree paths can join and fork at most once.
 147 A non-empty intersection of $w \xrightarrow{*} v$ and $u \xrightarrow{*} x$ must, therefore, be a
 148 treepath itself. \square

149 For two overlapping cycles, the last edge $u \rightarrow v$ on the shared tree path
 150 together with the succeeding edge $e_1 = (v, w_1), e_2 = (v, w_2)$ on each cycle is
 151 called their *fork*, and v its *branching point*. We will see that finding a planar
 152 combinatorial embedding reduces to finding an appropriate ordering of all
 153 triplets of edges that form a fork.

154 It will be convenient to fix a linearization of the cyclic ordering of outgoing
 155 edges around a vertex. Since every vertex v has at most one incoming tree
 156 edge, the clockwise order of outgoing edges is split at the incoming tree edge,
 157 or between any two consecutive outgoing edges if v is the root of a DFS tree.

158 In the next section, two simple observations help understand how cycle orien-
 159 tations impose fork orderings.

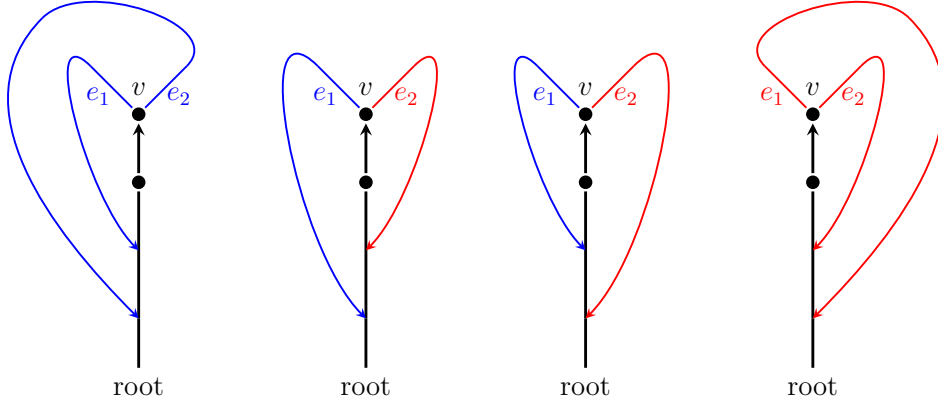


Fig. 3. In a planar drawing, overlapping fundamental cycles are nested, if and only if they are oriented alike. If the root is on the outer face, the lowest vertex of their union is contained in the outer of the two cycles.

160 *3.2 Orientation and nesting of fundamental cycles*

161 Recall that there are two classes of directed cycles in a planar drawing, because
 162 each is oriented either clockwise or counterclockwise. Since the intersection of
 163 overlapping fundamental cycles is a tree path containing at least one edge, the
 164 four possible configurations in Figure 3 can be summarized as follows.

165 **Observation 1** *In a planar drawing of a DFS-oriented graph $G = (V, T \uplus B)$,*
 166 *two overlapping cycles are nested, if and only if they are oriented alike.*

167 By assigning orientations we essentially determine whether the inside is to the
 168 left or to the right of a directed cycle, but the above observation does not
 169 specify which of two nested cycles is enclosed by the other.

170 For disambiguation we use the convention that roots of DFS trees are incident
 171 to the outer face and measure the nesting depth of a fundamental cycle using
 172 the following concepts.

173 The *return points* of a tree edge $v \rightarrow w \in T$ are the ancestors u of v with
 174 $u \xrightarrow{+} v \rightarrow w \xrightarrow{*} x \hookrightarrow u$ for some descendant x of w . A back edge $v \hookrightarrow w$ has
 175 exactly one return point, its target w . The return points of a vertex $v \in V$ are
 176 formed by the union of all return points of outgoing edges $(v, w) \in E^+(v) \subseteq$
 177 $T \uplus B$. A back edge $x \hookrightarrow u$ is a *return edge* for every tree edge $v \rightarrow w$ with
 178 $u \xrightarrow{+} v \rightarrow w \xrightarrow{*} x \hookrightarrow u$, and for itself.

179 The *lowpoint* of an edge is its lowest return point, if any, or its source if
 180 none exists. Note that the lowpoint of a back edge is the lowest vertex of its
 181 fundamental cycle, and therefore also called the lowpoint of that cycle.

182 The second important observation establishes nesting constraints induced by
183 lowpoints of cycles. It is justified by noting that if the root is on the outer
184 face and there is a proper tree path from the lowpoint of one cycle to that of
185 another cycle, this path can not be part of the inner cycle.

186 **Observation 2** *In a planar drawing of a DFS-oriented graph $G = (V, T \uplus B)$
187 with all roots of DFS trees on the outer face, overlapping fundamental cycles
188 are nested according to their lowpoint order.*

189 3.3 Relation to planar embeddings

190 The above two observations about orientations have immediate consequences
191 for planar embeddings. This becomes obvious by considering the single fork
192 in each of the four configurations in Figure 3.

193 First consider the fork of a pair of differently oriented cycles. Clearly, the
194 outgoing edge of the left cycle is before the outgoing edge of the right cycle in
195 the linearized order at branching point v .

196 Next consider the fork of a pair of likewise oriented cycles. In case they are right
197 cycles and one contains a vertex that is strictly lower than those in the other
198 cycle, the cycles outgoing edge (e_1 in Figure 3) comes first in the linearized
199 order at branching point v . The converse is true when v is the branching point
200 of left cycles.

201 A vertex may be the branching point for several pairs of overlapping cycles.
202 Combining both observations yields a (for now partial) embedding at branch-
203 ing points: outgoing edge of left cycles need to be before those of right cycles,
204 and the internal ordering in each subset is determined by lowpoints. Note that
205 there may be ties, and that outgoing tree edges may be part of several, differ-
206 ently oriented cycles. We will have to resolve these ambiguities, but otherwise
207 the whole approach rests entirely on Observations 1 and 2.

208 4 The Left-Right Planarity Criterion

209 With the above motivation in mind, we say that the *side* of a back edge in
210 a planar drawing is *right*, if its fundamental cycle is oriented clockwise, and
211 *left* otherwise. Assigning a side to a back edge thus corresponds to orienting
212 a fundamental cycle, and this will be all that needs to be done.

213 The following definition summarizes all constraints resulting from sets of over-
214 lapping fundamental cycles in terms of their respective back edge. It is worth




Fig. 4. LR constraints associated with $e = u \rightarrow v$.

215 noting that all constraints are generated by a single type of configuration
 216 associated with forks.

217 **Definition 2 (LR partition)** Let $G = (V, T \uplus B)$ be a DFS-oriented graph.
 218 A partition $B = L \uplus R$ of its back edges into two classes, referred to as left
 219 and right, is called left-right partition, or LR partition for short, if for every
 220 fork $u \rightarrow v \in T$ and $e_1, e_2 \in E^+(v)$

- 221 (1) all return edges of e_1 ending strictly higher than $\text{lowpt}(e_2)$
 222 belong to one class and
 223 (2) all return edges of e_2 ending strictly higher than $\text{lowpt}(e_1)$
 224 to the other.

225 The LR partition constraints are illustrated in Figure 4. Each group of con-
 226 straints is *associated* with a tree edge $u \rightarrow v$, and the system of constraints
 227 can be broken down into two sets of pairwise requirements: *same-constraints*
 228 forcing two back edges to be on the same side, and *different-constraints* forc-
 229 ing them to be on opposite sides. Note that two back edges are subject to
 230 a constraint only if their fundamental cycles overlap.¹ It is rather striking
 231 that these partition constraints (based on an arbitrary DFS orientation) are
 232 equivalent to planarity.

233 **Theorem 3 (Left-Right Planarity Criterion)** A graph is planar if and
 234 only if it admits an LR partition.

235 While necessity of the LR constraints is straightforward, we prove sufficiency
 236 in the next section by constructing a planar embedding from a given LR parti-
 237 tion. The construction is guided by the constraints that orientation and nesting
 238 of fundamental cycles impose on an embedding.

¹ The configurations inducing either type of constraint are considered in Section 6.

239 Removing the following ambiguity will simplify both argumentation and al-
240 gorithm. An LR partition is called *consistent*, if all back edges of a tree edge
241 that end at its lowpoint are on the same side.

242 **Lemma 4** *Any LR partition can be made consistent.*

243 **PROOF.** Consider two return edges b_1, b_2 of a tree edge $e = u \rightarrow v$ that end
244 at $\text{lowpt}(e)$. If one of them is involved in any LR constraint as specified in
245 Definition 2, this constraint must be associated with a tree edge $e' = u' \rightarrow v'$
246 such that $v' \xrightarrow{*} v$ and $\text{lowpt}(e')$ is strictly lower than $\text{lowpt}(e)$. Since b_1, b_2
247 originate from a common subtree entered by e and have the same return
248 point, actually both are involved in this constraint and even required to be on
249 the same side. Thus, consistency does not lead to contradictions. \square

250 4.1 Combinatorial embedding

251 Consider Figure 3 again and recall that the orientation of overlapping funda-
252 mental cycles induces a partial ordering of edges around forks.

253 We linearize clockwise cyclic orderings of edges around non-root vertices by
254 starting from the unique incoming tree edge. Outgoing edges belonging to
255 a counterclockwise cycle then need to appear before those belonging to a
256 clockwise cycle branching at the respective fork. Moreover, outgoing edges of
257 clockwise (counterclockwise) cycles must be ordered outside in (inside out)
258 around their branching point.

259 Given a DFS-oriented graph $G = (V, T \uplus B)$ together with an LR partition of
260 all back edges, a planar embedding can be obtained from extending the par-
261 tition to cover tree edges as well, and a linear order defined on the outgoing
262 edges of each vertex. This order will represent the nesting of cycles outside in
263 (with the root fixed to be on the outer face). It is used without modification as
264 the embedding order for right outgoing edges, and reversed for left outgoing
265 edges when flipping them to appear before any right edges. In an implemen-
266 tation, this can be realized by assigning its order rank to each edge, changing
267 the sign of left edges to minus, and a final sorting.

268 Extension of LR partitions to tree edges is straightforward. If a tree edge has
269 any return edges (i.e., its source is neither the root nor a cut vertex), it is
270 assigned to the same side as one of its return edges ending at the highest
271 return point (i.e., according to an innermost fundamental cycle it is part of).
272 Otherwise, the side is arbitrary.


273 To define a partial *nesting order* \prec , assume for a moment that all edges are

274 on the right side and consider a fork consisting of $u \rightarrow v$ and outgoing edges
 275 e_1, e_2 of v . If both have return edges, v is a branching point of overlapping
 276 fundamental cycles sharing $u \rightarrow v$. Since both cycles are clockwise for now,
 277 we must properly nest them to avoid edge crossings. Since we fixed the root
 278 of the DFS tree to be in the outer face, we have to define $e_1 \prec e_2$ if and
 279 only if the lowpoint of e_1 is strictly lower than that of e_2 . If both have the
 280 same lowpoint, but, say, only e_2 has another return point, we say that e_2 is
 281 *chordal* and let $e_1 \prec e_2$, because cycles containing e_2 and a return edge ending
 282 higher than $\text{lowpt}(e_2)$ can only lie inside of cycles containing e_1 and a return
 283 edge ending at $\text{lowpt}(e_1) = \text{lowpt}(e_2)$. If both e_1 and e_2 are chordal, the tie
 284 is broken arbitrarily, because eventually these two edges must be on different
 285 sides anyway.

286 In the planarity testing algorithm, \prec will be mimicked by defining the *nesting*
 287 *depth* of an edge e to be twice the height of the highest lowpoint of any cycle
 288 containing e , plus one if e is chordal.

289 The partial nesting order \prec is extended to a combinatorial embedding by
 290 *LR ordering*, i.e. by flip-reversing left edges before right ones and placing
 291 incoming back edges on the appropriate side of the tree edge leading into the
 292 subtree of their source. Some care is needed to avoid crossings of back edges,
 293 but we will see that, algorithmically, this embedding is almost trivial to realize.

294 **Definition 5 (LR Ordering)** *Given an LR partition, let $e_1^L \prec \dots \prec e_\ell^L$ be*
 295 *the left outgoing edges of a vertex v , and $e_1^R \prec \dots \prec e_r^R$ its right outgoing edges.*
 296 *If v is not the root, let u be its parent. The clockwise left-right ordering, or*
 297 *LR ordering for short, of the edges around v is defined as follows:*



299 where (u, v) is absent if v is the root, and $L(e)$ and $R(e)$ denote the left and
 300 right incoming back edges whose cycles share e . For two back edges $b_1 = x_1 \hookrightarrow$
 301 $v, b_2 = x_2 \hookrightarrow v \in R(e)$ let $z \rightarrow x, (x, y_1), (x, y_2)$ be the fork of $C(b_1)$ and $C(b_2)$.
 302 Then, b_1 comes after b_2 in $R(e)$ if and only if $(x, y_1) \prec (x, y_2)$. If $b_1, b_2 \in L(e)$,
 303 the order is reversed.

304 **Lemma 6** *Given any LR partition, LR ordering yields a planar embedding.*

305 **PROOF.** Let $G = (V, T \uplus B)$ be a DFS-oriented graph with an LR partition
 306 $B = L \uplus R$. We assume that it is consistent and extend it to also cover the tree
 307 edges as described above. Now consider the embedding defined by LR ordering




Fig. 5. Two types of crossings in proof of Lemma 6.

308 the edges around each vertex.

309 Since a graph with a spanning tree can always be drawn in such a way that
 310 a given embedding is respected, no two edges cross more than once, and none
 311 of the crossings involves a tree edge, the embedding is either planar, or any
 312 such drawing yields a simple crossing of two back edges (a crossing of more
 313 than two edges can be resolved into pairwise crossings). Only two cases are
 314 possible.

315 **Case 1:** (crossing back edges in same class)

316 Assume $x_1 \hookrightarrow u_1, x_2 \hookrightarrow u_2 \in R$ cross (the other case is symmetric). If
 317 $u_1 = u_2$, the crossings contradicts our definition of LR ordering the edges
 318 around that vertex.

319 W.l.o.g. we may therefore assume that u_1 is strictly higher than u_2 , and u_2
 320 therefore outside of the clockwise cycle $u_1 \xrightarrow{+} x_1 \hookrightarrow u_1$. Since the crossing is
 321 simple, x_2 in turn must be inside this cycle, and $u_1 \xrightarrow{+} x_1$ and $u_2 \xrightarrow{+} x_2$ cannot
 322 be disjoint (because we must enter the cycle somewhere along $u_2 \xrightarrow{+} x_2$).
 323 Let v be their highest common vertex, and e_1, e_2 the first edges on $v \xrightarrow{*} x_1$
 324 and $v \xrightarrow{*} x_2$.

325 Since x_2 is inside of the clockwise cycle, e_1 comes before e_2 in the order
 326 around v . On the other hand, the LR partition requires that all return edges
 327 of e_1 ending higher than u_2 are on the same side as $x_1 \hookrightarrow u_1$, so that also e_1
 328 is a right edge. LR ordering at v then implies that e_2 must be a right edge
 329 as well with $e_1 \prec e_2$.

330 By definition of \prec , either $lowpt(e_1)$ is strictly lower than u_2 , or $lowpt(e_1) =$
 331 $u_2 = lowpt(e_2)$ and e_2 is chordal as well. In the former case, $x_1 \hookrightarrow u_1$ and
 332 $x_2 \hookrightarrow u_2$ had to be assigned different sides. In the latter case, the highest
 333 ending return edge of e_2 is right as is e_2 , but conflicting with $x_1 \hookrightarrow u_1$ which
 334 is also right. In either case a contradiction.

335 **Case 2:** (crossing back edges in different classes)

336 Assume $x_1 \leftrightarrow u_1 \in R$ and $x_2 \leftrightarrow u_2 \in L$ (the other case is symmetric).
337 Since the crossing is simple, the tree paths $u_1 \xrightarrow{+} x_1$ and $u_2 \xrightarrow{+} x_2$ cannot be
338 disjoint and we define v, e_1, e_2 as in Case 1.

339 Again, e_1 must be before e_2 in the LR ordering of v for the back edges to
340 cross. If $u_1 = \text{lowpt}(e_1) = \text{lowpt}(e_2) = u_2$, the LR partition is not consistent.

341 Otherwise, we may assume that $\text{lowpt}(e_1)$ is strictly lower than u_2 (the
342 case that $\text{lowpt}(e_2)$ is strictly lower than u_1 is symmetric). Then, all return
343 edges of e_2 ending at u_2 or higher must be on the same side as $x_2 \leftrightarrow u_2 \in L$,
344 so that e_2 is left as well. Since e_1 comes before e_2 , it must also be left and
345 $e_2 \prec e_1$.

346 Due to the way we define sides for tree edges, e_1 is left only if it has a
347 left return edge ending strictly higher than $\text{lowpt}(e_1)$ (because it must end
348 at least as high as $x_1 \leftrightarrow u_1 \in R$ and the LR partition is consistent). On
349 the other hand, $e_2 \prec e_1$ implies that $\text{lowpt}(e_2)$ is lower than or equal to
350 $\text{lowpt}(e_1)$. This is a contradiction, since the LR constraints rule out that e_1
351 and e_2 have return edges ending strictly higher than $\text{lowpt}(e_2)$ and $\text{lowpt}(e_1)$
352 that are both on the left.

353 Since both types of crossings contradict our assumptions, the embedding is
354 planar. \square

355 We have thus proved constructively the non-obvious implication of the Left-
356 Right Planarity Criterion (Theorem 3).

357 5 Algorithm

358 We can now give the linear-time algorithm for testing planarity and for de-
359 termining a planar embedding or a minimal non-planar subgraph. After a
360 high-level description of its three main phases shown in Algorithm 1, full im-
361 plementation details are provided for all operations but those concerning the
362 specific data structure used to represent a graph and its embedding.

363 **Orientation.** The algorithm is based on the Left-Right Planarity Criterion
364 and therefore starts with a depth-first search (DFS) to orient the input graph.
365 For each connected component, the root of its spanning DFS tree is stored in
366 a list, *Roots*. The tree-path distance of a vertex from its root is stored in an
367 array *height*, so that roots of unexplored components are identified by still
368 having the initial value 0. Different from other planarity algorithms, there is
369 no need to worry about biconnected components.

variable	type	purpose	initially
<i>height</i>	integer node array	tree-path distance from root	∞
<i>lowpt</i>	integer edge array	<i>height</i> of lowest return point	n.a.
<i>lowpt2</i>	integer edge array	<i>height</i> of next-to-lowest return point (tree edges only)	n.a.
<i>nesting_depth</i>	integer edge array	proxy for nesting order \prec given by twice <i>lowpt</i> (plus 1 if chordal)	n.a.

(a) orientation phase

variable	type	interpretation	initially
<i>ref</i>	edge array of edges	edge relative to which side is defined	\perp
<i>side</i>	edge array of signs $\{-1, 1\}$	side of edge, or modifier for side of reference edge	1
$I = [low, high]$	pair of edges	interval of return edges represented by its two edges with extremal lowpoints	n.a.
$P = (L, R)$	pair of intervals	overlapping intervals, i.e., a conflict pair	n.a.
S	stack of conflict pairs	conflicting intervals formed by current return edges	\emptyset
<i>stack_bottom</i>	edge array of conflict pairs	top of stack S when traversing the edge (tree edges only)	n.a.
<i>lowpt_edge</i>	edge array of edges	next back edge in traversal (with lowest return point)	n.a.

(b) testing phase

variable	type	interpretation
<i>leftRef</i>	vertex array of edges	leftmost back edge from current DFS subtree (i.e. after next incoming left back edge)
<i>rightRef</i>	vertex array of edges	tree edge leading into current DFS subtree (i.e. before next incoming right back edge)

(c) embedding phase

Fig. 6. Main variables used in the algorithm.

Algorithm 1: Left-Right Planarity Algorithm

input: simple, undirected graph $G = (V, E)$ **output:** planar embedding (halts if graph is not planar)**if** $|E| > 3|V| - 6$ **then HALT: not planar****▼ orientation**

```
for  $s \in V$  do
  if  $height[s] = \infty$  then
     $height[s] \leftarrow 0$ ; append  $Roots \leftarrow s$ 
    DFS1( $s$ ) /* see Algorithm 2 */
```

▼ testing

```
sort adjacency lists according to non-decreasing  $nesting\_depth$ 
for  $s \in Roots$  do DFS2( $s$ ) /* see Algorithm 3 */
```

▼ embedding

```
for  $e \in E$  do  $nesting\_depth[e] = \underline{sign}(e) \cdot nesting\_depth[e]$ 
sort adjacency lists according to non-decreasing  $nesting\_depth$ 
for  $s \in Roots$  do DFS3( $s$ ) /* see Algorithm 6 */
```

where

integer sign(edge e)

```
if  $ref[e] \neq \perp$  then
   $side[e] \leftarrow side[e] \cdot \underline{sign}(ref[e])$ 
   $ref[e] \leftarrow \perp$ 
return  $side[e]$ 
```

370 During DFS, the partial nesting order \prec is determined by assigning to each
371 edge an integer value $nesting_depth$ such that $e_1 \prec e_2 \implies nesting_depth[e_1] <$
372 $nesting_depth[e_2]$.

373 **Testing.** To determine whether there exists a consistent LR partition, the
374 DFS forest is traversed for a second time. The traversal is modified, however,
375 such that outgoing edges are visited in the order induced by $nesting_depth$.
376 The second traversal halts if the graph is not planar, and we discuss at the
377 end of Section 5.2 how to extract one or more Kuratowski subgraphs in that
378 case.

379 The tentative side of edges may change often during the test, so that the
380 bipartition is returned only implicitly for efficiency reasons. An edge array ref
381 specifies for each edge a reference edge relative to which its side is defined,
382 and in an edge array $side$ a value of 1 or -1 indicates whether the side of
383 the edge is the same as, or different from, the side of its reference edge. If the
384 reference edge of e is undefined, i.e. $ref[e] = \perp$, the value of $side[e]$ specifies

385 the side directly, where -1 is for left and 1 is for right.

386 **Embedding.** Given an LR partition, flip-reversal of left edges is performed
387 by sorting the outgoing edges in all adjacency lists once again according to
388 their nesting order, though now modified by the side sign. Since the multi-
389 plication of *nesting_depth* with *side* only changes the sign of left edges to
390 negative, they are effectively placed before all right edges, in reverse order. To
391 complete the LR ordering, incoming edges are placed during a third traversal
392 of the DFS forest that is guided by the order of outgoing edges.

393 For each of the three main phases, we provide detailed pseudo-code with ample
394 comments in the subsequent sections.

395 5.1 Phase 1 – Orientation

396 The purpose of the first DFS is to orient the graph, and to determine lowpoints
397 and nesting order \prec . It is therefore a standard DFS computing the auxiliary
398 variables given in Table 6(a). Except for *height*, all are determined during
399 backtracking.

400 Our use of lowpoints is slightly non-standard in two ways. Firstly, we have
401 defined lowpoints for edges rather than vertices, and, secondly, we do not
402 assign DFS numbers, but heights. The latter induce the same ordering of
403 ancestors as DFS numbers, but are related to the tree more intuitively and in
404 general results in a smaller range of values which may in turn speed up the
405 subsequent sorting of adjacency lists according to *nesting_depth*.

406 Second lowpoints stored in *lowpt2* only serve to determine whether an edge
407 has more than one return point (i.e., it is chordal), and are not needed by
408 themselves.

409 The rationale for representing \prec via *nesting_depth* is two-fold: firstly, we can
410 apply a linear-time sorting algorithm such as counting sort on the set of edges,
411 because the range of values is linear in the size of the graph, and secondly, flip-
412 reversal of left edges after the second phase can be performed by sign changes
413 with subsequent re-sorting.

Algorithm 2: Phase 1 – DFS orientation and nesting order

DFS1(vertex v)

```
 $e \leftarrow \text{parent\_edge}[v]$ 
while there exists some non-oriented  $\{v, w\} \in E$  do
  orient  $\{v, w\}$  as  $(v, w)$ 
   $\text{lowpt}[(v, w)] \leftarrow \text{height}[v]$ ;  $\text{lowpt2}[(v, w)] \leftarrow \text{height}[v]$ 
  if  $\text{height}[w] = \infty$  then /* tree edge */
     $\text{parent\_edge}[w] \leftarrow (v, w)$ 
     $\text{height}[w] \leftarrow \text{height}[v] + 1$ 
    DFS1( $w$ )
  else /* back edge */
     $\text{lowpt}[(v, w)] \leftarrow \text{height}[w]$ 
    ▼ determine nesting depth
     $\text{nesting\_depth}[(v, w)] \leftarrow 2 \cdot \text{lowpt}[(v, w)]$ 
    if  $\text{lowpt2}[(v, w)] < \text{height}[v]$  then /* chordal */
       $\text{nesting\_depth}[(v, w)] \leftarrow \text{nesting\_depth}[(v, w)] + 1$ 
    ▼ update lowpoints of parent edge  $e$ 
    if  $e \neq \perp$  then
      if  $\text{lowpt}[(v, w)] < \text{lowpt}[e]$  then
         $\text{lowpt2}[e] \leftarrow \min\{\text{lowpt}[e], \text{lowpt2}[(v, w)]\}$ 
         $\text{lowpt}[e] \leftarrow \text{lowpt}[(v, w)]$ 
      else if  $\text{lowpt}[(v, w)] > \text{lowpt}[e]$  then
         $\text{lowpt2}[e] \leftarrow \min\{\text{lowpt2}[e], \text{lowpt}[(v, w)]\}$ 
      else
         $\text{lowpt2}[e] \leftarrow \min\{\text{lowpt2}[e], \text{lowpt2}[(v, w)]\}$ 
```

414 5.2 Testing

415 The second phase is the working horse of the algorithm. It determines a consis-
416 tent LR partition extended to all edges, if one exists; the code can be extended
417 to otherwise identify fundamental cycles whose union yields a Kuratowski sub-
418 graph as sketched in Section 6.

419 The main challenge is to detect and maintain all pairwise constraints among
420 back edges. Recall that constraints are associated with a tree edge and that
421 there are only two types of constraints. According to Definition 2, a pair of
422 back edges with overlapping fundamental cycles is placed on the same or on
423 different sides.

424 The following observation may help in building a better intuition. Consider a
425 signed *constraint graph*, in which vertices represent back edges of the original
426 DFS-oriented graph, and edges are introduced and labeled +1 or -1 when two

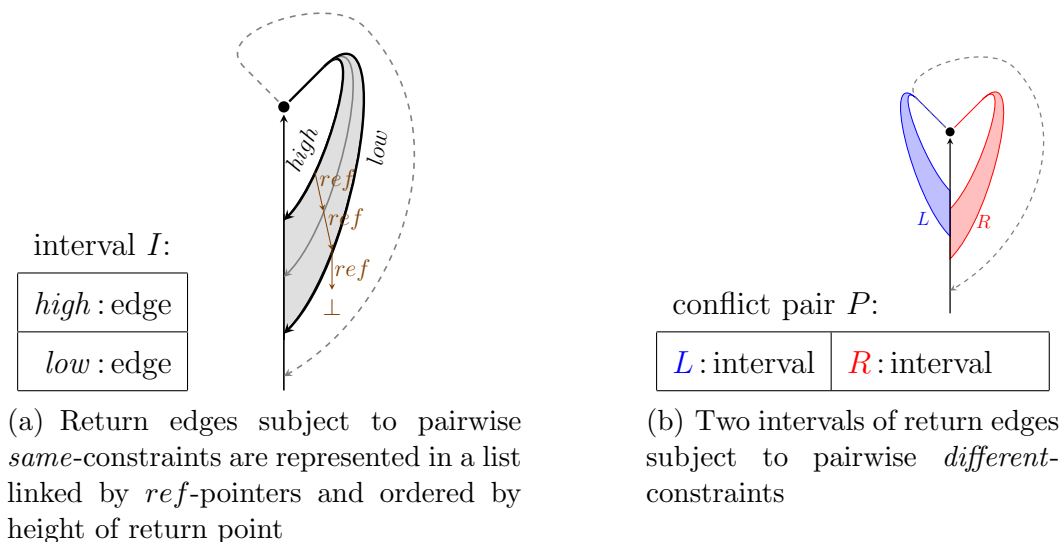


Fig. 7. The main data structure is a stack S storing overlapping pairs of intervals.

427 back edges are constrained to be on the same or on different sides. Finding an
 428 LR partition that satisfies all LR constraints is equivalent to testing whether
 429 this graph is *balanced*, i.e. there is a bipartition such that the corresponding
 430 cut is crossed exactly by the edges labeled -1 . Balancedness of signed graphs
 431 is introduced in Harary (1953).

432 Clearly, we cannot afford to maintain all pairwise constraints explicitly, be-
 433 cause their number may be quadratic in the size of the original graph. We will
 434 hence represent them implicitly to test for contradictions. To be able to real-
 435 ize a bipartition it is sufficient to maintain a spanning forest of the constraint
 436 graph. We therefore construct a rooted tree for each of its component using a
 437 reference pointer *ref* for every edge. This is including tree edges, since their
 438 side is determined by reference to a return edge ending at the highest return
 439 point, anyway.

440 A second array, *side*, is used to store the side of all edges that are roots
 441 in the spanning forest of the constraint graph. For all other edges the array
 442 holds the sign of the unique outgoing constraint-graph edge linking them to
 443 their corresponding reference edge. As indicated earlier, values $+1$ and -1 will
 444 therefore be interpreted either as *right* and *left*, or as *same* and *different*.

445 To reduce the number of constraints that need to be represented explicitly,
 446 observe that the *same*-constraints induced by a fork $u \rightarrow v$, $e_1, e_2 \in E^+(v)$
 447 in Definition 2 involve two sets of return edges with a simple structure. For,
 448 say, e_1 let $h = x_h \leftrightarrow u_h$ and $\ell = x_\ell \leftrightarrow u_\ell$ be the two (possibly equal) return
 449 edges ending at the highest and lowest return point of e_1 that is strictly higher
 450 than $lowpt(e_2)$. Then we know that h_ℓ and all return edges $x' \leftrightarrow u'$ of e_1
 451 with a return point in $u_\ell \xrightarrow{*} u_h$ are in the same group of *same*-constraints. This

452 *interval* of edges can thus be represented by its two extreme members, h and ℓ ,
453 as shown in Figure 7(a). Return edges belonging to an interval are maintained
454 in a singly-linked list, from highest to lowest return point, using the *ref*-array.

455 *Different*-constraints can be summarized similarly, because by transitivity
456 they always involve all pairs of edges in a pair of intervals. A *conflict pair*
457 therefore consists of two intervals of edges subject to *different*-constraints as
458 shown in Figure 7(b). It represents their tentative assignment to the left and
459 right, and thus a partial bipartition.

460 The second DFS traversal is designed to build the entire bipartition of edges
461 incrementally by merging conflict pairs. Its main data structure is a stack S
462 of conflict pairs representing all constraints associated with a tree edge that
463 has been traversed, but not yet backtracked over. Note that these constraints
464 involve only back edges that have already been traversed, but return to a
465 vertex below the current one. In other words, each back edge in the stack is a
466 return edge for at least one tree edge in the current DFS path.

467 By processing DFS trees bottom-up, the constraints associated with an edge
468 can be determined by merging those associated with its outgoing edges. Two
469 main invariants are maintained. We will not prove them explicitly, but rather
470 let them serve as an orientation for understanding the implementation. The
471 first invariant eventually yields correctness of the implementation,

472 **Partitioning Invariant:** The additional conflict pairs accumulated
at the top of the stack between traversing a tree edge and backtrack-
ing over it represent a partial bipartition satisfying all non-crossing
constraints associated with that edge.

473 and the second one ensures that constraint merging can be carried out effi-
474 ciently.

475 **Ordering Invariant:** Return edges forming an interval are repre-
sented in singly-linked lists ordered from highest to lowest return point,
and the lowest edge in a conflict pair is not lower than the highest edge
in another conflict pair deeper in S .

476 5.2.1 Ordered traversal

477 Pseudo-code for the second DFS is given in Algorithms 3–5. Since all edges
478 have been oriented during the first DFS, they are again traversed in the same
479 direction. The traversal order differs, though, since adjacency lists have been
480 rearranged according to *nesting_depth*, so that outgoing edges with lower low-
481 points are traversed first. This reordering is crucial for the ordering invariant.

Algorithm 3: Phase 2 – Testing for LR partition

```

DFS2(vertex  $v$ )
   $e \leftarrow \text{parent\_edge}[v]$ 
  for  $e_i \in E^+(v) = \langle e_1, \dots, e_d \rangle$  do /* ordered by nesting_depth */
     $\text{stack\_bottom}[e_i] \leftarrow \text{top}(S)$ 
    if  $e_i = \text{parent\_edge}[\text{target}(e_i)]$  then /* tree edge */
      |  $\text{DFS2}(\text{target}(e_i))$ 
    else /* back edge */
      |  $\text{lowpt\_edge}[e_i] \leftarrow e_i$ ;  $\text{push}(\emptyset, [e_i, e_i]) \rightarrow S$ 
    if  $\text{lowpt}[e_i] < \text{height}[v]$  then /*  $e_i$  has return edge */
      | if  $e_i = e_1$  then
        | |  $\text{lowpt\_edge}[e] \leftarrow \text{lowpt\_edge}[e_1]$ 
      | else
        | |  $\blacktriangleright$  add constraints of  $e_i$  (Algorithm 4)
    if  $e \neq \perp$  then /*  $v$  is not root */
      |  $u \leftarrow \text{source}(e)$ 
      |  $\blacktriangleright$  trim back edges ending at parent  $u$  (Algorithm 5)
      |  $\blacktriangledown$  side of  $e$  is side of a highest return edge
      | | if  $\text{lowpt}[e] < \text{height}[u]$  then /*  $e$  has return edge */
        | | |  $h_L \leftarrow \text{top}(S).L.\text{high}$ ;  $h_R \leftarrow \text{top}(S).R.\text{high}$ 
        | | | if  $h_L \neq \perp$  and  $(h_R = \perp \text{ or } \text{lowpt}[h_L] > \text{lowpt}[h_R])$  then
          | | | |  $\text{ref}[e] \leftarrow h_L$ 
        | | | else
          | | | |  $\text{ref}[e] \leftarrow h_R$ 

```

482 When visiting a vertex v during the DFS traversal, the high-level task is to
 483 recursively determine the constraints for all outgoing edges and integrate them
 484 into those associated with parent edge $e = u \rightarrow v$ (if v is not a DFS root).

485 Before traversing an outgoing edge $e_i \in E^+(v)$, we therefore remember the top
 486 conflict pair $\text{stack_bottom}[e_i]$ on S (where $\text{top}(S) = \perp$ if S is empty). If e_i was a
 487 tree edge in the first traversal, all constraints associated with e_i are recursively
 488 determined and pushed onto S . If e_i is a back edge, it is pushed onto S in a
 489 conflict pair of its own because it may be involved in later constraints. Recall
 490 that our goal is to determine a *consistent* LR partition. We therefore store in
 491 an edge array lowpt_edge the first back edge not traversed earlier. For edges
 492 that have return edges, this is the first return edge to their lowpoint and can
 493 thus be used as a reference for other return edges that have to be assigned to
 494 the same side to meet the consistency requirement. A back edge e_i is its own
 495 unique return edge to its lowpoint so that we let $\text{lowpt_edge}[e_i] = e_i$.

496 From the partitioning invariant we know that when returning from the traver-

497 sal of e_i , the conflict pairs above $stack_bottom[e_i]$ represent a partial LR parti-
 498 tion of all return edges of e_i . While processing the first outgoing edge e_1 we sim-
 499 ply leave them on the stack, if any, and pass on $lowpt_edge[e_1]$ to $lowpt_edge[e]$.
 500 Note that, since e_1 has a return edge, $parent_edge[v] = e \neq \perp$, i.e. v is not
 501 a root. For each of the other outgoing edges $e_i = e_2, \dots, e_d \in E^+(v)$, the
 502 constraints above $stack_bottom[e_i]$ are merged into those which have already
 503 been accumulated for e and are directly beneath in S . Constraint integration
 504 is the most essential step and described separately in Algorithm 4 and below.

505 After all outgoing edges have been traversed, we trim all those back edges
 506 from the top of S that are return edges of some $e_i \in E^+(v)$, but not of e ,
 507 i.e. which end at u . This requires some annoyingly lengthy but simple case
 508 distinctions given in Algorithm 5 and explained below. Observe that, if v is
 509 a DFS root, then there is no parent edge $e = u \rightarrow v$, but there are also no
 510 remaining constraint pairs on S , since a DFS root does not have outgoing back
 511 edges and there is more than one outgoing tree edge only if each leads into a
 512 different biconnected component.

513 If existent, parent edge e is finally assigned to the side of a back ending at the
 514 highest return point as suggested by the LR ordering procedure of Section 4.1.
 515 By the ordering invariant, this edge is the highest return edge in one of the two
 516 intervals in the top conflict pair, and we have already removed all non-return
 517 edges. Observe that the stack cannot be empty if there is a return edge.

518 5.2.2 Adding constraints associated with the next outgoing edge

519 We have to merge all constraints associated with the next outgoing edge, e_i ,
 520 with those already accumulated from e_1, \dots, e_{i-1} . The involved intervals are
 521 therefore gathered one by one in an initially empty conflict pair P as illustrated
 522 in Figure 8.

523 **Merge return edges of e_i into right interval.** All return edges of e_i have
 524 been traversed since traversing e_i , and they are represented in the top conflict
 525 pairs on stack S down to, but not including, $stack_bottom[e_i]$. All of these
 526 intervals have to be merged on one side because of the fundamental cycle of
 527 $lowpt_edge[e]$. If there is a conflict pair with two non-empty intervals, merging
 528 on one side violates an earlier constraint and the graph is not planar.

529 There is at least one conflict pair above $stack_bottom[e_i]$ for otherwise we
 530 would not have entered this section. The non-empty interval of each such pair
 531 is merged in the right interval $P.R$ of P without changing their order by having
 532 the lowest edge of $P.R$ refer to the highest edge of the next conflict pair and
 533 replacing it accordingly. An exception is the interval containing a return edge

Algorithm 4: Adding constraints associated with e_i (part of Alg. 3)

```
▼ add constraints of  $e_i$ 
   $P \leftarrow (\emptyset, \emptyset)$ 
  ▼ merge return edges of  $e_i$  into  $P.R$ 
    repeat
       $Q \leftarrow \text{pop}(S)$ 
      if  $Q.L \neq \emptyset$  then swap  $Q.L, Q.R$ 
      if  $Q.L \neq \emptyset$  then
        | HALT: not planar
      else
        | if  $\text{lowpt}[Q.R.\text{low}] > \text{lowpt}[e]$  then /* merge intervals */
        |   | if  $P.R = \emptyset$  then  $P.R.\text{high} \leftarrow Q.R.\text{high}$ 
        |   | else  $\text{ref}[P.R.\text{low}] \leftarrow Q.R.\text{high}$ 
        |   |  $P.R.\text{low} \leftarrow Q.R.\text{low}$ 
        | else /* make consistent */
        |   |  $\text{ref}[Q.R.\text{low}] \leftarrow \text{lowpt\_edge}[e]$ 
      until  $\text{top}(S) = \text{stack\_bottom}[e_i]$ 
  ▼ merge conflicting return edges of  $e_1, \dots, e_{i-1}$  into  $P.L$ 
    while  $\text{conflicting}(\text{top}(S).L, e_i)$  or  $\text{conflicting}(\text{top}(S).R, e_i)$  do
       $Q \leftarrow \text{pop}(S)$ 
      if  $\text{conflicting}(Q.R, e_i)$  then swap  $Q.L, Q.R$ 
      if  $\text{conflicting}(Q.R, e_i)$  then
        | HALT: not planar
      else /* merge interval below  $\text{lowpt}(e_i)$  into  $P.R$  */
        |  $\text{ref}[P.R.\text{low}] \leftarrow Q.R.\text{high}$ 
        | if  $Q.R.\text{low} \neq \perp$  then  $P.R.\text{low} \leftarrow Q.R.\text{low}$ 
      if  $P.L = \emptyset$  then  $P.L.\text{high} \leftarrow Q.L.\text{high}$ 
      else  $\text{ref}[P.L.\text{low}] \leftarrow Q.L.\text{high}$ 
       $P.L.\text{low} \leftarrow Q.L.\text{low}$ 
    if  $P \neq (\emptyset, \emptyset)$  then push  $P \rightarrow S$ 
```

where

```
boolean conflicting(interval  $I$ , edge  $b$ )
  | return ( $I \neq \emptyset$  and  $\text{lowpt}[I.\text{high}] > \text{lowpt}[b]$ )
```

534 to the lowpoint of e ; to make the LR partition consistent, we make it refer to
535 the *lowpt_edge* directly.

536 **Merge conflicting return edges of e_1, \dots, e_{i-1} into left interval.** Re-
537 turn edges of e_1, \dots, e_{i-1} with lowpoints higher than $\text{lowpt}[e_i]$ are subject
538 to pairwise *same*-constraints and to a *different*-constraint with respect to
539 some return edge of e_i . (If $\text{lowpt}[e_i] = \text{lowpt}[e]$ this is not *lowpt_edge* $[e_i]$ but,

540 e.g., a back edge returning to $lowpt2[e_i]$ which must exist, because apparently
 541 $lowpt2[e_{i-1}]$ exists as well by the way outgoing edges are ordered).

542 So while there are conflict pairs on the stack that contain return edges with
 543 lowpoints higher than $lowpt[e_i]$, these have to be merged on one side.

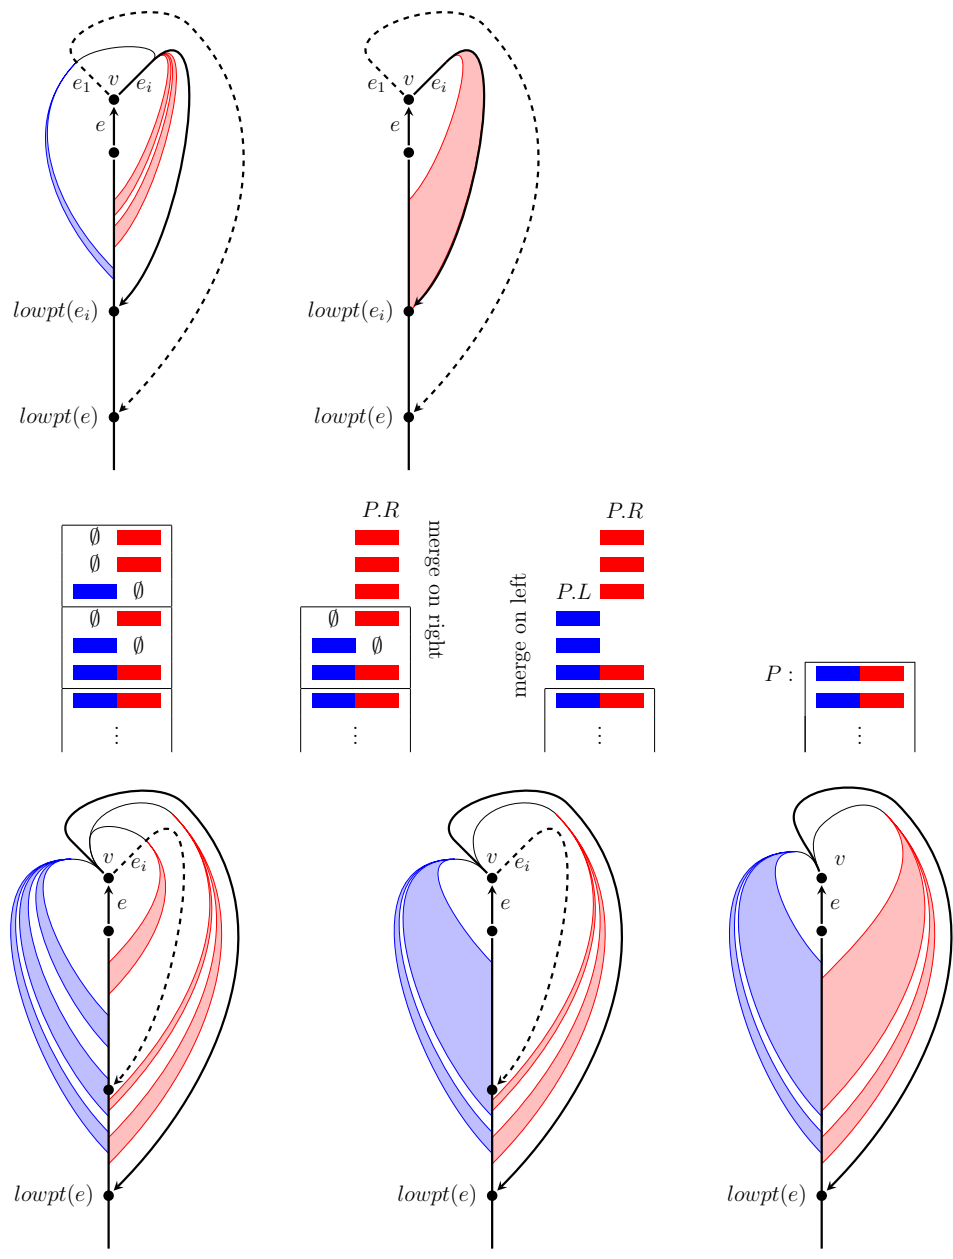


Fig. 8. In the core step of the algorithm, the constraints of e_i are merged into those of e_1, \dots, e_{i-1} . Horizontal lines indicate where the top of stack S is divided by $stack_bottom[e_i]$ and the topmost pair that is not conflicting with $lowpt_edge[e_i]$. If $lowpt(e_i) = lowpt(e)$, the pair containing only $lowpt_edge[e_i]$ is discarded and the bipartition is made consistent by assigning $ref[lowpt(e_i)] \leftarrow lowpt(e)$. Note that in this case, $P.R$ is not growing when merging on the left.

544 a pair contains two intervals ending above $lowpt[e_i]$, we again have a contra-
 545 diction with a previous constraint and thus non-planarity. If only one side
 546 ends above $lowpt[e_i]$, we merge the other into $P.R$ (effectively closing these
 547 constraints under transitivity).

548 The actual merging of intervals is performed in the same way as above, and
 549 the final pair can be placed on the stack.

550 5.2.3 Trimming back edges

551 The purpose of Algorithm 5 is to remove all those back edges from conflict
 552 pairs on the stack that have the parent of the current tree edge $e = u \rightarrow v$ as
 553 their lowpoint, because they are no return edges of e or any lower tree edge,
 554 and therefore not subject to any constraint associated with a tree edge still to
 555 be processed.

Algorithm 5: Removing back edges ending at parent u (part of Alg. 3)

▼ trim back edges ending at parent u

▼ drop entire conflict pairs

while $S \neq \emptyset$ and $lowest(top(S)) = height[u]$ do

$P \leftarrow pop(S)$

 if $P.L.low \neq \perp$ then $side[P.L.low] \leftarrow -1$

if $S \neq \emptyset$ then /* one more conflict pair to consider */

$P \leftarrow pop(S)$

 ▼ trim left interval

 while $P.L.high \neq \perp$ and $target(P.L.high) = u$ do

$P.L.high \leftarrow ref[P.L.high]$

 if $P.L.high = \perp$ and $P.L.low \neq \perp$ then /* just emptied */

$ref[P.L.low] \leftarrow P.R.low$

$side[P.L.low] \leftarrow -1$

$P.L.low \leftarrow \perp$

 ▶ trim right interval

 push $P \rightarrow S$

where

integer lowest(conflictpair P)

 if $P.L = \emptyset$ then return $lowpt[P.R.low]$

 if $P.R = \emptyset$ then return $lowpt[P.L.low]$

 return $\min\{lowpt[P.L.low], lowpt[P.R.low]\}$

556 **Dropping entire conflict pairs.** If the lowest lowpoint on either side of a
 557 conflict pair P is the source of the current tree edge $u \rightarrow v$, all lowpoints of

558 back edges in P are the same and the edges will not be involved in any future
559 constraints. The pair is finalized by assigning the lowest back edge of the left
560 interval to the left side. Since $side$ is initialized with 1, the lowest back edge
561 in the right interval $P.R$ is already assigned correctly to the right side, and all
562 other back edges b in P to the same side as $ref[b]$.

563 **Trimming a left interval.** Since back edges in an interval are concatenated
564 by ref -pointers in an order monotonic in the $height$ of their lowpoints, we
565 can simply remove back edges from the upper end of the left interval until
566 the highest lowpoint is no longer u , or the interval has become empty. In the
567 latter case the lower end of the interval is still defined and made to refer to
568 an edge on the other side, setting its $side$ to -1 accordingly. Note that the
569 right interval cannot be empty for otherwise the entire conflict pair had been
570 removed in the first while loop. All other removed back edges still refer to a
571 back edge on the same side, so that the initial 1 of their $side$ -entry must not
572 be changed.

573 **Trimming a right interval.** This is symmetric to the previous operation.
574 Note, however, that the assigned $side$ in case the right interval becomes empty
575 is -1 as well, because this indicates that the side of the lowest back edge is
576 different from the side of the lowest back edge in the left interval. Again, the
577 left interval cannot be empty.

578 **Assigning a side to a tree edge.** After trimming all back edges ending
579 at source u of the current tree edge $e = u \rightarrow v$ in Algorithm 3, the side of
580 e is determined by reference to a highest return edge. There is a return edge
581 only if $lowpt[e] < height[u]$. Otherwise the u is a cutvertex or root and it
582 does not matter, which side e is assigned to. Since the existence of a return
583 edge renders S non-empty, the ordering invariant asserts that we have to
584 compare the lowpoints of the highest back edges in the two intervals of the
585 top constraint pair on S (checking for existence).

586 At the end of the testing phase, a non-crossing LR partition is given implicitly
587 by edge arrays ref and $side$, if and only if the graph is planar. These define the
588 side of an edge e relative to another, where $side[e]$ indicates whether the side
589 is the same or different from that of $ref[e]$. Since $ref[e]$ always has a strictly
590 lower target than e , the referrals are acyclic and form a rooted spanning forest
591 of the constraint graph. The roots of that forest refer to \perp , and their side is
592 determined explicitly by $side$. After dereferencing all referrals at the beginning
593 of the embedding phase, the LR partition is known explicitly.




Fig. 9. The algorithm testing $K_{3,3}$ and K_5 for planarity. In either case, the status before starting the second DFS is depicted in the middle, and the algorithm halts in the configuration on the right while processing e .

594 Two small examples are shown in Figure 9. Even though both graphs are non-
 595 planar, the workings of the algorithm are nicely illustrated, since coloring and
 596 embedding correspond to the current (implicitly represented) bipartition and
 597 LR ordering.

598 5.3 Embedding

599 Compared to other planarity algorithms, the embedding phase is extremely
 600 simple. LR ordering the outgoing edges of the DFS-oriented graph is achieved
 601 by sorting them according to their *nesting_depth* on both sides. It is known
 602 that a partial embedding like this already determines a complete combinatorial
 603 embedding (see, e.g., Cai 1993), but for completeness we provide full details
 604 in Algorithm 6.

605 The DFS forest is traversed for the third time. Since outgoing edges are already
 606 ordered in the desired way, back edges are encountered exactly as required in
 607 the definition of LR ordering. As described in Table 6(c) we therefore maintain,
 608 for each vertex v , the two positions next to which the next left or right back
 609 edge needs to be inserted.

610 Observe that the incoming back edges from the same subtree actually appear
 611 in counterclockwise order. If the data structure available for embedded graphs
 612 does not provide a constant-time method for direct neighbor insertion, the
 613 now obsolete array *ref* can be used to build a singly-linked list of all edges
 614 incident to a vertex in counterclockwise order.

Algorithm 6: Phase 3 – Embedding

DFS3(vertex v)

```
  for  $e_i \in E^+(v) = \langle e_1, \dots, e_d \rangle$  do
     $w \leftarrow \text{target}(e_i)$ 
    if  $e_i = \text{parent\_edge}[w]$  then /* tree edge */
      make  $e_i$  first edge in adjacency list of  $w$ 
       $\text{leftRef}[v] \leftarrow e_i$ ;  $\text{rightRef}[v] \leftarrow e_i$ 
      DFS3( $w$ )
    else /* back edge */
      if  $\text{side}[e_i] = 1$  then
        place  $e_i$  directly behind  $\text{rightRef}[w]$  in adjacency list of  $w$ 
      else
        place  $e_i$  directly before  $\text{leftRef}[w]$  in adjacency list of  $w$ 
         $\text{leftRef}[w] \leftarrow e_i$ 
```

615 5.4 Running time and implementation

616 **Theorem 7** Algorithm 1 can be implemented to test in $\mathcal{O}(n)$ time whether a
617 graph is planar and return a planar combinatorial embedding if it is.

618 **PROOF.** We have argued throughout this section that the algorithm cor-
619 rectly yields an LR ordering if the graph admits an LR partition. Hence,
620 correctness is established by the Left-Right Planarity Criterion (Theorem 3).
621 Since a graph cannot be planar if $m > 3n - 6$, we may assume that the number
622 of edges is at most linear in the number of vertices.

623 The algorithm performs three DFS traversals, and rearranges the edges twice
624 in between. Both rearrangements are obtained from sorting the edges accord-
625 ing to *nesting_depth*, which can be done in linear time using bucket sort
626 because all entries are integers with absolute value less than $2n$.

627 The first DFS clearly requires constant time per edge traversal and backtrack-
628 ing step, and hence linear time overall.

629 During the second traversal, every back edge is pushed onto the stack ex-
630 actly once (when it is traversed), so that the number of newly generated con-
631 straint pairs is bounded by the number of back edges. If more than a constant
632 number of constraint pairs is inspected during the addition of constraints, a
633 corresponding number of them is merged. Since also the total time spent on
634 trimming back edges that return to the parent is linear in the number of edges,
635 the overall running time is linear.

636 Dereferencing *ref*-pointers takes linear time, because it is performed only once
637 before the third DFS traversal, which also requires linear time if the graph
638 data structure provides a constant-time operation to move an edge next to
639 another in the embedding order. If it does not, the algorithm can be altered
640 to re-use *ref*-pointers for the embedding as described in Section 5.3. \square

641 The left-right approach can be implemented as described above² and our expe-
642 riences with its performance essentially confirm the results of Boyer, Cortese,
643 Patrignani, and Di Battista (2004). A special edge numbering scheme used
644 in PIGALE (de Fraysseix and Ossona de Mendez, 2002) serves to avoid re-
645 peated DFS traversals, but we have found the sorting of adjacency lists to
646 be even more costly. Note, however, that both sorting and DFS traversal can
647 be avoided during the testing phase by splitting the stack into singly-linked
648 lists associated with edges and processing edges (i.e., merging their final list
649 of constraints into that of another edge) in the order given by *nesting_depth*.
650 This order is determined by creating two buckets for each *height* and adding
651 an edge to its respective bucket when its *lowpt* is known during the initial
652 DFS, i.e. when it is backtracked over. Since *lowpt* is determined bottom-up,
653 edges added to the same bucket end up being in the desired order.

654 6 Non-Planarity

655 The algorithm halts, if and only if the input graph is non-planar. The initial
656 test on the number of edges is justified by the fact that in any planar graph,
657 $m \leq 3n - 6$ (a well-known consequence of Euler's formula; see, e.g., Chiba and
658 Nishizeki 1988).

659 If the algorithm does not halt, a planar embedding constructed as described
660 in Section. 5.3 can serve as a certificate for planarity. If it does halt after
661 the initial condition, though, the graph is non-planar and we would like a
662 certificate for this case as well. The earliest characterization of planarity is
663 due to Kuratowski (1930) and provides for exactly that. It is described in the
664 following.

665 Recall the two non-planar graphs in Figure 9 and let us call a *subdivision* of a
666 graph G any graph that can be obtained from G by repeatedly splitting edges
667 using new vertices of degree two.

668 **Theorem 8 (Kuratowski's Planarity Criterion)** *A graph is planar, if and*
669 *only if it does not contain a subdivision of $K_{3,3}$ or K_5 .*

² The version described here has been implemented almost literally in C/C++
using LEDA by Daniel Kaiser, and in Java using yFiles by Martin Mader.

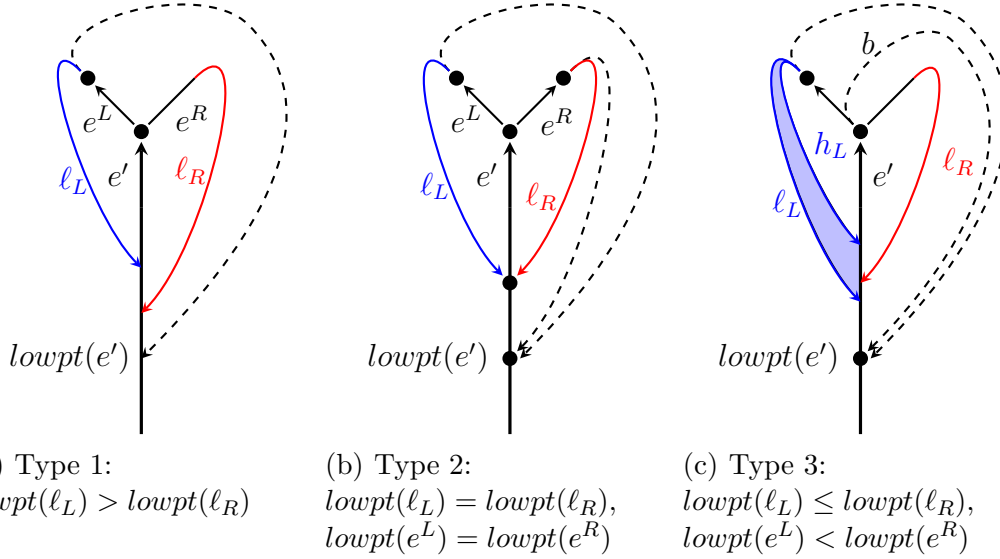


Fig. 10. A *different*-constraint between the two lowest edges ℓ_L, ℓ_R of a new constraint pair associated with e' is caused by one of three inclusion-minimal configurations (assuming w.l.o.g. that $lowpt(e^L) \leq lowpt(e^R)$).

670 Each subgraph that is a subdivision of $K_{3,3}$ or K_5 is called a *Kuratowski sub-*
 671 *graph* (or *planarity obstruction*), and we show how to extract such a subgraph
 672 if the planarity test of the previous section fails.

673 Algorithm 1 halts when a pair $P = (L, R)$ at the top of S contains two
 674 conflicting intervals $L \neq \emptyset \neq R$ that prevent merging on the right or left
 675 side. In other words, a *same*-constraint associated with the current tree edge
 676 e contradicts a previously found *different*-constraint involving the same two
 677 intervals. Let $\ell_L = P.L.low$ and $P.R.low = \ell_R$ be the lowest edges in L and R .
 678 A Kuratowski subgraph can be determined from the union of the fundamental
 679 cycles of ℓ_L, ℓ_R and at most four others.

680 ***Different*-constraint.** Since both ℓ_L, ℓ_R are return edges of e , their funda-
 681 mental cycles are overlapping. Let e', e^L, e^R be the respective fork as shown
 682 in Figure 10, and let $lowpt(e^L) \leq lowpt(e^R)$ (otherwise exchange names). The
 683 *different*-constraint between ℓ_L, ℓ_R is introduced while backtracking over e' ,
 684 and we have $lowpt(e') < lowpt(\ell_R)$ because $lowpt(e') = lowpt(\ell_R)$ implies
 685 that ℓ_R is not involved in a constraint associated with e' . We distinguish three
 686 types of *different*-constraints (depicted in Figure 10) and show how to extract
 687 an inclusion-minimal subgraph preserving them.

688 **Type 1:** If $lowpt[\ell_L] > lowpt[\ell_R]$, then $lowpt_edge[e^L] \neq \ell_L$ since $lowpt(e^L) <$
 689 $lowpt(\ell_R) \leq lowpt(\ell_L)$. Moreover, $lowpt(\ell_R) > lowpt(e^L) = lowpt(e')$, be-
 690 cause if $lowpt(e^L) > lowpt(e')$, then ℓ_L and $lowpt_edge[e^L]$ are subject to a
 691 *same*-constraint and thus ℓ_L is not the lowest edge in $P.L$. Hence, the funda-

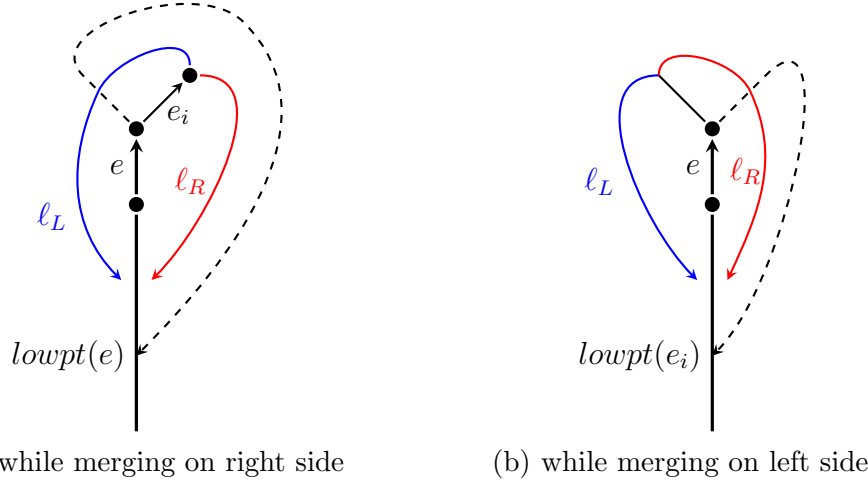


Fig. 11. One of these inclusion-minimal configurations is present when a subsequent *same*-constraint on ℓ_L, ℓ_R causes the planarity test to halt during constraint merging.

692 mental cycles of ℓ_L, ℓ_R and $lowpt_edge[e^L]$ preserve the constraint forcing ℓ_L
 693 and ℓ_R on different sides; another cycle is added due to the *same*-constraint
 694 to form a subdivision of $K_{3,3}$.

695 **Type 2:** If $lowpt[\ell_L] = lowpt[\ell_R]$ and $lowpt(e^L) = lowpt(e^R)$, arguments sym-
 696 metric to Case 1 imply that $\ell_R \neq lowpt_edge[e^R]$ and that, in addition to the
 697 three cycles above, $C(lowpt_edge[e^R])$ is needed to preserve the constraint.
 698 The overlap of $C(\ell_L)$ and $C(\ell_R)$ is removed to obtain a subdivision of $K_{3,3}$
 699 when adding the final cycle.

700 **Type 3:** The final case is $lowpt[\ell_L] \leq lowpt[\ell_R]$ and $lowpt(e^L) < lowpt(e^R)$,
 701 since $lowpt[e^L] = lowpt[e^R]$ yields Case 2 if $lowpt(\ell_L) = lowpt(\ell_R)$, or Case 1
 702 with the roles of L and R exchanged if $lowpt[\ell_L] < lowpt[\ell_R]$. It follows
 703 that $\ell_R = lowpt_edge[e^R]$ for otherwise both would be subject to a *same*-
 704 constraint and ℓ_R would not be lowest in $P.R$. Hence, ℓ_L is in indirect
 705 conflict with ℓ_R via an edge $h_L \neq \ell_L$ which we can choose to be the highest
 706 in $P.L$. For ℓ_L and h_L to be forced into an interval, however, there must
 707 have been another back edge returning to $lowpt(e')$ (a back edge to a return
 708 point between $lowpt(e')$ and $lowpt(\ell_L)$ would replace ℓ_R as lowest in $P.R$).
 709 In addition to the four fundamental cycles of $\ell_L, \ell_R, lowpt_edge[e']$ and h_L
 710 we therefore include $C(b)$, where b is the last edge that was made consistent
 711 with $lowpt_edge[e']$. With the additional cycle added later due to the *same*-
 712 constraint we obtain a K_5 minor, i.e. a subgraph that can be turned into K_5
 713 by contracting edges. Note that every K_5 minor that is not a subdivision of
 714 K_5 contains a subdivision of $K_{3,3}$.

715 All conditions distinguishing the three types can be tested in constant time
 716 using information obtained during the testing phase.

717 **Same-constraint.** There are only two straightforward cases illustrated in
718 Figure 11. Either the constraint is induced by the fundamental cycle of $lowpt_edge[e]$
719 while merging the constraints associated with e_i on the right, or by $lowpt_edge[e_i]$
720 while merging on the left. In either case only one cycle needs to be added to
721 the subgraph extracted for the *different*-constraint.

722 The fundamental cycles forming a Kuratowski subgraph are extracted easily
723 by traversing the DFS tree from the source of each of the at most six critical
724 back edges downwards using *parent_edge* pointers. Forks of overlapping cycles
725 are found along the way, and only the cycle with lower lowpoint needs to
726 be continued. The time needed to extract a Kuratowski subgraph after the
727 planarity test determined a constraint violation is thus proportional to the
728 size of the union of at most six fundamental cycles.

729 Note that the failure configurations in Figure 9 are the prototypical combi-
730 nation of Type 1 and Type 3 *different*-constraints with a subsequent *same*-
731 constraint that initiates merging on the left.

732 7 Discussion

733 We have reviewed the Left-Right Planarity Criterion (Theorem 3) and de-
734 scribed a simple linear-time algorithm (Algorithm 1) based on it. While this
735 is not a review of graph planarity, and many important references and devel-
736 opments are left out, some notes on closely related work seem in place.

737 7.1 Characterization

738 Historically, Kuratowski (1930) provides the first characterization of graph
739 planarity. Given the discussion in Section 6, this characterization in terms of
740 minimal forbidden subgraphs can be re-interpreted as identifying the cycle
741 structures of $K_{3,3}$ and K_5 as the two minimal overlap configurations that
742 prevent planar drawing.

743 Among various later characterizations, the seemingly most closely related cri-
744 terion is due to Mac Lane (1937), because it is also formulated in terms of
745 a representative set of cycles. Consider the set of all undirected cycles of a
746 graph, and define the sum of two cycles as the symmetric difference of their
747 edge sets. These together form a vector space, called *cycle space*. A *basis*
748 of the cycle space is a minimum-cardinality set of cycles such that every cycle is
749 the sum of some basis cycles.

750 **Theorem 9 (Mac Lane’s Planarity Criterion)** *A graph is planar, if and*
751 *only if it has a cycle basis in which every edge appears at most twice.*

752 For a better intuition, consider a planar drawing of a connected planar graph.
753 Traversing each face in the drawing (say, inner faces clockwise, the outer face
754 counterclockwise) yields the set of (directed) *facial cycles* forming a basis of
755 the cycle space. As required, every edge is traversed exactly twice (once in
756 each direction).

757 Any cycle basis for a graph G has cardinality $\mu(G) = m - n + \kappa(G)$, where
758 $\kappa(G)$ is the number of connected components of G and $\mu(G)$ is called the
759 *cyclomatic number* of G . This is exactly the number of non-tree edges of a
760 spanning forest and, in fact, the fundamental cycles of any spanning forest
761 induce a cycle basis.

762 The left-right criterion thus also asks for a cycle basis with a special property,
763 namely that its elements, the (directed) fundamental cycles of a DFS orienta-
764 tion, can be bipartitioned such that all constraints associated with forks are
765 satisfied.

766 The cycle bases considered in these criteria are therefore maximally distinct.
767 While the basis cycles in Mac Lane’s criterion are as different as possible (with
768 each edge in at most two cycles), the basis cycles in the left-right criterion are
769 as concentrated as possible (with their overlap forming a spanning forest).

770 7.2 Development

771 The earliest precursor of the left-right approach is a planarity characterization
772 of Wu (1955), which states that a graph is planar, if and only if a certain
773 system of linear equations has a solution. It was complemented by the concept
774 of crossing chains in Tutte (1970), and refined to Boolean variables and fewer
775 equations in the 1970s (see Wu 1985, 1986; Liu 1990). The variables in this
776 smaller system are associated with the edges, and the equations represent
777 constraints generated from configurations of overlapping cycles obtained from
778 a spanning DFS forest. An alternative interpretation of the existence of a
779 solution is that of balancing a constraint graph as in Section 5.2. Rosenstiehl
780 (1980) gives an algebraic proof for this characterization.

781 This work was further developed in several papers, but the descriptions are
782 rather incomplete, in particular with respect to linear-time implementation
783 (de Fraysseix and Rosenstiehl, 1982, 1985; Xu, 1989; Cai, Han, and Tarjan,
784 1993).

785 Finally, de Fraysseix, Ossona de Mendez, and Rosenstiehl (2006) simplified
786 the approach even further by concentrating on the single constraint-inducing
787 configuration of Definition 2. While this paper is still incomplete and difficult
788 to read, the linear-time implementation is described in just enough detail to
789 provide a basis for replication. Among the differences to the present description
790 is the maintenance and merging of constraints, since intervals are described as
791 stacks rather than their extreme pairs of edges and there is a constraint stack
792 for each edge rather than our global stack S . It turns out, however, that the
793 most recent implementation in PIGALE (de Fraysseix and Ossona de Mendez,
794 2002) uses a similar representations.

795 The characterization of Kuratowski subgraphs in terms of configurations along
796 a DFS spanning tree given in de Fraysseix and Ossona de Mendez (2003) and
797 de Fraysseix (2008) leads to a linear-time extraction algorithm associated with
798 the left-right approach. To the best of my knowledge, the version given here
799 is original, and it is more directly based on the left-right characterization.

800 7.3 Algorithms

801 The first published polynomial-time planarity testing algorithm is due to Aus-
802 lander and Parter (1961). It is based on an observation already noted above,
803 namely that in a planar drawing of a graph every simple cycle forms a closed
804 curve partitioning the plane into an inside and an outside region.

805 Consider the graph obtained by removing the edges of some simple cycle, but
806 retaining copies of vertices on the cycle for every incident non-cycle edge. These
807 vertices are called *attachments* and the connected components of the resulting
808 graph are called *segments*. Clearly, each segment must be drawn completely
809 inside or completely outside of the removed cycle, but a pair of segments must
810 not be placed in the same region if their attachments interleave on the cycle.
811 Planarity can thus be tested by recursively choosing cycles and sides. The
812 related algorithm of Demoucron, Malgrange, and Pertuiset (1964) also starts
813 from a simple cycle, but then iteratively chooses a path that can be added into
814 one of the current faces. The algorithm is not only simple, but also has the
815 unusual property to eagerly maintain a partial embedding that is not changed
816 later on. Both algorithms require $\Omega(n^2)$ time, though.

817 In a graph-algorithmic milestone, the first linear-time planarity test was pre-
818 sented by Hopcroft and Tarjan (1974). Their approach is called *path-addition*
819 because it refines that of Auslander and Parter (1961) by adding paths rather
820 than segments, and in an order determined from a depth-first search of the
821 graph. It took many years, though, until finally Mehlhorn and Mutzel (1996)
822 complemented the algorithm with an $\mathcal{O}(n)$ embedding phase.

823 Recall how we observed in Section 3 that likewise-oriented cycles are nested
824 if they overlap. Maybe because de Fraysseix and Rosenstiehl (1982) phrased
825 the notions of left and right in terms of angles with the DFS tree rather
826 than orientations of fundamental cycles, it has gone almost unnoticed that
827 the left-right approach is yet another refinement of Auslander and Parter
828 (1961) and Hopcroft and Tarjan (1974), progressing from segments to paths
829 to edges. Together with Canfield and Williamson (1990) and Haeupler and
830 Tarjan (2008) this observation instills hope that there may be a useful and
831 elegant unification of path- and vertex-addition approaches including the two
832 most efficient versions of de Fraysseix, Ossona de Mendez, and Rosenstiehl
833 (2006) and Boyer and Myrvold (2004).

834 References

- 835 Auslander, L., Parter, S. V., 1961. On imbedding graphs in the sphere. *Journal*
836 *of Mathematics and Mechanics* 10 (3), 517–523.
- 837 Booth, K. S., Lueker, G. S., 1976. Testing for the consecutive ones property,
838 interval graphs, and graph planarity using PQ-tree algorithms. *Journal of*
839 *Computer and System Sciences* 13, 335–379.
- 840 Boyer, J. M., Cortese, P.-F., Patrignani, M., Di Battista, G., 2004. Stop mind-
841 ing your P’s and Q’s: Implementing a fast and simple DFS-based planarity
842 testing and embedding algorithm. In: Liotta, G. (Ed.), *Proc. Intl. Symp.*
843 *Graph Drawing (GD ’03)*. Vol. 2912 of LNCS. Springer-Verlag, pp. 25–36.
- 844 Boyer, J. M., Myrvold, W. J., 2004. On the cutting edge: Simplified $\mathcal{O}(n)$
845 planarity by edge additon. *Journal of Graph Algorithms and Applications*
846 8 (3), 241–273.
- 847 Cai, J., 1993. Counting embeddings of planar graphs using DFS trees. *SIAM*
848 *Journal on Discrete Mathematics* 6 (3), 335–352.
- 849 Cai, J., Han, X., Tarjan, R. E., 1993. An $\mathcal{O}(m \log n)$ -time algorithm for the
850 maximal planar subgraph problem. *SIAM Journal on Computing* 22 (6),
851 1142–1162.
- 852 Canfield, E. R., Williamson, S. G., 1990. The two basic linear time planarity
853 algorithms: Are they the same? *Linear and Multilinear Algebra* 26, 243–265.
- 854 Chiba, N., Nishizeki, T., 1988. *Planar Graphs: Theory and Algorithms*. North-
855 Holland.
- 856 de Fraysseix, H., 2008. Trémaux trees and planarity. *Electronic Notes in Dis-*
857 *crete Mathematics* 31, 169–180.
- 858 de Fraysseix, H., Ossona de Mendez, P., 2002. Pigale: Public implemen-
859 tation of a graph algorithm library and editor, software project at
860 pigale.sourceforge.net (GPL License).
- 861 de Fraysseix, H., Ossona de Mendez, P., 2003. On cotree-critical and DFS
862 cotree-critical graphs. *Journal of Graph Algorithms and Applications* 7 (4),
863 411–427.

- 864 de Fraysseix, H., Ossona de Mendez, P., Rosenstiehl, P., 2006. Trémaux trees
865 and planarity. *International Journal of Foundations of Computer Science*
866 17 (5), 1017–1029.
- 867 de Fraysseix, H., Rosenstiehl, P., 1982. A depth-first characterization of pla-
868 narity. *Annals of Discrete Mathematics* 13, 75–80.
- 869 de Fraysseix, H., Rosenstiehl, P., 1985. A characterization of planar graphs by
870 Trémaux orders. *Combinatorica* 5 (2), 127–135.
- 871 Demoucron, G., Malgrange, Y., Pertuiset, R., 1964. Graphes planaires: Recon-
872 naissance et construction de représentations planaires topologiques. *Revue*
873 *Fran çais Recherche Opérationnelle* 8 (30), 33–47.
- 874 Haeupler, B., Tarjan, R. E., 2008. Planarity algorithms via PQ -trees. *Elec-*
875 *tronic Notes in Discrete Mathematics* 31, 143–149.
- 876 Harary, F., 1953. On the notion of balance of a signed graph. *Michigan Math-*
877 *ematical Journal* 2, 143–146 (with addendum).
- 878 Hopcroft, J. E., Tarjan, R. E., 1974. Efficient planarity testing. *Journal of the*
879 *ACM* 21 (4), 549–568.
- 880 Kuratowski, K., 1930. Sur le problème des courbes gauches en Topologie. *Fun-*
881 *damenta Mathematicae* 15, 271–283.
- 882 Lempel, A., Even, S., Cederbaum, I., 1967. An algorithm for planarity testing
883 of graphs. In: Rosenstiehl, P. (Ed.), *Proc. Intl. Symp. Theory of Graphs*
884 (Rome, July 1966). Gordon and Breach, pp. 215–232.
- 885 Liu, Y., 1990. A Boolean characterization of planarity and planar embeddings
886 of graphs. *Annals of Operations Research* 24, 165–174.
- 887 Mac Lane, S., 1937. A combinatorial condition for planar graphs. *Fundamenta*
888 *Mathematicae* 28, 22–32.
- 889 Mehlhorn, K., Mutzel, P., 1996. On the embedding phase of the hopcroft and
890 tarjan planarity testing algorithm. *Algorithmica* 16 (2), 233–242.
- 891 Nishizeki, T., Rahman, M. S., 2004. *Planar Graph Drawing*. Vol. 12 of *Lecture*
892 *Notes Series on Computing*. World Scientific.
- 893 Rosenstiehl, P., 1980. Preuve algébrique du critère de planarite de Wu-Liu.
894 *Annals of Discrete Mathematics* 9, 67–78.
- 895 Tutte, W. T., 1970. Toward a theory of crossing numbers. *Journal of Combi-*
896 *natorial Theory* 8, 45–53.
- 897 Wu, W., 1955. On the realization of complexes in Euclidean space I. *Acta*
898 *Mathematica Sinica* 5, 505–552, English version in *American Mathematical*
899 *Society Translations, Series 2* 78:137–184, 1968.
- 900 Wu, W., 1985. On the planar imbedding of linear graphs. *Journal of System*
901 *Sciences and Mathematical Sciences* 5 (4), 290–302.
- 902 Wu, W., 1986. On the planar imbedding of linear graphs (continued). *Journal*
903 *of System Sciences and Mathematical Sciences* 6 (1), 23–35.
- 904 Xu, W., 1989. An improved algorithm for planarity testing based on Wu-Liu’s
905 criterion. *Annals of the New York Academy of Sciences* 576, 641–652.