

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 drawn in the plane without crossings. I give a detailed exposition of a simple and efficient, yet poorly known algorithm for planarity testing and embedding based on the left-right characterization of planarity. This is complemented by a new Kuratowski subgraph extraction algorithm.

Key words: Graph algorithms, planarity testing, planar embedding, Kuratowski subgraphs

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 article is not a review of the exciting history of planarity testing at large, however, but of the lesser known *left-right approach*, which is generally considered to be different from the above and associated with de Fraysseix and Rosenstiehl (1982). And even though the development from its origins in Wu (1955) to its latest version in de Fraysseix, Ossona de Mendez, and Rosenstiehl (2006) and de Fraysseix (2008) appears to be another interesting story, my only goal here is to meet the apparent demand for an accessible

^{*} I would like to thank Sabine Cornelsen, Giuseppe Di Battista, Bernd Gaertner, Daniel Kaiser, Martin Mader, and Maurizio Patrignani for helpful comments, suggestions, and corrections. I am particularly grateful to an anonymous reviewer who pointed out a gap in the first version of the Kuratowski subgraph extraction method.

24 exposition. As it turns out, the left-right approach simplifies and improves
25 upon the approach of Hopcroft and Tarjan (1974).

26 The left-right approach is remarkably elementary and does not require tricky
27 data structures (e.g., Booth and Lueker 1976), a complicated embedding phase
28 (e.g., Mehlhorn and Mutzel 1996), or even special treatment of biconnected
29 components. Moreover, it was found to be extremely fast (Boyer, Cortese,
30 Patrignani, and Di Battista, 2004) and can be augmented easily to return a
31 Kuratowski subgraph if the input is not planar.

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

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

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

50 The remainder is organized such that readers solely interested in understand-
51 ing the left-right approach can stop reading after Section 6. Therefore, only
52 minimal background on graph planarity and the associated algorithmic prob-
53 lems is provided in Section 2. A new motivation for the left-right approach is
54 given in Section 3, and the planarity characterization on which it is based in
55 Section 4. We derive a simple polynomial-time planarity test in Section 5, be-
56 fore giving the linear-time left-right algorithm for planarity testing and planar
57 embedding in Section 6, including detailed pseudo-code. Kuratowski subgraph
58 extraction for non-planar graphs is treated separately in Section 7, and the
59 relation to other planarity criteria and algorithms as well as some notes on
60 the history of the left-right approach conclude the paper in Section 8.

61 2 Planarity

62 We consider simple undirected graphs $G = (V, E)$, since directions, loops and
63 multiple edges have no effect on planarity, and denote $n = n(G) = |V|$ and
64 $m = m(G) = |E| \leq \frac{n(n-1)}{2}$ throughout. Bollobás (1998) and Diestel (2005) are
65 excellent textbooks on graph theory.

66 A *drawing* of a graph is a mapping of its vertices onto points in the plane (or,
67 equivalently, the surface of a sphere), and its edges onto curves connecting
68 their endpoints. Where possible without confusion, we neglect the distinction
69 between vertices, edges, etc., and their drawings. A drawing of a graph is
70 *planar*, if edges do not intersect except at common endpoints. A graph is
71 planar, if it admits a planar drawing.

72 A planar drawing divides the plane into connected regions, called *faces*. Each
73 bounded face is an *inner* face, and the single unbounded one is called the *outer*
74 face. We use $f = f(G)$ to denote the number of faces in a planar drawing of
75 G . The following classic result assures, in particular, that the number of faces
76 $f = f(G)$ is the same for every planar drawing of G . For a proof see, e.g.,
77 Aigner and Ziegler (2009, Chapter 12).

78 **Theorem 1 (Euler's Formula)** *For connected planar graphs, $n - m + f = 2$.*

79 As an immediate consequence, we can rule out graphs that are too dense to
80 be planar from further consideration.

81 **Corollary 2** *For planar graphs with $n > 2$ vertices, $m \leq 3n - 6$.*

82 **PROOF.** In a planar drawing of a planar graph, an edge is in the boundary
83 of at most two faces, while every face has at least three edges. It follows that
84 $f \leq \frac{2m}{3}$ and hence $m \leq 3n - 6$ by Euler's Formula. \square

85 A (*combinatorial*) *embedding* consists of cyclic orderings of the incident edges
86 at each vertex. An embedding is *realized* by a drawing, if the clockwise order-
87 ing of the edges around each vertex in the drawing agrees with the embedding.
88 Note that an embedding represents an equivalence class of drawings that re-
89 alize it. An embedding is planar, if it can be realized by a planar drawing.

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

- 92 (1) Decide whether G is planar.
- 93 (2) If G is planar, determine a planar embedding.

- 94 (3) If G is not planar, determine a Kuratowski subgraph.
95 (4) Given a planar embedding of G , determine a planar drawing realizing it.

96 A Kuratowski subgraph is an inclusion-minimal non-planar subgraph. Since
97 the problem of determining a Kuratowski subgraph it is of less general interest,
98 the topic is deferred to Section 7.

99 Moreover, we will not consider the fourth problem. Note, however, that real-
100 izations of a given planar embedding may be subject to various criteria such as
101 integer coordinates, straight-line edges, small area, polygonal edges with few
102 bends and/or slopes, etc., and there are many algorithms for drawing planar
103 graphs according to such criteria (see, e.g., Nishizeki and Rahman 2004).

104 The linear-time testing and embedding algorithm described in Section 6 is
105 based on a rather intuitive criterion that is motivated and established in the
106 next two sections, respectively.

107 **3 Motivation**

108 Since planarity is about the absence of crossings, cycles are the root cause of
109 difficulties: cycles yield closed curves that disconnect regions of the plane, so
110 that care has to be taken which other parts of the graph are placed inside or
111 outside such cycles.

112 There are only two significantly different ways to draw a simple cycle pla-
113 narily, namely clockwise or counterclockwise. Fixing an orientation, however,
114 may impose constraints on the orientation of other, overlapping cycles via
115 the ordering of edges around vertices. In fact, testing planarity amounts to
116 deciding whether there is a consistent simultaneous orientation of all cycles.
117 Despite a potentially exponential number of cycles, this can be done efficiently,
118 because we will see that constraints need to be resolved only for a small set
119 of cycles representing the entire cycle structure.

120 This small set of representative cycles is determined from a depth-first search
121 as described next. We then motivate how apparent orientation constraints can
122 be used to relate cycle orientations to embeddings. In Section 4, this is made
123 more precise in order to characterize planar graphs via cycle orientations.
124 The proof is constructive and yields a planar combinatorial embedding, if one
125 exists.

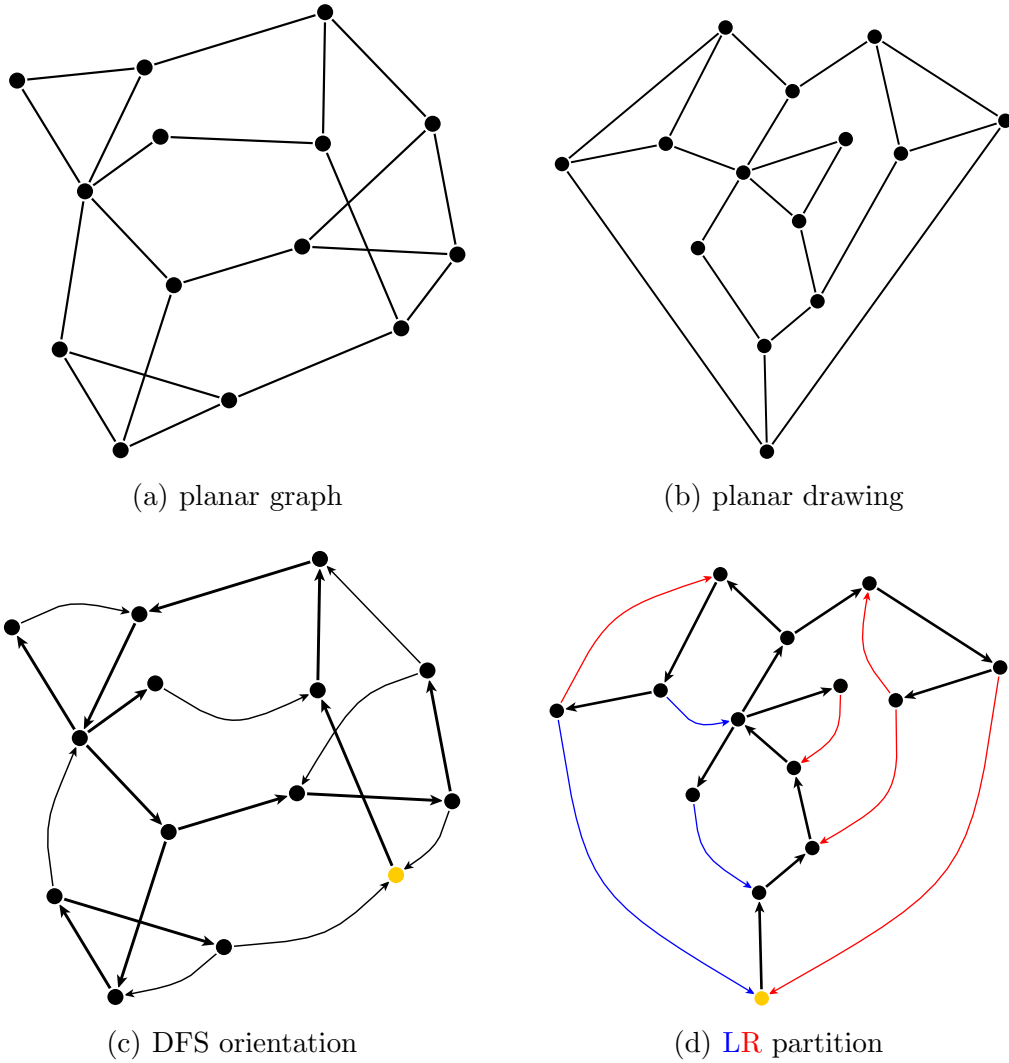


Fig. 1. Example of a planar graph (from Cai, Han, and Tarjan 1993). In both the planar and non-planar drawing, the same depth-first search (DFS) orientation is shown with thick tree edges and curved back edges. In any planar drawing the back edges can be partitioned into **left** and **right**, depending 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.

126 *3.1 Depth-first search*

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

130 Recall that a depth-first search on a connected undirected graph $\bar{G} = (V, \bar{E})$
 131 yields a *DFS orientation* of \bar{G} , i.e., a directed graph $\vec{G} = (V, \vec{E})$ in which
 132 each undirected edge is oriented according to its traversal direction. Once

133 the graph is oriented, we will only work with its directed version and hence
 134 neglect the distinction between \bar{E} and \vec{E} . In the oriented graph, we denote by
 135 $E^+(v) = \{(v, w) \in E : w \in V\}$ the set of all outgoing edges of $v \in V$, so that
 136 $E = \bigcup_{v \in V} E^+(v)$.

137 In addition to an orientation, a DFS traversal yields a bipartition $E = T \uplus B$ of
 138 the edges, where those in T are called *tree edges* and induce a rooted spanning
 139 tree (the *DFS tree*), and the non-tree edges in B are called *back edges*. See
 140 Figure 3. We write $u \rightarrow v$ and $v \hookrightarrow w$ for $(u, v) \in T$ and $(v, w) \in B$. Also,
 141 we use $\xrightarrow{+}$ for the transitive and $\xrightarrow{*}$ for the reflexive and transitive closure of
 142 \rightarrow . The unique sequence of edges inducing $u \xrightarrow{*} v$ is called a *tree path*. While
 143 the tree path is empty for $u = v$, $u \xrightarrow{+} v$ induces a *proper* tree path implying
 144 $u \neq v$. We occasionally make use of straightforward generalizations to edges
 145 such as $(u, v) \xrightarrow{*} w$ in case $v \xrightarrow{*} w$.

146 If $v \xrightarrow{*} w$ ($v \xrightarrow{+} w$), v is said to be (strictly) *lower* than w , and w (strictly)
 147 *higher* than v . A vertex is *lowest* (*highest*) in a set of vertices, if no other
 148 member of that set is lower (higher). The *height* of a vertex v is its distance
 149 from the root. With these definitions we adopt the convention that the root
 150 is indeed the lowest vertex in a tree, and illustrations are drawn accordingly.

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

157 **Lemma 3** *Let $G = (V, T \uplus B)$ be a DFS-oriented graph.*

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

161 **PROOF.**

- 162 (1) All fundamental cycles are simple and, because of DFS, directed. Now
 163 consider any simple directed cycle and let $v \in V$ be lowest on that cycle.
 164 Since every cycle contains at least one back edge, let $x \hookrightarrow u$ be the first
 165 back edge after v . Vertex v is lowest, so that u must be in $v \xrightarrow{*} x$. Since
 166 the cycle is simple, $u = v$ and there are no more edges.
 167 (2) Let $w \xrightarrow{*} v \hookrightarrow w$ and $u \xrightarrow{*} x \hookrightarrow u$ be two fundamental cycles. Since they
 168 are distinct, $v \hookrightarrow w \neq x \hookrightarrow u$. Since there is exactly one path between
 169 any pair of vertices in a tree, two tree paths can join and fork at most




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.

170 once. A non-empty intersection of $w \overset{*}{\rightarrow} v$ and $u \overset{*}{\rightarrow} x$ must, therefore, be
 171 a tree path itself. \square

172 For two overlapping cycles, the last edge $u \rightarrow v$ on the shared tree path
 173 together with the succeeding edges $e_1 = (v, w_1), e_2 = (v, w_2)$ on each cycle is
 174 called their *fork*, and v its *branching point*. We will see that finding a planar
 175 combinatorial embedding reduces to finding an appropriate ordering of all
 176 triplets of edges that form a fork. Since all forks at the same branching point
 177 share the incoming tree edge, it will be convenient to consider a *linearization*
 178 of the cyclic ordering of the outgoing edges around that vertex. It is defined
 179 by splitting the clockwise order restricted to outgoing edges at the incoming
 180 tree edge, or between any two consecutive outgoing edges if v is the root of a
 181 DFS tree.

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

184 3.2 Orientation and nesting of fundamental cycles

185 Recall that there are two classes of simple directed cycles in a planar drawing,
 186 those oriented clockwise and those oriented counterclockwise. Since the inter-
 187 section of overlapping fundamental cycles is a tree path containing at least
 188 one edge, the four possible configurations in Figure 3 can be summarized as
 189 follows, where two overlapping fundamental cycles are called *nested*, if the part




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

190 of one cycle that is not common to both is drawn completely inside the other
 191 cycle.

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

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

197 For disambiguation we use the convention that the root of a DFS tree is
 198 incident to the outer face and define the nesting depth of a fundamental cycle
 199 using the following concepts.

200 The *return points* of a tree edge $v \rightarrow w \in T$ are the ancestors u of v with
 201 $u \xrightarrow{+} v \rightarrow w \xrightarrow{*} x \hookrightarrow u$ for some descendant x of w . A back edge $v \hookrightarrow w$ has
 202 exactly one return point, its target w . The return points of a vertex $v \in V$ are
 203 formed by the union of all return points of outgoing edges $(v, w) \in E^+(v) \subseteq$
 204 $T \uplus B$. A back edge $x \hookrightarrow u$ is a *return edge* for every tree edge $v \rightarrow w$ with
 205 $u \xrightarrow{+} v \rightarrow w \xrightarrow{*} x \hookrightarrow u$, and for itself.

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

209 Our second important observation establishes nesting constraints induced by
 210 lowpoints of cycles. It is justified by noting that if the root is on the outer
 211 face and there is a proper tree path from the lowpoint of one cycle to that of
 212 another cycle, this path can not be part of the inner of the two cycles.

213 **Observation 2** *In a planar drawing of a connected DFS-oriented graph $G =$
214 $(V, T \uplus B)$ with the root of the DFS tree on the outer face, overlapping funda-
215 mental cycles are nested according to their lowpoint order.*

216 3.3 Relation to planar embeddings

217 The above two observations about orientations have immediate consequences
218 for planar embeddings which become evident by considering the single fork in
219 each of the four configurations in Figure 3.

220 Considering the fork of a pair of differently oriented cycles, we see that the
221 outgoing edge of the left cycle is before the outgoing edge of the right cycle in
222 the linearized order at branching point v .

223 For a pair of cycles oriented alike, on the other hand, the order depends on
224 their orientation. In case they are right cycles and one contains a vertex that
225 is strictly lower than those in the other cycle, the lower cycle's outgoing edge
226 (e_1 in Figure 3) comes first in the linearized order at branching point v . The
227 converse is true when v is the branching point of left cycles.

228 A vertex may be the branching point for several pairs of overlapping cycles.
229 Combining both observations yields a (partial) embedding at branching points:
230 outgoing edge of left cycles need to be before those of right cycles, and the
231 internal ordering in each subset is determined by lowpoints. Note that there
232 may be ties, and that outgoing tree edges may be part of several, differently
233 oriented cycles. We will have to resolve these ambiguities, but otherwise the
234 approach rests entirely on Observations 1 and 2.

235 4 The Left-Right Planarity Criterion

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

240 The following, crucial definition summarizes all constraints resulting from sets
241 of overlapping fundamental cycles in terms of their respective back edge. It is
242 worth noting that all constraints are generated by a single type of configuration
243 associated with forks.

244 **Definition 4 (LR partition)** *Let $G = (V, T \uplus B)$ be a DFS-oriented graph.
245 A partition $B = L \uplus R$ of its back edges into two classes, referred to as left*

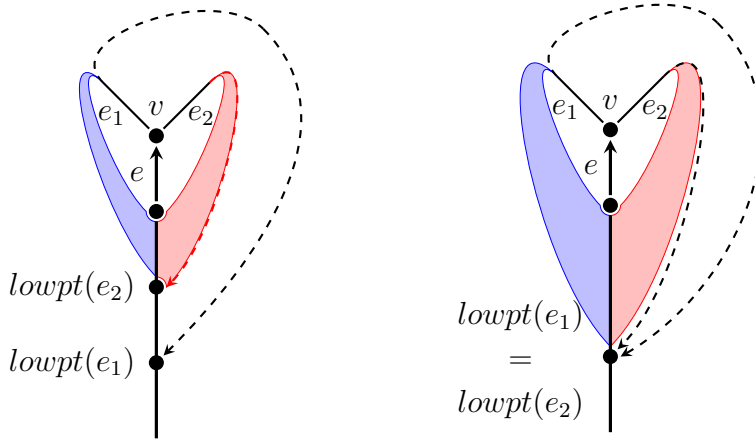


Fig. 4. LR constraints associated with $e = u \rightarrow v$.

246 and right, is called left-right partition, or LR partition for short, if for every
 247 fork consisting of $u \rightarrow v \in T$ and $e_1, e_2 \in E^+(v)$

- 248 (1) all return edges of e_1 ending strictly higher than $lowpt(e_2)$
 249 belong to one class and
 250 (2) all return edges of e_2 ending strictly higher than $lowpt(e_1)$
 251 to the other.

252 The LR partition constraints are illustrated in Figure 4. They can be broken
 253 down into two sets of pairwise constraints: *same*-constraints forcing two back
 254 edges to be on the same side, and *different*-constraints forcing them to be on
 255 opposite sides; furthermore, each constraint is *associated* with a unique tree
 256 edge ($e = u \rightarrow v$ in Figure 4).

257 Note that two back edges are subject to a constraint only if their fundamental
 258 cycles overlap, and the minimal configurations inducing a constraint are char-
 259 acterized in Section 7. It is rather striking that the above partition constraints
 260 (based on an arbitrary DFS orientation) represent a condition equivalent to
 261 planarity.

262 **Theorem 5 (Left-Right Planarity Criterion)** *A graph is planar if and*
 263 *only if it admits an LR partition.*

264 While necessity of the LR constraints is straightforward, we prove sufficiency
 265 in the next section by constructing a planar embedding from a given LR par-
 266 tition. The construction is guided by those constraints that orientation and
 267 nesting of fundamental cycles impose on an embedding.

268 Removing the following ambiguity will simplify both, argumentation and al-
 269 gorithm. An LR partition is called *aligned*, if all return edges of a tree edge e
 270 that return to $lowpt(e)$ are on the same side.

271 **Lemma 6** *Any LR partition can be turned into an aligned LR partition.*

272 **PROOF.** Consider two return edges b_1, b_2 of a tree edge $e = u \rightarrow v$ that end
273 at $\text{lowpt}(e)$. If one of them is involved in any LR constraint as specified in
274 Definition 4, this constraint is associated with a tree edge $e' = u' \rightarrow v'$ such
275 that $v' \xrightarrow{*} v$ and $\text{lowpt}(e')$ is strictly lower than $\text{lowpt}(e)$. Since b_1, b_2 originate
276 from a common subtree entered by e and have the same return point, actually
277 both are involved in this constraint and even required to be on the same side.
278 Thus, alignments do not lead to contradictions. \square

279 4.1 Combinatorial embedding

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

282 As motivated Section 3, the clockwise cyclic orderings of edges around non-root
283 vertices are linearized by starting from the unique incoming tree edge. Then,
284 outgoing edges belonging to a counterclockwise cycle need to appear before
285 those belonging to a clockwise cycle. Moreover, outgoing edges of clockwise
286 (counterclockwise) cycles must be ordered outside in (inside out) around their
287 branching point.

288 Given a DFS-oriented graph $G = (V, T \uplus B)$ together with an LR partition of
289 all back edges, we show that a planar embedding can be obtained by extending
290 the partition to cover tree edges as well and defining a linear nesting order on
291 the outgoing edges of each vertex. If the root is incident to the outer face, the
292 order determines an outside-in nesting of the cycles. The order is used without
293 modification as the embedding order for right outgoing edges, but reversed for
294 left outgoing edges by flipping them to appear before any right edges. In an
295 implementation, this can be realized by assigning each edge its rank in the
296 nesting order, changing the sign of left edges to minus, and sorting the edges
297 according to the signed ranks.

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


303 To determine the partial *nesting order* \prec , assume for a moment that all edges
304 are on the right side and consider a fork consisting of $u \rightarrow v$ and outgoing edges
305 e_1, e_2 of v . If both have return edges, v is a branching point of overlapping
306 fundamental cycles sharing $u \rightarrow v$. Since both cycles are clockwise for now,

307 we must properly nest them to avoid edge crossings. Since we fixed the root
 308 of the DFS tree to be part of the outer face, we have to define $e_1 \prec e_2$ if and
 309 only if the lowpoint of e_1 is strictly lower than that of e_2 . If both have the
 310 same lowpoint, but, say, only e_2 has another return point, we say that e_2 is
 311 *chordal* and let $e_1 \prec e_2$, because cycles containing e_2 and a return edge ending
 312 higher than $\text{lowpt}(e_2)$ can only lie inside of cycles containing e_1 and a return
 313 edge ending at $\text{lowpt}(e_1) = \text{lowpt}(e_2)$. If both e_1 and e_2 are chordal, the tie
 314 is broken arbitrarily, because eventually these two edges must be on different
 315 sides anyway.

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

319 The partial nesting order \prec is extended to a combinatorial embedding by
 320 *LR ordering*, i.e. by flip-reversing left edges before right ones and placing
 321 incoming back edges on the appropriate side of the tree edge leading to them.
 322 Some care is needed to avoid crossings of back edges, but we will see that,
 323 algorithmically, this embedding is almost trivial to obtain.

324 **Definition 7 (LR Ordering)** *Given an LR partition, let $e_1^L \prec \dots \prec e_\ell^L$ be*
 325 *the left outgoing edges of a vertex v , and $e_1^R \prec \dots \prec e_r^R$ its right outgoing edges.*
 326 *If v is not the root, let u be its parent. The clockwise left-right ordering, or*
 327 *LR ordering for short, of the edges around v is defined as follows:*



329 where (u, v) is absent if v is the root, and $L(e)$ and $R(e)$ denote the left and
 330 right incoming back edges whose cycles share e . For two back edges $b_1 = x_1 \hookrightarrow$
 331 $v, b_2 = x_2 \hookrightarrow v \in R(e)$ let $z \rightarrow x, (x, y_1), (x, y_2)$ be the fork of $C(b_1)$ and $C(b_2)$.
 332 Then, b_1 comes after b_2 in $R(e)$ if and only if $(x, y_1) \prec (x, y_2)$. If $b_1, b_2 \in L(e)$,
 333 the order is reversed.

334 **Lemma 8** *Given an LR partition, LR ordering yields a planar embedding.*

335 **PROOF.** Let $G = (V, T \uplus B)$ be a DFS-oriented graph with an LR partition
 336 $B = L \uplus R$. We assume that the partition is aligned and extend it to cover
 337 also the tree edges as described above. Now consider the embedding defined
 338 by LR ordering the edges around each vertex.

339 Since a graph with a spanning tree can always be drawn in such a way that
 340 a given embedding is respected, no two edges cross more than once, and none




Fig. 5. Two types of crossings in proof of Lemma 8.

341 of the crossings involves a tree edge, the embedding is either planar, or any
 342 such drawing yields a simple crossing of two back edges (a crossing of more
 343 than two edges can be resolved into pairwise crossings). Only two cases are
 344 possible.

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

346 Assume $x_1 \hookrightarrow u_1, x_2 \hookrightarrow u_2 \in R$ cross (the other case is symmetric). If
 347 $u_1 = u_2$, the crossings contradicts our definition of LR ordering the edges
 348 around that vertex.

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

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

360 By definition of \prec , either $lowpt(e_1)$ is strictly lower than u_2 , or $lowpt(e_1) =$
 361 $u_2 = lowpt(e_2)$ and e_2 is chordal as well. In the former case, $x_1 \hookrightarrow u_1$ and
 362 $x_2 \hookrightarrow u_2$ had to be assigned different sides. In the latter case, the highest
 363 ending return edge of e_2 is right as is e_2 , but conflicting with $x_1 \hookrightarrow u_1$ which
 364 is also right. In either case a contradiction.

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

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

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

371 Otherwise, we may assume that $\text{lowpt}(e_1)$ is strictly lower than u_2 (the
 372 case that $\text{lowpt}(e_2)$ is strictly lower than u_1 is symmetric). Then, all return
 373 edges of e_2 ending at u_2 or higher must be on the same side as $x_2 \leftrightarrow u_2 \in L$,
 374 so that e_2 is left as well. Since e_1 comes before e_2 , it must also be left and
 375 $e_2 \prec e_1$.

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

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

385 We have thus proved constructively the non-obvious implication of the left-
 386 right planarity criterion (Theorem 5).

387 5 Straightforward Algorithm

388 As an intermediate exercise, we derive a polynomial-time planarity test di-
 389 rectly from the characterization in the previous section. It mainly serves to
 390 build a better intuition for the subsequent linear-time algorithms.

391 Let $G = (V, T \uplus B)$ be a DFS-oriented graph. According to Theorem 5, testing
 392 planarity amounts to testing for the existence of an LR partition $B = L \uplus R$ of
 393 its back edges. Such a partition exists, if and only if the LR constraints of all
 394 forks can be satisfied simultaneously. This can be tested using the following
 395 immediate consequence of Definition 4, which is illustrated in Figure 6.

396 **Corollary 9** *Let $G = (V, T \uplus B)$ be a DFS-oriented graph. For a pair of back*
 397 *edges $b_1, b_2 \in B$ with overlapping fundamental cycles, let $v_1 \rightarrow \dots \rightarrow v_k$ be the*
 398 *tree path of their intersection and $(v_{k-1}, v_k), e_1, e_2$ the corresponding fork with*
 399 *$e_1 \xrightarrow{*} b_1$ and $e_2 \xrightarrow{*} b_2$. Then, b_1 and b_2 are subject to*

- 400 • a different-constraint, if and only if $\text{lowpt}(e_2) < \text{lowpt}(b_1)$ and $\text{lowpt}(e_1) <$
 401 $\text{lowpt}(b_2)$.
- 402 • a same-constraint, if and only if $\text{lowpt}(e') < \min\{\text{lowpt}(b_1), \text{lowpt}(b_2)\}$ for
 403 some $e' = (v_i, w) \in T \uplus B$, $1 < i < k$, $w \neq v_{i+1}$.




Fig. 6. The constraints between pairs of back edges b_1, b_2 summarized in the definition of LR constraints are induced by three types of minimal configurations (de Fraysseix and Rosenstiehl 1985; cf. Corollary 9).




Fig. 7. The constraint graph for the example from Figure 3 consists of eight (square) vertices, one *same*-constraint, and three *different*-constraints. Note that the LR partition is not unique because, e.g., the lower isolates are not aligned.

404 With this observation and precomputed lowpoints, we can test whether two
 405 given back edges are subject to a constraint by traversing their fundamental
 406 cycles, determining the branching point in case they overlap, and comparing
 407 a few lowpoints (possibly including those of edges incident to vertices in the
 408 intersection). Note that a *different*-constraint can be associated with only one
 409 fork, whereas a *same*-constraint may be induced repeatedly.

410 Pairwise constraints can be represented in a graph that has back edges as its
 411 vertices and an edge between two of them, if they are subject to a constraint.
 412 To distinguish the type of constraint, we use *signed edges* that carry labels
 413 “+1” or “-1” as indicated in Figure 7.

Definition 10 Let $G = (V, T \uplus B)$ be a DFS-oriented graph such that each pair of back edges $b_1, b_2 \in B$ is subject to at most one type of constraint. The signed graph $\mathcal{C}(G) = (B, E(\mathcal{C}); \sigma : E(\mathcal{C}) \rightarrow \{-1, +1\})$ with

$$\sigma(b_1, b_2) = \begin{cases} -1 & \text{if } b_1, b_2 \in B \text{ are subject to a different-constraint} \\ +1 & \text{if } b_1, b_2 \in B \text{ are subject to a same-constraint} \end{cases}$$

414 is called constraint graph of G .

415 If any pair of back edges is subject to both a *same*-constraint and a *different*-
416 constraint, no LR partition can exist and hence the graph is non-planar. This
417 is noticed during the construction of $\mathcal{C}(G)$, and we may therefore assume in
418 the following that each pair of back edges is subject to at most one type of
419 constraint.

420 Finding an LR partition that satisfies all LR constraints is then equivalent
421 to testing whether the constraint graph is *balanced* (Harary and Cartwright,
422 1956), i.e. whether there is a bipartition such that each edge labeled “+1”
423 connects two vertices in the same set, and each edge labeled “−1” connects
424 vertices in different sets. Balancedness of signed graphs is equivalent to the
425 absence of cycles with an odd number of edges labeled “−1,” and can hence
426 be tested in linear time using a variant of breadth-first search (Harary and
427 Kabell, 1980). The reader is encouraged to fill in the details.

428 6 Linear-Time Algorithm

429 The straightforward partition approach of the previous section can be refined
430 into a linear-time algorithm for planarity testing and embedding. Extraction
431 of a minimal non-planar subgraph is treated only in the next section. After
432 a high-level description of its three main phases shown in Algorithm 1, full
433 implementation details are provided for all operations but those concerning
434 the specific data structure representing the graph and its embedding.

435 **Orientation.** The algorithm is based on the left-right planarity criterion
436 and therefore starts with a depth-first search to orient the input graph. For
437 each connected component, the root of its spanning DFS tree is appended to a
438 list, *Roots*. The tree-path distance of a vertex from the root of its component
439 is stored in an array *height*, so that roots of unexplored components are identified
440 by still having the initial value ∞ . Different from most other planarity
441 algorithms, there is no need to worry about biconnected components.

442 During DFS, lowpoints of edges are computed and the partial nesting order \prec

variable	type	interpretation	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	intervals with conflicting edges, i.e., a conflict pair	n.a.
S	stack of conflict pairs	conflict pairs consisting of 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 (i.e. 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. 8. 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** $|V| > 2$ **and** $|E| > 3|V| - 6$ **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  $side[e]$ 
```

443 is determined by assigning to each edge an integer value $nesting_depth$ such
444 that $e_1 \prec e_2$ implies $nesting_depth[e_1] < nesting_depth[e_2]$.

445 **Testing.** To determine whether there exists an aligned LR partition, the
446 DFS trees are traversed for a second time. The traversal order is modified,
447 however, by visiting outgoing edges in the order given by $nesting_depth$. This
448 second traversal halts if the graph is not planar, and we discuss in Section 7
449 how to extract a Kuratowski subgraph in that case.

450 The tentative side of edges may change often during the test, so that the
451 bipartition is maintained only implicitly for efficiency reasons. An edge array
452 ref specifies for each edge a reference edge relative to which its side is defined,
453 and in an edge array $side$ a value of $+1$ or -1 indicates whether the side of
454 the edge is the same as, or different from, the side of its reference edge. If the
455 reference edge of e is undefined, i.e. $ref[e] = \perp$, the value of $side[e]$ specifies
456 the side directly, where -1 is for left and $+1$ is for right.

457 **Embedding.** Given an LR partition, flip-reversal of left edges is performed
458 by sorting the outgoing edges in all adjacency lists once again according to
459 their nesting order, though now modified by the signs in *side*. Since the mul-
460 tiplication of *nesting_depth* with *side* only changes the sign of left edges to
461 negative, they are effectively placed before all right edges, and in reverse or-
462 der. To complete the LR ordering, incoming edges are rearranged during a
463 third traversal of the DFS forest that is guided once again by the new order
464 of outgoing edges.

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

467 6.1 Orientation

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
                     $\text{lowpt2}[e] \leftarrow \min\{\text{lowpt2}[e], \text{lowpt2}[(v, w)]\}$ 

```

468 The purpose of the first DFS is to orient the graph, and to determine lowpoints
469 and nesting order \prec . It is therefore a standard DFS computing the auxiliary
470 variables listed in Table 8(a). Except for *height*, all of them are determined
471 during backtracking.

472 Our use of lowpoints is slightly unusual in two ways. Firstly, we determine
473 lowpoints for edges rather than vertices, and, secondly, we do not assign DFS
474 numbers, but heights. The latter induce the same ordering of ancestors as
475 do DFS numbers, but are related to the tree more intuitively, and in general
476 result in a smaller range of values which may in turn speed up the subsequent
477 sorting of adjacency lists according to *nesting_depth*.

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

481 The rationale for representing \prec via *nesting_depth* is two-fold: firstly, we can
482 sort the edges in linear time using, e.g., bucket sort, because the range of
483 values is linear in the size of the graph, and secondly, flip-reversal of left edges
484 after the second phase can be performed by changing their sign and sorting
485 again.

486 6.2 Testing

487 The second phase is the working horse of the algorithm. It determines an
488 aligned LR partition including all tree edges, if one exists. With this, LR or-
489 dering can be carried out as described in Section 6.3, otherwise the code can
490 be augmented to identify fundamental cycles whose union yields a Kuratowski
491 subgraph as described in Section 7.

492 Our strategy will be to implement the straightforward algorithm of Section 5
493 without constructing the constraint graph explicitly. Recall that constraints
494 are associated with a tree edge and that there are only two types of constraints:
495 according to Definition 4, a pair of back edges with overlapping fundamental
496 cycles can be required to be placed either on the same side or on different
497 sides.

498 Clearly, we cannot afford to detect the edges of the signed constraint graph
499 individually, because their number may already be quadratic in the size of the
500 original graph. Since our actual goal is a bipartition certifying that the con-
501 straint graph is balanced, we will eagerly maintain bipartitions of its connected
502 components and represent constraints only implicitly to test for contradictions.

503 To represent a bipartition it is sufficient to have a signed spanning forest of




Fig. 9. The main data structure is a stack S storing conflicting pairs of intervals with consecutively returning back edges.

504 the constraint graph available. We will therefore construct a rooted tree for
 505 each component using a reference pointer ref for every edge. Such a pointer is
 506 stored not only for back edges, but also for tree edges in the DFS orientation,
 507 since in the extended LR partition required for LR ordering, their sides are
 508 determined by reference to a return edge ending at the highest return point,
 509 anyway.

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

515 To grow the partial bipartition, we need to keep track of all constraints en-
 516 countered during an ordered examination of all forks, but instead of storing
 517 constraints individually, a compact data structure is used to represent their
 518 transitive closure. Observe that the *same*-constraints induced by a fork $u \rightarrow v$,
 519 $e_1, e_2 \in E^+(v)$ in Definition 4 involve two sets of return edges with a simple
 520 structure. For, say, e_1 let $h = x_h \leftrightarrow u_h$ and $\ell = x_\ell \leftrightarrow u_\ell$ be the two (pos-
 521 sibly equal) return edges ending at the highest and lowest return point of e_1
 522 that is also a return point of $u \rightarrow v$ (i.e., $u_h \neq u$) and strictly higher than
 523 $lowpt(e_2)$. Then we know that h_ℓ and all return edges $x' \leftrightarrow u'$ of e_1 with a re-
 524 turn point in $u_\ell \xrightarrow{*} u_h$ are in the same group of *same*-constraints. This *interval*
 525 of edges can thus be represented by its two bounding members, h and ℓ , as
 526 shown in Figure 9(a). Return edges belonging to an interval are maintained in
 527 a singly-linked list, from highest to lowest return point, using the ref -array.

528 The closure of *different*-constraints can be summarized similarly, because by

529 transitivity it always involves all pairs of edges in a pair of intervals. A *conflict*
530 *pair* therefore consists of two intervals of edges subject to at least one *different-*
531 *constraint* as shown in Figure 9(b). It represents their tentative assignment to
532 the left and right, and thus a partial bipartition.

533 The second DFS traversal is designed to build an extended LR partition of
534 edges incrementally by merging conflict pairs. Its main data structure is a stack
535 S of conflict pairs representing all constraints associated with a tree edge that
536 has been traversed, but not yet backtracked over. Note that these constraints
537 involve only back edges that have already been traversed, but return to a
538 vertex below the current one. In other words, each back edge in the stack is a
539 return edge for at least one tree edge in the current DFS path.

540 By processing the DFS trees bottom-up, the constraints associated with an
541 edge can be determined by merging those associated with its outgoing edges.
542 Two main invariants are maintained. Clearly, we can not prove them before
543 the algorithm is described, but since they provide an orientation for under-
544 standing the implementation, they are stated already here and the reader is
545 encouraged to check that they are maintained. The first invariant eventually
546 yields correctness of the implementation,

547 **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.

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

550 **Ordering Invariant:** For any two conflict pairs P, Q where P is
above Q in the stack, no edge in P has a return point below that
of an edge in Q . Each interval in a conflict pair is represented as a
singly-linked list of return edges that is ordered from highest to lowest
return point as well.

551 6.2.1 Ordered traversal

552 Pseudo-code for the second DFS is given in Algorithms 3–5. All edges have
553 been oriented during the first DFS, and they are traversed again in the same
554 direction. The traversal order differs, though, since adjacency lists have been
555 rearranged according to *nesting_depth*, so that outgoing edges with lower low-
556 points are traversed first. This reordering is crucial for the ordering invariant.

557 When visiting a vertex v during the DFS traversal, the high-level task is to

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 */
    |  $\text{lowpt\_edge}[e_i] \leftarrow e_i$ ;  $\text{push}(\emptyset, [e_i, e_i]) \rightarrow S$ 
    ▼ integrate new return edges
      | 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
          | | ► add constraints of  $e_i$  (Algorithm 4)
    ▼ remove back edges returning to parent
      | if  $e \neq \perp$  then /*  $v$  is not root */
        |  $u \leftarrow \text{source}(e)$ 
        | ► trim back edges ending at parent  $u$  (Algorithm 5)
        | ▼ 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 ( $h_R = \perp$  or  $\text{lowpt}[h_L] > \text{lowpt}[h_R]$ ) then
              | | |  $\text{ref}[e] \leftarrow h_L$ 
            | | else
              | | |  $\text{ref}[e] \leftarrow h_R$ 
```

558 recursively determine the constraints for all outgoing edges and integrate them
559 into those associated with parent edge $e = u \rightarrow v$ (if v is not a DFS root).

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

571 From the partitioning invariant we know that when returning from the traversal of e_i , the conflict pairs above $stack_bottom[e_i]$ represent a partial LR partition of all return edges of e_i . While processing the first outgoing edge e_1 we simply leave them on the stack, if any, and pass on $lowpt_edge[e_1]$ to $lowpt_edge[e]$. Note that, since e_1 has a return edge, $parent_edge[v] = e \neq \perp$, i.e. v is not a root. For each of the other outgoing edges $e_i = e_2, \dots, e_d \in E^+(v)$, the constraints above $stack_bottom[e_i]$ are merged into those which have already been accumulated for e and are directly beneath in S . Constraint integration is the most essential step and described separately in Algorithm 4 and below.

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

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

593 6.2.2 Adding constraints associated with the next outgoing edge

594 We have to merge all constraints associated with the next outgoing edge, e_i , with those already accumulated from e_1, \dots, e_{i-1} . The involved intervals are therefore gathered one by one in an initially empty conflict pair P as illustrated in Figure 10.

598 **Merge return edges of e_i into right interval.** All return edges of e_i have been traversed since traversing e_i , and they are represented in the top conflict pairs on stack S down to, but not including, $stack_bottom[e_i]$. All of these intervals have to be merged on one side because of the *same*-constraints induced by the fundamental cycle of $lowpt_edge[e]$ according to Definition 4. If there is a conflict pair with two non-empty intervals, merging on one side violates an earlier constraint and the graph is not planar.

605 There is at least one conflict pair above $stack_bottom[e_i]$ for otherwise we would not have entered this section. The non-empty interval of each such pair is merged in the right interval PR of P without changing their order by

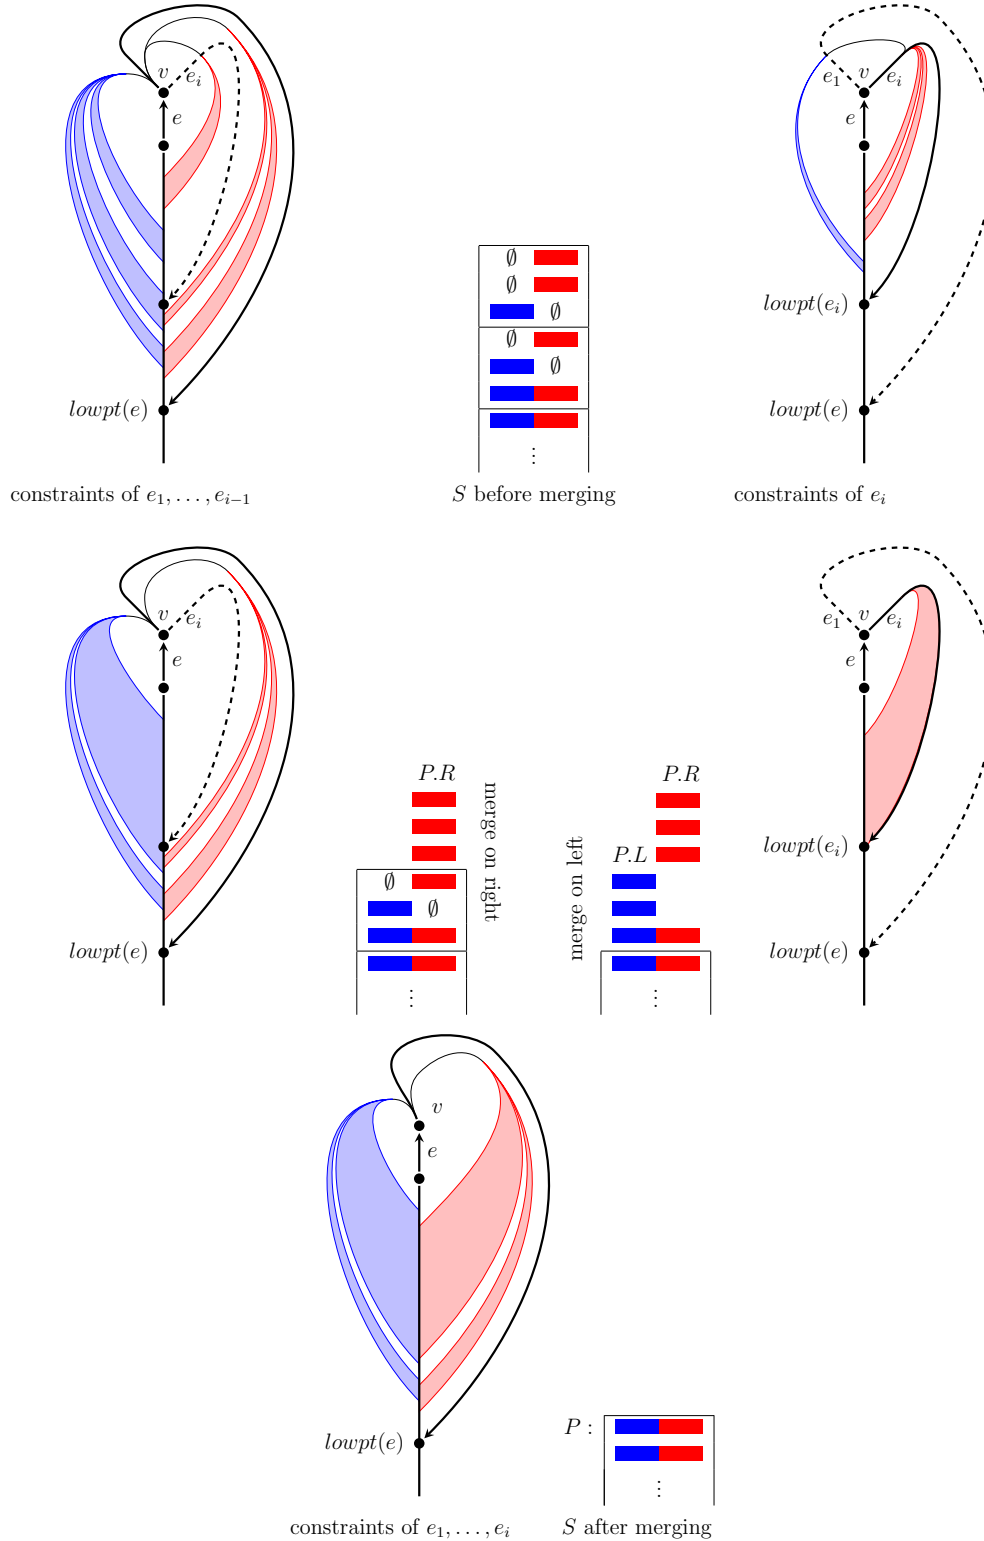


Fig. 10. In the core step of the algorithm, the constraints of e_i are merged into those of e_1, \dots, e_{i-1} . Horizontal lines indicate where the top of stack S is divided by $stack_bottom[e_i]$ and the topmost pair not in conflict with $lowpt_edge[e_i]$. If $lowpt(e_i) = lowpt(e)$, the pair containing only $lowpt_edge[e_i]$ is not merged into $P.R$, but the bipartition is aligned by assigning $ref[lowpt(e_i)] \leftarrow lowpt(e)$.

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 /* topmost interval */
            |  $P.R.\text{high} \leftarrow Q.R.\text{high}$ 
          else
            |  $\text{ref}[P.R.\text{low}] \leftarrow Q.R.\text{high}$ 
            |  $P.R.\text{low} \leftarrow Q.R.\text{low}$ 
          else /* align */
            |  $\text{ref}[Q.R.\text{low}] \leftarrow \text{lowpt\_edge}[e]$ 
        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 /* topmost interval */
        |  $P.L.\text{high} \leftarrow Q.L.\text{high}$ 
      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 $\text{lowpt}[I.\text{high}] > \text{lowpt}[b]$)

608 having the lowest edge of $P.R$ refer to the highest edge of the next conflict pair
609 and replacing it accordingly. An exception is the interval containing a return
610 edge to the lowpoint of e ; to align the LR partition, we make it refer to the
611 lowpt_edge directly.

612 **Merge conflicting return edges of e_1, \dots, e_{i-1} into left interval.** Re-
613 turn edges of e_1, \dots, e_{i-1} with lowpoints higher than $lowpt[e_i]$ are subject
614 to pairwise *same*-constraints and to a *different*-constraint with respect to
615 some return edge of e_i . (If $lowpt[e_i] = lowpt[e]$ this is not $lowpt_edge[e_i]$ but,
616 e.g., a back edge returning to $lowpt2[e_i]$ which must exist, because apparently
617 $lowpt2[e_{i-1}]$ exists as well by the way outgoing edges are ordered).

618 So while there are conflict pairs on the stack that contain return edges with
619 lowpoints higher than $lowpt[e_i]$, these have to be merged on one side. If such
620 a pair contains two intervals ending above $lowpt[e_i]$, we again have a contra-
621 diction with a previous constraint and thus non-planarity. If only one side
622 ends above $lowpt[e_i]$, we merge the other into $P.R$ (effectively closing these
623 constraints under transitivity).

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

626 6.2.3 *Trimming back edges*

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

632 **Dropping entire conflict pairs.** If the lowest lowpoint on either side of a
633 conflict pair P is the source of the current tree edge $u \rightarrow v$, all lowpoints of
634 back edges in P are the same and the edges will not be involved in any future
635 constraints. The pair is finalized by assigning the lowest back edge of the left
636 interval to the left side. Since *side* is initialized with 1, the lowest back edge
637 in the right interval $P.R$ is already assigned correctly to the right side, and all
638 other back edges b in P to the same side as $ref[b]$.

639 **Trimming a left interval.** Since back edges in an interval are concatenated
640 by *ref*-pointers in an order monotonic in the *height* of their lowpoints, we
641 can simply remove back edges from the upper end of the left interval until
642 the highest lowpoint is no longer u , or the interval has become empty. In the
643 latter case the lower end of the interval is still defined and made to refer to
644 an edge on the other side, setting its *side* to -1 accordingly. Note that the
645 right interval cannot be empty for otherwise the entire conflict pair had been
646 removed in the first while loop. All other removed back edges still refer to a

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  $\text{lowest}(\text{top}(S)) = \text{height}[u]$  do
  |   |    $P \leftarrow \text{pop}(S)$ 
  |   |   if  $P.L.\text{low} \neq \perp$  then  $\text{side}[P.L.\text{low}] \leftarrow -1$ 
  |   if  $S \neq \emptyset$  then /* one more conflict pair to consider */
  |   |    $P \leftarrow \text{pop}(S)$ 
  |   |   ▼ trim left interval
  |   |   |   while  $P.L.\text{high} \neq \perp$  and  $\text{target}(P.L.\text{high}) = u$  do
  |   |   |   |    $P.L.\text{high} \leftarrow \text{ref}[P.L.\text{high}]$ 
  |   |   |   |   if  $P.L.\text{high} = \perp$  and  $P.L.\text{low} \neq \perp$  then /* just emptied */
  |   |   |   |   |    $\text{ref}[P.L.\text{low}] \leftarrow P.R.\text{low}; \text{side}[P.L.\text{low}] \leftarrow -1$ 
  |   |   |   |   |    $P.L.\text{low} \leftarrow \perp$ 
  |   |   |   ► trim right interval
  |   |   |   push  $P \rightarrow S$ 
```

where

```
integer  $\text{lowest}(\text{conflictpair } P)$ 
  |   if  $P.L = \emptyset$  then return  $\text{lowpt}[P.R.\text{low}]$ 
  |   if  $P.R = \emptyset$  then return  $\text{lowpt}[P.L.\text{low}]$ 
  |   return  $\min\{\text{lowpt}[P.L.\text{low}], \text{lowpt}[P.R.\text{low}]\}$ 
```

647 back edge on the same side, so that the initial 1 of their *side*-entry must not
648 be changed.

649 **Trimming a right interval.** This is symmetric to the previous operation.
650 Note, however, that the assigned *side* in case the right interval becomes empty
651 is -1 as well, because this indicates that the side of the lowest back edge is
652 different from the side of the lowest back edge in the left interval. Again, the
653 left interval cannot be empty.

654 **Assigning a side to a tree edge.** After trimming all back edges ending
655 at source u of the current tree edge $e = u \rightarrow v$ in Algorithm 3, the side of
656 e is determined by reference to a highest return edge. There is a return edge
657 only if $\text{lowpt}[e] < \text{height}[u]$. Otherwise, u is a cutvertex or root and it does
658 not matter which side e is assigned to. Since the existence of a return edge
659 renders S non-empty, the ordering invariant asserts that the highest return
660 point is found by comparing lowpoints of the highest return edges in the two
661 intervals of the top constraint pair on S (checking for existence).




Fig. 11. The algorithm testing $K_{3,3}$ and K_5 for planarity. For both cases, 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 .

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

670 Two small examples are shown in Figure 11. Even though both graphs are
 671 non-planar, the workings of the algorithm are nicely illustrated, since coloring
 672 and embedding correspond to the current (implicitly represented) bipartition
 673 and LR ordering.

674 6.3 Embedding

675 Compared to other planarity algorithms, the embedding phase is extremely
 676 simple. LR ordering the outgoing edges of the DFS-oriented graph is achieved
 677 by sorting them according to their *nesting_depth* on both sides. Such an em-
 678 bedding of outgoing edges is already sufficient for a planar combinatorial em-
 679 bedding (see, e.g., Cai 1993), but for completeness we provide full details in
 680 Algorithm 6.

681 The DFS forest is traversed for the third time. Since, after sorting, outgoing
 682 edges are already ordered in the desired way, back edges are encountered
 683 exactly as required in the definition of LR ordering. As described in Table 8(c)
 684 we therefore maintain, for each vertex v , the two positions next to which the

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 after  $\text{rightRef}[w]$  in adjacency list of  $w$ 
      else
        | place  $e_i$  directly before  $\text{leftRef}[w]$  in adjacency list of  $w$ 
        |  $\text{leftRef}[w] \leftarrow e_i$ 
```

685 next left or right incoming back edge is to be inserted.

686 Observe that incoming back edges from the same subtree actually appear in
687 in counterclockwise order. If the data structure available for embedded graphs
688 does not provide a constant-time method for direct neighbor insertion, we
689 may therefore use the now obsolete array *ref* to build a singly-linked list of
690 all edges incident to a vertex in counterclockwise order instead.

691 *6.4 Running time and implementation*

692 **Theorem 11** *Algorithm 1 can be implemented to test in $\mathcal{O}(n)$ time whether*
693 *a graph is planar and return a planar combinatorial embedding if it is.*

694 **PROOF.** We have argued throughout this section that the algorithm cor-
695 rectly yields an LR ordering if the graph admits an LR partition. Hence, cor-
696 rectness is established by the left-right planarity criterion (Theorem 5). Recall
697 that the initial test is justified by Corollary 2, and we may hence assume that
698 $m \in \mathcal{O}(n)$.

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

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

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

712 Dereferencing *ref*-pointers takes linear time, because it is performed only once
713 before the third DFS traversal, which also requires linear time if the graph
714 data structure provides a constant-time operation to move an edge next to
715 another in the embedding order. If it does not, the algorithm can be altered
716 to re-use *ref*-pointers for the embedding as described in Section 6.3. \square

717 The left-right approach can be implemented as described above and our expe-
718 riences with its performance essentially confirm the favorable results of Boyer,
719 Cortese, Patrignani, and Di Battista (2004). A special edge numbering scheme
720 used in the PIGALE implementation (de Fraysseix and Ossona de Mendez,
721 2002) serves to avoid repeated DFS traversals, but it seems that most of the
722 running time in our implementations is actually spend on the sorting of adja-
723 cency lists.

724 Note, however, that both sorting and DFS traversal can be avoided during
725 the testing phase by splitting the stack into singly-linked lists associated with
726 edges and processing edges (i.e., merging their final list of constraints into
727 that of another edge) in the order given by *nesting_depth*. This order is es-
728 tablished by creating two buckets for each *height* and adding an edge to its
729 respective bucket when its *lowpt* is known during the initial DFS, i.e. when it
730 is backtracked over. Since *lowpt* is determined bottom-up, edges added to the
731 same bucket end up being in the desired order.

732 **7 Non-Planarity**

733 **[This section is under revision]**

734 **8 Discussion**

735 We have reviewed the left-right planarity criterion (Theorem 5) and described
736 a simple linear-time algorithm (Algorithm 1) based on it. While this is not a

737 review of graph planarity, and many important references and developments
738 are left out, some notes on closely related work seem in place.

739 8.1 Characterization

740 In Section 7 we made use of a characterization of planar graphs in terms
741 of forbidden subgraphs (Kuratowski, 1930). This characterization can be re-
742 interpreted as identifying the overlapping cycle structures of $K_{3,3}$ and K_5 as
743 the two minimal configurations that can not be drawn planarly.

744 Therefore, among the various later characterizations, the criterion due to
745 Mac Lane (1937) appears to be related most closely, because it is also for-
746 mulated in terms of a representative set of cycles. Consider the set of all
747 undirected cycles of a graph, and define the sum of two cycles as the sym-
748 metric difference of their edge sets. These together form a vector space, called
749 *cycle space*. A *basis* of the cycle space is a minimum-cardinality set of cycles
750 such that every cycle is the sum of some basis cycles.

751 **Theorem 12 (Mac Lane’s Planarity Criterion)** *A graph is planar, if and*
752 *only if it has a cycle basis in which every edge appears at most twice.*

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

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

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

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

771 8.2 *Development*

772 The earliest precursor of the left-right approach is a planarity characterization
773 of Wu (1955), which states that a graph is planar, if and only if a certain
774 system of linear equations has a solution. It was complemented by the concept
775 of crossing chains in Tutte (1970), and refined to Boolean variables and fewer
776 equations in the 1970s (see Wu 1985, 1986; Liu 1990). The variables in this
777 smaller system are associated with the edges, and the equations represent
778 constraints generated from configurations of overlapping cycles obtained from
779 a spanning DFS forest. An alternative interpretation of the existence of a
780 solution is that of balancing a constraint graph as in Section 6.2. Rosenstiehl
781 (1980) gives an algebraic proof for this characterization.

782 This work was further developed in several papers, but the descriptions are
783 rather incomplete, in particular with respect to linear-time implementation
784 (de Fraysseix and Rosenstiehl, 1982, 1985; Xu, 1989; Cai, Han, and Tarjan,
785 1993).

786 Finally, de Fraysseix, Ossona de Mendez, and Rosenstiehl (2006) simplified
787 the approach even further by concentrating on the single constraint-inducing
788 configuration of Definition 4. While this paper is still incomplete and difficult
789 to read, the linear-time implementation is described in just enough detail to
790 provide a basis for replication. Among the differences to the present description
791 is the maintenance and merging of constraints, since intervals are described as
792 stacks rather than their extreme pairs of edges and there is a constraint stack
793 for each edge rather than our global stack S . It turns out, however, that the
794 most recent implementation in PIGALE (de Fraysseix and Ossona de Mendez,
795 2002) uses a similar representation.

796 The characterization of Kuratowski subgraphs in terms of configurations in-
797 duced by a DFS spanning tree given in de Fraysseix and Ossona de Mendez
798 (2003) and de Fraysseix (2008) led to a linear-time extraction algorithm as-
799 sociated with the left-right approach. Because of principal commonalities it
800 is likely that similarities can be unveiled, but the extraction algorithm given
801 here is original and more intimately related to the left-right characterization.

802 8.3 *Algorithms*

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

808 copies of vertices on the cycle for every incident non-cycle edge. These ver-
809 tices are called *attachments* and the connected components of the resulting
810 graph are called *segments*. Clearly, each segment must be drawn completely
811 inside or completely outside of the removed cycle, but a pair of segments must
812 not be placed in the same region if their attachments interleave on the cycle.
813 Planarity can thus be tested by recursively choosing cycles and sides.

814 The related algorithm of Demoucron, Malgrange, and Pertuiset (1964) also
815 starts from a simple cycle, but then iteratively chooses a path that can be
816 added into one of the current faces. The algorithm is not only simple, but also
817 has the unusual property to eagerly maintain a partial embedding that is not
818 changed later on. Both algorithms require $\Omega(n^2)$ time, though.

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

825 Recall how we observed in Section 3 that likewise-oriented cycles are nested
826 if they overlap. Maybe because de Fraysseix and Rosenstiehl (1982) phrased
827 the notions of left and right in terms of angles with the DFS tree rather
828 than orientations of fundamental cycles, it has gone almost unnoticed that
829 the left-right approach is yet another refinement of Auslander and Parter
830 (1961) and Hopcroft and Tarjan (1974), progressing from segments to paths
831 to edges. Together with Canfield and Williamson (1990) and Haeupler and
832 Tarjan (2008) this observation instills hope that there may be a useful and
833 elegant unification of path- and vertex-addition approaches including the two
834 most efficient versions of de Fraysseix, Ossona de Mendez, and Rosenstiehl
835 (2006) and Boyer and Myrvold (2004).

836 References

- 837 Aigner, M., Ziegler, G. M., 2009. Proofs from THE BOOK, 4th Edition.
838 Springer.
- 839 Auslander, L., Parter, S. V., 1961. On imbedding graphs in the sphere. Journal
840 of Mathematics and Mechanics 10 (3), 517–523.
- 841 Bollobás, B., 1998. Modern Graph Theory, 2nd Edition. No. 184 in Graduate
842 Texts in Mathematics. Springer.
- 843 Booth, K. S., Lueker, G. S., 1976. Testing for the consecutive ones property,
844 interval graphs, and graph planarity using PQ-tree algorithms. Journal of
845 Computer and System Sciences 13, 335–379.
- 846 Boyer, J. M., Cortese, P.-F., Patrignani, M., Di Battista, G., 2004. Stop mind-

- 847 ing your P's and Q's: Implementing a fast and simple DFS-based planarity
848 testing and embedding algorithm. In: Liotta, G. (Ed.), Proc. Intl. Symp.
849 Graph Drawing (GD '03). Vol. 2912 of LNCS. Springer-Verlag, pp. 25–36.
- 850 Boyer, J. M., Myrvold, W. J., 2004. On the cutting edge: Simplified $\mathcal{O}(n)$
851 planarity by edge additon. Journal of Graph Algorithms and Applications
852 8 (3), 241–273.
- 853 Cai, J., 1993. Counting embeddings of planar graphs using DFS trees. SIAM
854 Journal on Discrete Mathematics 6 (3), 335–352.
- 855 Cai, J., Han, X., Tarjan, R. E., 1993. An $\mathcal{O}(m \log n)$ -time algorithm for the
856 maximal planar subgraph problem. SIAM Journal on Computing 22 (6),
857 1142–1162.
- 858 Canfield, E. R., Williamson, S. G., 1990. The two basic linear time planarity
859 algorithms: Are they the same? Linear and Multilinear Algebra 26, 243–265.
- 860 de Fraysseix, H., 2008. Trémaux trees and planarity. Electronic Notes in Dis-
861 crete Mathematics 31, 169–180.
- 862 de Fraysseix, H., Ossona de Mendez, P., 2002. Pigale: Public implemen-
863 tation of a graph algorithm library and editor, software project at
864 pigale.sourceforge.net (GPL License).
- 865 de Fraysseix, H., Ossona de Mendez, P., 2003. On cotree-critical and DFS
866 cotree-critical graphs. Journal of Graph Algorithms and Applications 7 (4),
867 411–427.
- 868 de Fraysseix, H., Ossona de Mendez, P., Rosenstiehl, P., 2006. Trémaux trees
869 and planarity. International Journal of Foundations of Computer Science
870 17 (5), 1017–1029.
- 871 de Fraysseix, H., Rosenstiehl, P., 1982. A depth-first characterization of pla-
872 narity. Annals of Discrete Mathematics 13, 75–80.
- 873 de Fraysseix, H., Rosenstiehl, P., 1985. A characterization of planar graphs by
874 Trémaux orders. Combinatorica 5 (2), 127–135.
- 875 Demoucron, G., Malgrange, Y., Pertuiset, R., 1964. Graphes planaires: Recon-
876 naissance et construction de représentations planaires topologiques. Revue
877 Français Recherche Opérationnelle 8 (30), 33–47.
- 878 Diestel, R., 2005. Graph Theory, 3rd Edition. No. 173 in Graduate Texts in
879 Mathematics. Springer.
- 880 Haeupler, B., Tarjan, R. E., 2008. Planarity algorithms via PQ -trees. Elec-
881 tronic Notes in Discrete Mathematics 31, 143–149.
- 882 Harary, F., Cartwright, D., 1956. Structural balance: A generalization of Hei-
883 der's theory. Psychological Review 63 (5), 277–293.
- 884 Harary, F., Kabell, J. A., 1980. A simple algorithm to detect balance in signed
885 graphs. Mathematical Social Sciences 1, 131–136.
- 886 Hopcroft, J. E., Tarjan, R. E., 1974. Efficient planarity testing. Journal of the
887 ACM 21 (4), 549–568.
- 888 Kuratowski, K., 1930. Sur le problème des courbes gauches en Topologie. Fun-
889 damenta Mathematicae 15, 271–283.
- 890 Lempel, A., Even, S., Cederbaum, I., 1967. An algorithm for planarity testing
891 of graphs. In: Rosenstiehl, P. (Ed.), Proc. Intl. Symp. Theory of Graphs

- 892 (Rome, July 1966). Gordon and Breach, pp. 215–232.
- 893 Liu, Y., 1990. A Boolean characterization of planarity and planar embeddings
894 of graphs. *Annals of Operations Research* 24, 165–174.
- 895 Mac Lane, S., 1937. A combinatorial condition for planar graphs. *Fundamenta*
896 *Mathematicae* 28, 22–32.
- 897 Mehlhorn, K., Mutzel, P., 1996. On the embedding phase of the Hopcroft and
898 Tarjan planarity testing algorithm. *Algorithmica* 16 (2), 233–242.
- 899 Nishizeki, T., Rahman, M. S., 2004. *Planar Graph Drawing*. Vol. 12 of *Lecture*
900 *Notes Series on Computing*. World Scientific.
- 901 Rosenstiehl, P., 1980. Preuve algèbrique du critère de planarité de Wu-Liu.
902 *Annals of Discrete Mathematics* 9, 67–78.
- 903 Tutte, W. T., 1970. Toward a theory of crossing numbers. *Journal of Combi-*
904 *natorial Theory* 8, 45–53.
- 905 Wu, W., 1955. On the realization of complexes in Euclidean space I. *Acta*
906 *Mathematica Sinica* 5, 505–552, English version in *American Mathematical*
907 *Society Translations, Series 2* 78:137–184, 1968.
- 908 Wu, W., 1985. On the planar imbedding of linear graphs. *Journal of System*
909 *Science and Mathematical Sciences* 5 (4), 290–302.
- 910 Wu, W., 1986. On the planar imbedding of linear graphs (continued). *Journal*
911 *of System Science and Mathematical Sciences* 6 (1), 23–35.
- 912 Xu, W., 1989. An improved algorithm for planarity testing based on Wu-Liu’s
913 criterion. *Annals of the New York Academy of Sciences* 576, 641–652.