



An Improved Upper Bound on the Queue Number of Planar Graphs

Michael Bekos¹ · Martin Gronemann² · Chrysanthi N. Raftopoulou³

Received: 8 December 2021 / Accepted: 8 September 2022 / Published online: 28 September 2022
© The Author(s) 2022

Abstract

A k -queue layout is a special type of a linear layout, in which the linear order avoids $(k + 1)$ -rainbows, that is, $k + 1$ independent edges that pairwise form a nested pair. The optimization goal is to determine the *queue number* of a graph, which is defined as the minimum value of k for which a k -queue layout is feasible. Recently, Dujmović et al. [J. ACM, 67(4), 22:1–38, 2020] showed that the queue number of planar graphs is at most 49, thus settling in the positive a long-standing conjecture by Heath, Leighton and Rosenberg. To achieve this breakthrough result, their approach involves three different techniques: (1) an algorithm to obtain 2-queue layouts of outerplanar graphs, (2) an algorithm to obtain 5-queue layouts of planar 3-trees, and (3) a decomposition of a planar graph into so-called tripods. In this work, we push further each of these techniques to obtain the first non-trivial improvement of the upper bound on the queue number of planar graphs from 49 to 42.

Keywords Linear layouts · Queue number · Planar graphs

✉ Michael Bekos
bekos@uoi.gr

Martin Gronemann
mgronemann@ac.tuwien.ac.at

Chrysanthi N. Raftopoulou
crisraft@mail.ntua.gr

¹ Department of Mathematics, University of Ioannina, University Campus, 45110 Ioannina, Greece

² Algorithms and Complexity Group, TU Wien, Favoritenstraße 9–11, 1040 Vienna, Austria

³ School of Applied Mathematics and Physical Sciences, National Technical University of Athens, Iroon Polytechniou 9, Zografou, 15780 Athens, Greece

1 Introduction

Linear layouts of graphs have a long tradition of study in different contexts, including graph theory and graph drawing, as they form a framework for defining different graph-theoretic parameters with several applications; see, e.g., [1]. In this regard, one seeks to find a total order of the vertices of a graph that reaches a certain optimization goal [2–4]. In this work, we focus on a well-studied type of linear layouts, called *queue layout* [5–8], in which the goal is to minimize the size of the largest *rainbow*, namely, a set of independent edges that are pairwise nested. Equivalently, the problem asks for a linear order of the vertices and a partition of the edges into a minimum number of queues (called *queue number*), such that no two independent edges in the same queue are nested [7]; see Fig. 1b.

Queue layouts of graphs were introduced by Heath and Rosenberg [7] in 1992 as the counterpart of *stack layouts* (widely known also as *book embeddings*), in which the edges must be partitioned into a minimum number of stacks (called *stack number*), such that no two edges in the same stack cross [9]; see Fig. 1c. Since their introduction, queue layouts of graphs have been a fruitful subject of intense research with several important algorithmic and combinatorial results proposed over the years [6, 8, 10–14]; for an introduction, we refer the reader to [15].

The most intriguing problem in this research field is undoubtedly the problem of specifying the queue number of planar graphs, that is, the maximum queue number over all planar graphs. This problem dates back to 1992, when Heath, Leighton and Rosenberg conjectured that the queue number of planar graphs is bounded [6]. Notably, despite the different efforts [11, 16, 17], this conjecture remained unanswered for more than two decades. That only changed in 2019 with a breakthrough result of Dujmović, Joret, Micek, Morin, Ueckerdt and Wood [5, 18], who managed to settle in the positive the conjecture, as they showed that the queue number of planar graphs is at most 49. The best-known corresponding lower bound is 4 due to Alam et al. [10].

It is immediate to see, however, that the gap between the currently best-known lower and upper bounds is rather large, which implies that the exact queue number of planar graphs is, up to the point of writing, still unknown. Note that this is in contrast with the stack number of planar graphs, which was recently shown to be exactly 4 [19, 20]. Also, the existing gap in the bounds on the queue number of planar graphs gives the intuition that it is unlikely the upper bound of 49 by Dujmović et al. [5] is tight, even though in the last two years no improvement appeared in the literature.




Fig. 1 a A planar straight-line drawing of the octahedron graph and b–c different linear layouts of it

1.1 Our Contribution

We verify the aforementioned intuition by reducing the upper bound on the queue number of planar graphs from 49 to 42 (see Theorem 7 in Sect. 5). This improvement requires non-trivial adjustments to each of the three algorithms that are involved in the approach by Dujmović et al. [5], namely: (i) an algorithm to obtain 2-queue layouts of outerplanar graphs by first computing straight-line drawings of them in which the y -distance of any two adjacent vertices is either 1 or 2 [21]; see Sect. 4.1, (ii) an algorithm to obtain 5-queue layouts of planar 3-trees [10]; see Sect. 4.2, and (iii) a decomposition of a planar graph into so-called tripods [5]; see Sect. 4.3. Although we assume familiarity with these algorithms, in Sect. 3 we outline their most important aspects. For preliminary notions and standard terminology refer to Sect. 2.

2 Preliminaries

A *drawing* of a graph maps each vertex to a distinct point of the Euclidean plane and each edge to a Jordan curve connecting its endpoints. A drawing of a graph is *planar*, if no two edges cross (except at common endpoints). A graph that admits a planar drawing is called *planar*; see Fig. 1a. A planar drawing partitions the plane into topologically connected regions called *faces*; its unbounded one is called *outer face*. A *combinatorial embedding* of a planar graph is an equivalence class of planar drawings that are pairwise topologically equivalent, that is, they define the same set of faces up to an orientation-preserving homeomorphism of the plane. A planar graph together with a combinatorial embedding is a *plane graph*. A graph that admits a planar drawing, in which all vertices are incident to its outer face, is called *outerplanar*. It is known that outerplanar graphs have treewidth at most 2. A planar graph with treewidth at most 3 is commonly referred to as a *planar 3-tree*. Planar 3-trees are subgraphs of *maximal planar 3-trees*, which are recursively defined as follows. In the base case, the graph is a 3-cycle. Otherwise, there exists a vertex of degree 3 whose removal yields a maximal planar 3-tree with one vertex less.

A linear layout of a graph is defined by a total order of its vertices and an objective over its edges. We denote the (total) vertex order of a simple undirected graph G by \prec , such that for any two vertices u and v of G , $u \prec v$ if and only if u precedes v in the order. Let F be a set of $k \geq 2$ independent edges (u_i, v_i) of G , that is, $F = \{(u_i, v_i) : i = 1, \dots, k\}$. If $u_1 \prec \dots \prec u_k \prec v_k \prec \dots \prec v_1$, then we say that the edges of F form a *k-rainbow*, while if $u_1 \prec v_1 \prec \dots \prec u_k \prec v_k$, then the edges of F form a *k-necklace*. The edges of F form a *k-twist*, if $u_1 \prec \dots \prec u_k \prec v_1 \prec \dots \prec v_k$; see Fig. 2. Finally, two independent edges that form a 2-rainbow (2-necklace, 2-twist) are referred to as *nested* (*disjoint*, *crossing*, respectively).

There are several types of linear layouts of graphs in the literature, with queue layouts being among the most studied of them. Formally, a *k-queue layout* of a graph G consists of a vertex order \prec of G and a partition of the edges of G into k sets of pairwise non-nested edges with respect to \prec , called *queues*. A preliminary result by Heath and Rosenberg [7] states that a graph admits a k -queue layout if and only if it admits a vertex order in which no $(k + 1)$ -rainbow is formed. The *queue number* of a




Fig. 2 Illustration of **a** a 3-rainbow, **b** a 3-necklace, and **c** a 3-twist




Fig. 3 Introducing a degree-2 vertex v along the upper envelope, when its two neighbors u and w are connected with **a** a span-1 edge, and **b** a span-2 edge

graph G , denoted by $qn(G)$, is the minimum k , such that G admits a k -queue layout. Accordingly, the queue number of a class of graphs is the maximum queue number over all its members.

3 Outline of the Involved Techniques

In this section, we summarize the main aspects of the three algorithms mentioned in Sect. 1 that are involved in the approach by Dujmović et al. [5] to achieve the bound of 49 on the queue number of planar graphs. The first is an algorithm by Dujmović, Pór, and Wood [21] to compute 2-queue layouts of outerplanar graphs (Sect. 3.1). The second one is by Alam et al. [10] to compute 5-queue layouts of planar 3-trees (Sect. 3.2). The last one is the actual algorithm by Dujmović et al. [5] to compute 49-queue layouts of planar graphs (Sect. 3.3).

3.1 Outerplanar Graphs

As already stated, the main ingredient of the algorithm by Dujmović, Pór and Wood [21] is the construction of a straight-line drawing $\Gamma(G)$ of a maximal outerplane graph G that can be transformed into a 2-queue layout of G . The recursive construction of $\Gamma(G)$ maintains the following invariant properties:

- (O.1) The cycle delimiting the outer face consists of two strictly x -monotone paths, referred to as *upper* and *lower envelopes*, respectively.
- (O.2) The y -coordinates of the endvertices of each edge differ by either 1 (*span-1* edge) or 2 (*span-2* edge).

To maintain Invariants (O.1) and (O.2), Dujmović et al. adopt an approach according to which, at each recursive step, a vertex of degree 2 is added to the already constructed

drawing; see Fig. 3. The base of the recursion consists of a triangle¹ that can be trivially drawn with two span-1 edges and one span-2 edge satisfying Invariants (O.1) and (O.2). Assume that G has $n > 3$ vertices. Since G is biconnected outerplane, it contains a vertex v of degree 2. Removing v yields a biconnected outerplane graph G' with $n - 1$ vertices, which admits a drawing $\Gamma(G')$ satisfying Invariants (O.1) and (O.2). By Invariant (O.1), none of the edges on the outer face is drawn vertically in $\Gamma(G')$. To obtain drawing $\Gamma(G)$ of G , vertex v is inserted into $\Gamma(G')$ as follows. Let u and w be the neighbors of v in G . Since G is maximal outerplane, (u, w) is an edge of G that lies on the outer face of $\Gamma(G')$. By Invariant (O.2), (u, w) is either a span-1 or a span-2 edge. Assume that (u, w) is along the upper envelope of $\Gamma(G')$; the case where it is along the lower envelope is symmetric. Assume first that (u, w) has span 1 in $\Gamma(G')$ and without loss of generality that $y(u) = y(w) + 1$. Then, vertex v is placed such that $y(v) = y(u) + 1$ and $x(v) = \frac{1}{2}(x(u) + x(w))$; see Fig. 3a. Hence, (u, v) and (v, w) have span 1 and 2, respectively, which implies that Invariant (O.2) is maintained. Since (the non-vertical) edge (u, w) of the upper envelope of $\Gamma(G')$ is replaced by the x -monotone path $u \rightarrow v \rightarrow w$ in $\Gamma(G)$, Invariant (O.1) is also maintained. Assume now that (u, w) has span 2 in $\Gamma(G')$ and without loss of generality that $y(u) = y(w) + 2$. Then, vertex v is placed such that $y(v) = y(u) - 1$ and $x(v) = \frac{3x(w) + x(u)}{4}$; see Fig. 3b. This implies that both edges (u, v) and (v, w) have span 1. Similarly to the previous case, one can argue that Invariants (O.1) and (O.2) are maintained.

Finally, drawing $\Gamma(G)$ is transformed into a 2-queue layout of G as follows:

- (i) for any two vertices u and v of G , $u < v$ if and only if either $y(u) > y(v)$, or $y(u) = y(v)$ and $x(u) < x(v)$ in $\Gamma(G)$, and
- (ii) the edges assigned to the first (second) queue in the layout are those with span 1 (span 2, respectively) in $\Gamma(G)$.

3.2 Planar 3-trees

Alam et al. [10] adopt a *peeling-into-levels* approach [22] to produce a 5-queue layout of a maximal plane 3-tree H . Initially, the vertices of H are partitioned into *levels* L_0, \dots, L_λ with $\lambda \geq 1$, such that L_0 -vertices are incident to the outer face of H , while L_{i+1} -vertices are in the outer face of the subgraph of H obtained by the removal of all vertices in L_0, \dots, L_i , for $i = 0, \dots, \lambda - 1$. The edges of H are partitioned into *level* and *binding* edges, depending on whether their endpoints are on the same or on consecutive levels respectively. Let H_i be the subgraph of H induced by the edges of level L_i , where $i = 0, \dots, \lambda$. Since H is a maximal plane 3-tree, each connected component of H_i is an internally triangulated outerplane graph. Therefore, it is embeddable in two queues. This implies that each connected component c of H_{i+1} (which is outerplane) lies in the interior of a triangular face of H_i . As a result there are exactly three vertices of H_i that are connected to vertices of c . The construction of the 5-queue layout of H satisfies the following invariant properties:

- (T.1) In the linear order $<_H$, all vertices of level L_i precede all vertices of level L_{i+1} for every $i = 0, \dots, \lambda - 1$;

¹ In [21], Dujmović et al. consider a single edge at the base of the recursion. To ease the presentation, we slightly modify their approach.

- (T.2) Vertices of each connected component of level L_i appear consecutively in \prec_H for every $i = 0, \dots, \lambda$;
- (T.3) All level edges of H are assigned to two queues denoted by Q_0 and Q_1 ;
- (T.4) For every $i = 0, \dots, \lambda - 1$, the binding edges between vertices of L_i and L_{i+1} are assigned to three queues Q_2, Q_3 and Q_4 as follows. For each connected component c of H_{i+1} , let x, y and z be its three neighbors in H_i so that $x \prec_H y \prec_H z$. Then, the binding edges between L_i and L_{i+1} with one endpoint in c are assigned to Q_2, Q_3 or Q_4 if their other endpoint is x, y or z , respectively.

3.3 General Planar Graphs

Dujmović et al. [5] present an algorithm to produce a 49-queue layout of a (general) plane graph. Central in their algorithm is the notion of an H -partition², defined as follows. Given a graph G , an H -partition of G is a partition of the vertices of G into sets A_x with $x \in V(H)$, called *bags*, such that for each edge (u, v) of G with $u \in A_x$ and $v \in A_y$ either $x = y$ holds or (x, y) is an edge of H . In the former case, (u, v) is called *intra-bag* edge, while in the latter case it is called *inter-bag* edge. A *BFS-layering* of G is a partition $\mathcal{L} = (V_0, V_1, \dots)$ of its vertices according to their distance from a specific vertex of G . In other words, it is a special type of H -partition, where H is a path and each bag V_i corresponds to a *layer*. In this special type of H -partition, an intra-bag edge is called *intra-layer*, while an inter-bag edge is called *inter-layer*³. An H -partition has *layered-width* ℓ with respect to a BFS-layering \mathcal{L} if each bag of H has at most ℓ vertices on each layer of \mathcal{L} .

The following lemma is the main tool of the algorithm by Dujmović et al. [5], since it yields a construction of a queue-layout of a graph G , whose number of queues depends on the layered-width of an H -partition of G and the queue number of H .

Lemma 1 (Dujmović et al. [5]) For all graphs G and H , if H admits a k -queue layout and G has an H -partition of layered-width ℓ with respect to some layering $\mathcal{L} = (V_0, V_1, \dots)$ of G , then G admits a $(3\ell k + \lfloor \frac{3}{2}\ell \rfloor)$ -queue layout using vertex order $\vec{V}_0, \vec{V}_1, \dots$, where \vec{V}_i is some order of V_i . In particular,

$$\text{qn}(G) \leq 3\ell \text{qn}(H) + \lfloor \frac{3}{2}\ell \rfloor. \tag{1}$$

In the original proof of Lemma 1 (see [5, Lemma 9]), the order of the vertices of G on each layer of \mathcal{L} is defined as follows. Let x_1, \dots, x_h be the vertices of H as they appear in a k -queue layout $QL(H)$ of H , that is, $x_1 \prec_H \dots \prec_H x_h$, and let A_{x_1}, \dots, A_{x_h} be the corresponding bags of the H -partition. Then, the order \vec{V}_i of each layer V_i of \mathcal{L} with $i \geq 0$ is:

$$\vec{V}_i = A_{x_1} \cap V_i, A_{x_2} \cap V_i, \dots, A_{x_h} \cap V_i$$

² To avoid confusion with notation used earlier, note that, in the scope of the algorithm by Dujmović et al. [5], graph H denotes a plane 3-tree, as we will shortly see.

³ Dujmović et al. [5] refer to the intra- and inter-layer edges as intra- and inter-level edges, respectively. We adopt the terms intra- and inter-layer edges to avoid confusion with the different type of leveling used in the algorithm of Alam et al. [10].

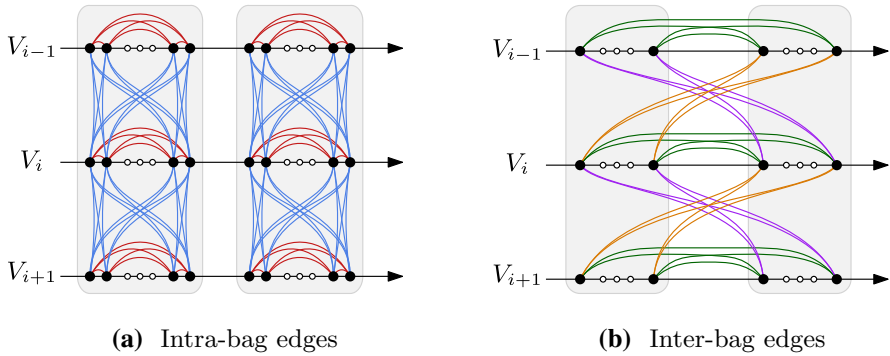


Fig. 4 Illustration of **a** Intra-bag edges; the intra-layer ones are red, while the inter-layer ones are blue, and **b** inter-bag edges; the intra-layer ones are green, while the inter-layer ones are purple (forward) and orange (backward). In both subfigures, a bag is illustrated as a gray box, in which vertices of different layers (that are part of the bag) are drawn at different horizontal lines; the white vertices indicate that the intersection of a layer with the content of a bag may consist of more than four vertices; the edges towards these vertices have not been illustrated (Color figure online)

where each subset $A_{x_j} \cap V_i$ is ordered arbitrarily, for $j = 1, \dots, h$. This gives the total order \prec_G for the vertices of G . The edge-to-queue assignment, which completes the construction of the queue layout $QL(G)$ of G , exploits the following two properties:

- (P.1) Two intra-bag edges of G nest in \prec_G , only if they belong to the same bag; see blue and red edges in Fig. 4a.
- (P.2) Two inter-layer edges of G nest in \prec_G , only if their endpoints belong to the same pair of consecutive layers of \mathcal{L} ; see blue, purple and orange edges in Fig. 4.

Recall that the edges of G are classified into four categories given by the bags of the H -partition and the layers of \mathcal{L} , namely, intra-layer intra-bag edges, inter-layer intra-bag edges, intra-layer inter-bag edges and inter-layer inter-bag edges. The edge-to-queue assignment is based on these four categories. We start with the intra-bag edges of G , that is, edges whose endpoints belong to the same bag; see Fig. 4a.

- (E.1) Intra-layer intra-bag edges of G are assigned to at most $\lfloor \frac{\ell}{2} \rfloor$ queues, as the queue number of K_ℓ is $\lfloor \frac{\ell}{2} \rfloor$ [22]; see red edges in Fig. 4a.
- (E.2) Inter-layer intra-bag edges of G are assigned to at most ℓ queues, as the queue number of $K_{\ell,\ell}$ is ℓ , when all vertices of the first bipartition precede those of the second; see blue edges in Fig. 4a.

The remaining edges of G are the inter-bag edges that connect vertices of different bags; see Fig. 4b. Inter-layer inter-bag edges are further partitioned into two categories. Let (u, v) be an inter-layer inter-bag edge with $u \in A_x \cap V_i$ and $v \in A_y \cap V_{i+1}$, for some $i \geq 0$ and $x \neq y$. Then (u, v) is *forward*, if $x \prec_H y$ holds in $QL(H)$, otherwise, (u, v) is *backward*; see purple and orange edges in Fig. 4b, respectively. For all inter-bag edges, in total, $3\ell k$ queues suffice.

- (E.3) Intra-layer inter-bag edges of G are assigned to at most ℓk queues; on each layer, an edge of H corresponds to a subgraph of $K_{\ell,\ell}$, where all vertices of the first bipartition precede those of the second; see green edges in Fig. 4b.

- (E.4) Forward inter-layer inter-bag edges of G are assigned to at most ℓk queues; for two consecutive layers, an edge of H corresponds to a subgraph of $K_{\ell, \ell}$, where all vertices of the first bipartition precede those of the second; see purple edges of Fig. 4b.
- (E.5) Symmetrically all backward inter-layer inter-bag edges of G are assigned to at most ℓk queues; see orange edges in Fig. 4b.

The next property follows from the edge-to-queue assignment of inter-bag edges of G . Details can be found in the proof of Lemma 9 in [5].

- (P.3) For $1 \leq i \leq r$, let (u_i, v_i) be an edge of G , such that $u_i \prec_G v_i, u_i \in A_{x_i}$ and $v_i \in A_{y_i}$. If all these r edges belong to one of (E.3)-(E.5) and form an r -rainbow in \prec_G , while edges $(x_1, y_1), \dots, (x_r, y_r)$ of H are assigned to the same queue in $QL(H)$, then $r \leq \ell$ and either u_1, \dots, u_r or v_1, \dots, v_r belong to the same bag of the H -partition of G .

In order to bound the parameters k and ℓ used in Lemma 1, the fact that G is maximal plane is being exploited. To this end, a few more ingredients are needed. A *vertical path* of G in a BFS-layering \mathcal{L} is a path $P = v_0, \dots, v_\kappa$ of G consisting only of edges of the BFS-tree of \mathcal{L} and such that if v_0 belongs to V_i in \mathcal{L} , then v_j belongs to V_{i+j} , with $j = 1, \dots, \kappa$. Further, we say that v_0 and v_κ are the *first* and *last* vertices of P . A *tripod* of G consists of up to three pairwise vertex-disjoint vertical paths in \mathcal{L} whose last vertices form a clique of size at most 3 in G . We refer to this clique as the *base* of the tripod. Dujmović et al. [5] describe a recursive algorithm that partitions the vertices of G into tripods, whose bases are triangular faces of G . In particular, they show that for any BFS-layering \mathcal{L} , G admits an H -partition with the following properties:

- (P.4) H is a planar 3-tree and thus $QL(H)$ is a k -queue layout with $k \leq 5$ [10].
- (P.5) Its layered-width ℓ is at most 3, since each bag induces a tripod in G .

Properties (P.4) and (P.5) along with Eq. (1) imply that the queue number of planar graphs is at most $3 \cdot 3 \cdot 5 + \lfloor \frac{3}{2} \cdot 3 \rfloor = 49$.

4 Refinements of the Involved Techniques

In this section, we present refinements of the algorithms outlined in Sect. 3 that will allow us to reduce the upper bound on the queue number of planar graphs from 49 to 42.

4.1 Outerplanar Graphs

The initial algorithm by Dujmović et al. [21], outlined in Sect. 3.1, produces outerplanar straight-line drawings that satisfy Invariants (O.1) and (O.2). Our modification ensures two additional properties that are stated in Lemma 2. To achieve this, besides (O.1) and (O.2), we maintain a third invariant:

- (O.3) The lower envelope consists of a single edge.

To maintain Invariant (O.3), we observe that a biconnected maximal outerplane graph G with at least four vertices, contains at least two non-adjacent degree-2 vertices. Let x be such a degree-2 vertex of G , which we assume to be fixed in the recursive construction of $\Gamma(G)$. In particular, by our previous observation, we can always remove a degree-2 vertex that is different from x at every recursive step. This ensures that x will eventually appear in the triangle \mathcal{T} that is drawn at the base of the recursion. We draw \mathcal{T} , such that x is its bottommost vertex while its two incident edges are of span-1 and span-2 as follows. Let y and z be the other two vertices of \mathcal{T} . We draw x , y and z at points $(2, 0)$, $(0, 1)$ and $(1, 2)$, respectively. This implies that (x, y) forms the lower envelope of \mathcal{T} , while (y, z) and (z, x) form the upper one. This guarantees that Invariant (O.3) is maintained at the base of the recursion. Let G' be the subgraph of G obtained by removing a degree-2 vertex v from G with $v \neq x$. We recursively compute a drawing $\Gamma(G')$ of G' that satisfies Invariant (O.3), such that the edge (x, y) of \mathcal{T} forms the lower envelope of $\Gamma(G')$. Since (x, y) belongs to the outer face of G , it follows that v is incident to two vertices of the upper envelope of $\Gamma(G')$. So, after the addition of v to $\Gamma(G')$ in order to obtain $\Gamma(G)$, the lower envelope of $\Gamma(G)$ consists only of edge (x, y) , as desired. This concludes the construction of $\Gamma(G)$ satisfying all three Invariants (O.1), (O.2) and (O.3).

In the following, we focus on proving two properties of drawing $\Gamma(G)$, that are summarized in Lemma 2. Let $\langle u, v, w \rangle$ be a face of $\Gamma(G)$ such that $y(u) - y(w) = 2$ and $y(u) - y(v) = y(v) - y(w) = 1$. We refer to vertices u , v and w as the *top*, *middle* and *bottom* vertex of the face, respectively⁴.

Lemma 2 *Let G be a biconnected maximal outerplane graph, and let $\Gamma(G)$ be an outerplanar drawing of G satisfying Invariants (O.1)–(O.3). Then, each vertex of G is*

- (a) *the top vertex of at most two triangular faces of $\Gamma(G)$, and*
- (b) *the middle vertex of at most two triangular faces of $\Gamma(G)$.*

Proof For (a), consider a vertex u of G . If u is the top vertex of a face, then u is incident to a span-2 edge (u, v) with $y(u) > y(v)$. By Invariant (O.3), u is a successor of v in the recursive approach by Dujmović et al. [21], namely, when u is placed in the recursively constructed drawing, vertex v belongs to its upper envelope. We now claim that u cannot be incident to two span-2 edges (u, v) and (u, v') with $y(u) > y(v)$ and $y(u) > y(v')$, that is, with the properties mentioned above; this claim implies part (a) of the lemma. Assuming the contrary, by Invariant (O.2), when u is placed in the drawing at most one edge incident to u has span 2. So, at most one of (u, v) and (u, v') is drawn when u is placed in the drawing, which implies that at least one of v or v' , say v' , is a successor of u . Thus, $y(u) < y(v')$ holds; a contradiction.

For (b), assume for a contradiction that G contains a vertex u , which is the middle vertex of three triangular faces, say \mathcal{T}_1 , \mathcal{T}_2 and \mathcal{T}_3 , of $\Gamma(G)$. For a significantly small constant $\varepsilon > 0$, let $p_1 = (x(u) - \varepsilon, y(u))$ and $p_2 = (x(u) + \varepsilon, y(u))$, and consider the two horizontal rays r_1 and r_2 emanating from vertex u , such that r_1 and r_2 contain points p_1 and p_2 , respectively. Since u is middle vertex for \mathcal{T}_1 , \mathcal{T}_2 and \mathcal{T}_3 , it follows by Invariant (O.2) that the span-2 edge of each of \mathcal{T}_1 , \mathcal{T}_2 and \mathcal{T}_3 crosses either r_1 or

⁴ Alam et al. [10] refer to the middle vertex of a triangular face in $\Gamma(G)$ as its *anchor*.

r_2 ; by the choice of ε , we may assume that if a span-2 edge of one of the faces $\mathcal{T}_1, \mathcal{T}_2$ and \mathcal{T}_3 crosses r_1 (r_2), then point p_1 (p_2 , respectively) lies in the interior of this face. On the other hand, by outerplanarity, each of the points p_1 and p_2 can be contained in at most one of $\mathcal{T}_1, \mathcal{T}_2$ and \mathcal{T}_3 ; a contradiction. \square

Note that part (a) of Lemma 2 requires Invariant (O.3), that is, it does not necessarily hold for all drawings obtained by the algorithm by Dujmović et al. [21]. On the other hand, part (b) of Lemma 2 is a property of the original algorithm as it holds for drawings that do not necessarily satisfy Invariant (O.3).

4.2 Planar 3-trees

To maintain Invariant (T.3), Alam et al. [10] use the algorithm by Dujmović et al. [21] to assign the level edges of L_0, \dots, L_λ of the input plane 3-tree H to two queues \mathcal{Q}_0 and \mathcal{Q}_1 , since on each level these edges induce a (not necessarily connected) outerplane graph. In our approach we adopt the modification for the algorithm by Dujmović et al. [21] introduced in Sect. 4.1. As Invariants (O.1) and (O.2) are preserved, queues \mathcal{Q}_0 and \mathcal{Q}_1 still suffice for the level edges of H . We exploit Invariant (O.3), that our modification additionally satisfies, to derive an important property of the produced queue layout, as stated in Lemma 3.

To maintain Invariant (T.4), Alam et al. [10] adopt the following assignment scheme for the binding edges between L_j and L_{j+1} to queues $\mathcal{Q}_2, \mathcal{Q}_3$ and \mathcal{Q}_4 , for each $j = 0, \dots, \lambda - 1$. Consider a binding edge (u, v) with $u \in L_j$ and $v \in L_{j+1}$. Then, vertex u belongs to a connected component C_u of the subgraph H_j of H induced by the level- L_j vertices. Similarly, vertex v belongs to a connected component C_v of the corresponding subgraph H_{j+1} of H . Furthermore, C_u is outerplane and its 2-queue layout has been computed by the algorithm by Dujmović et al. [21], while C_v resides in the interior of a triangular face \mathcal{T}_v of C_u in the embedding of H , such that u is on the boundary of \mathcal{T}_v . In the edge-to-queue assignment scheme by Alam et al. [10], edge (u, v) is assigned to $\mathcal{Q}_2, \mathcal{Q}_3$ or \mathcal{Q}_4 if and only if u is the middle, top or bottom vertex of \mathcal{T}_v , respectively [10]. For a vertex $u \in L_j$ with $j = 0, \dots, \lambda - 1$, we denote by $N_2(u), N_3(u)$ and $N_4(u)$ the neighbors of u in L_{j+1} such that the edges connecting them to u are assigned to queues $\mathcal{Q}_2, \mathcal{Q}_3$ and \mathcal{Q}_4 , respectively. The next property follows from Lemma 2.

Lemma 3 *For $j = 0, \dots, \lambda - 1$, let $u \in L_j$ be a vertex of H in our modification of the peeling-into-levels approach by Alam et al. [10]. Then, the vertices of $N_2(u)$ precede those of $N_3(u)$ in \prec_H . Also, the vertices of $N_2(u)$ belong to at most two connected components of H_{j+1} (residing within distinct faces of H_j); the same holds for the vertices of $N_3(u)$.*

Proof The first part of the lemma is proven in [10], although it is not explicitly stated as part of Invariants (T.1)–(T.4). For the second part, consider a binding edge (u, v) with $v \in N_2(u)$; a similar argument applies when $v \in N_3(u)$. By the definition of $N_2(u)$, vertex u is the top vertex of the triangular face \mathcal{T}_v of H_j , in which the connected component C_v of H_{j+1} that contains v resides. Since by Lemma 2a vertex u can be




Fig. 5 Tripod τ with parents τ_1, τ_2 and τ_3 . In **a** τ shares no vertex with τ_1, τ_2, τ_3 and in **b** τ shares a vertex with τ_3

the top vertex of at most two triangular faces of H_j , there exist at most two connected components of H_{j+1} , to which vertex v can belong. \square

4.3 General Planar Graphs


Dujmović et al. [5] recursively compute the bags (that is, the tripods) of the H -partition. In this recursive approach, each newly discovered tripod τ is adjacent to at most three other tripods τ_1, τ_2 and τ_3 already discovered. We say that τ_1, τ_2 and τ_3 are the *parents* of τ ; see Fig. 5a. Also, each non-empty vertical path of τ is connected to only one of its parents via an edge of the BFS-tree used to construct the BFS-layering \mathcal{L} (black in Fig. 5a). This property gives rise to at most three sub-instances (gray in Fig. 5a), which are processed recursively to compute the final tripod decomposition.

Note that, in general, one or more vertical paths of a tripod may have no vertices; see Fig. 5b. We refer to such a tripod as *degenerated*. In addition, we call a degenerated tripod *empty* if it shares a vertex with each of its parents. Hence, empty tripods have no vertices and as a result do not contribute to H . To avoid considering degenerated tripods that are not empty explicitly, we introduce an appropriate augmentation of the input graph. More precisely, let f be the base triangle of a non-empty degenerated tripod τ , and let T be the BFS tree that was used to compute τ . We augment G by adding a triangle f' in the interior of face f . By an appropriate triangulation, we can obtain a triangulated plane graph G' that contains G as a subgraph, such that all vertices of f' are leaves of T , which gives rise to a BFS tree T' of G' . In this way, face f of G gets three leaves of T' inside it, one attached to each vertex of f . Hence, the base triangle of the corresponding tripod in G' is not degenerated. This allows us to assume without loss of generality in the following that, in the tripod decomposition, either each tripod is empty or non-degenerated, that is, no vertical path of it is empty.

For $i = 1, 2, 3$, let p_i^1, p_i^2 and p_i^3 be the three vertical paths of tripod τ_i . Up to renaming, we assume that τ lies in the cycle bounded by (parts of) $p_1^1, p_1^2, p_2^1, p_2^2, p_3^1$ and p_3^2 as in Fig. 5a. The next properties follow by planarity and the BFS-layering:

- (P.6) There is no edge connecting a vertex of τ to a vertex of p_i^3 for $i = 1, 2, 3$; see the dotted edge in Fig. 6.

Fig. 6 Illustration for Properties (P.6) and (P.8). Dotted and dashed edges would introduce crossings



- (P.7) Let v_i^p be the vertex of vertical path p of τ on layer V_i of \mathcal{L} . For two vertical paths p and q of τ , edge (v_i^p, v_j^q) belongs to G only if $|i - j| \leq 1$.
- (P.8) For vertical paths p and q of τ , at most one of the edges (v_i^p, v_{i+1}^q) and (v_{i+1}^p, v_i^q) exists in G ; see the dashed edges of Fig. 6.

Note that Properties (P.6), (P.7), and (P.8) hold even if τ has less than three parents, or if the cycle bounding the region of τ does not contain two vertical paths of each parent tripod.

In the original algorithm by Dujmović et al. [5], each vertex v_τ in H corresponds to a non-empty tripod τ in G , and an edge $(v_\tau, v_{\tau'})$ exists in H , if τ is a parent of τ' in G , or vice versa. By construction, H is a connected partial planar 3-tree, which is arbitrarily augmented to a maximal planar 3-tree H' . Then, the algorithm by Alam et al. [10] is employed to compute a 5-queue layout of H' as required by Lemma 1. Here, we adopt a particular augmentation to guarantee an additional property for the graph H' (see Lemma 4). Similarly to the original approach, we contract the vertices of each non-empty tripod τ of G to a single vertex v_τ . However, in our modification, we keep self-loops that occur when an edge of G has both endpoints in τ (unless this edge belongs to one of its vertical paths), as well as, parallel edges that occur when two vertices of τ have a common neighbor not in τ . Two important properties of this contraction scheme that follow directly from planarity are given below; see Fig. 7.

- (P.9) The edges around each contracted vertex v_τ appear in the same clockwise cyclic order as they appear in a clockwise traversal along τ in G .
- (P.10) The edges having at least one endpoint on the same vertical path of τ appear consecutively around v_τ ; see Fig. 7b.

Note that the aforementioned contraction scheme creates self-loops and pairs of parallel edges. We focus on *homotopic* self-loops and pairs of parallel edges, which contain no vertex either in the interior or in the exterior region that they define. We remove such self-loops and keep one copy of such parallel edges. Then, we subdivide each (non-homotopic) self-loop twice, and for each edge with multiplicity $m > 1$, we subdivide all but one of its copies. In this way, each vertex v_τ corresponding to a tripod τ in G always lies in the interior of a separating 3-cycle C that contains all the vertices corresponding to its parent tripods on its boundary. To see this, observe that if τ has three parent τ_1, τ_2 and τ_3 , then C is formed by v_{τ_1}, v_{τ_2} and v_{τ_3} , such that v_τ is




Fig. 7 **a** A tripod τ in G where the edges incident to its three vertical paths are drawn dotted, and **b** the result after contracting τ to v_τ

connected to each of them. If τ has two parent τ_1 and τ_2 , then C is formed by v_{τ_1} , v_{τ_2} and a subdivision vertex, such that v_τ is connected to v_{τ_1} and v_{τ_2} . Finally, if τ has only one parent τ_1 , then C is formed by v_{τ_1} and two subdivision vertices, such that v_τ is connected to v_{τ_1} . Since subdivision vertices are of degree 2, the result is a simple (possibly not maximal) planar 3-tree, which is a supergraph of H . To derive H' , we augment it to maximal by adding edges, while maintaining its embedding [23]. We are now ready to state the additional property that H' has.

Lemma 4 *Let v_τ and v_{τ_p} be two vertices of H' that correspond to a tripod τ and to a parent tripod τ_p of τ in G . If L_i and L_j are the levels of v_τ and v_{τ_p} in the peeling-into-levels approach for H' , then $i \geq j$.*

Proof Let C be the inclusion-minimal separating 3-cycle of H' containing v_τ in its interior and all vertices that correspond to the parent tripods of τ on its boundary. Let L_k , L_l and L_m be the levels of the three vertices of C , with $k \leq l \leq m$, in the peeling-into-levels approach for H' . As τ_p is a parent of τ , $j \in \{k, l, m\}$ holds. Since C is a 3-cycle and since each edge in the peeling-into-level approach is either level or binding, $m \leq k + 1$ holds. The fact that v_τ lies in the interior of C and is connected to each of the vertices in H' corresponding to its parent tripods in G , implies that v_τ is on level L_{k+1} , that is, $i = k + 1$. So, $j \leq m \leq k + 1 = i$ holds. \square

As in the original algorithm by Dujmović et al. [5], we compute a 5-queue layout $QL(H')$ of H' . However, we use our modification of the algorithm by Alam et al. [10] described in Sect. 4.2. Denote by x_1, \dots, x_h the vertices of the subgraph H of H' as they appear in $QL(H')$, that is, we ignore subdivision vertices introduced when augmenting H to H' . Let also Q_0, \dots, Q_4 be the queues of $QL(H')$ as described in Invariants (T.1)–(T.4). To compute the linear layout $QL(G)$ of G , we use Lemma 1, which orders the vertices of each layer V_i of the BFS-layering \mathcal{L} , with $i \geq 0$, as:

$$\vec{V}_i = A_{x_1} \cap V_i, A_{x_2} \cap V_i, \dots, A_{x_h} \cap V_i,$$

where A_{x_1}, \dots, A_{x_h} are the bags, that is, the tripods, of the H -partition of G . Unlike in the original algorithm, we do not order the vertices in each subset $A_{x_j} \cap V_i$, with $j \in \{1, \dots, h\}$, arbitrarily. Instead, we carefully choose their order; we describe this

choice in the rest of this section. Let τ be the tripod of bag A_{x_j} . Then, $A_{x_j} \cap V_i$ contains at most one vertex of each vertical path of τ . We will order the three vertical paths of τ , which defines the order of the (at most three) vertices of $A_{x_j} \cap V_i$ for every $i \geq 0$.

Let L_l be the level of v_τ in the peeling-into-levels of H' , with $0 \leq l < \lambda$, that is, we assume that v_τ is not a vertex of the last level L_λ of H' . By Lemma 3, there are at most four connected components c_s^1, c_s^2, c_t^1 and c_t^2 of the subgraph H'_{l+1} of H' induced by the vertices of L_{l+1} , such that the edges connecting v_τ to vertices of c_s^1 and c_s^2 (c_t^1 and c_t^2) belong to \mathcal{Q}_2 (\mathcal{Q}_3 , respectively). Let c be one of c_s^1, c_s^2, c_t^1 and c_t^2 ; since c is a connected component of H'_{l+1} , it may contain vertices that correspond to tripods in G (that is, not subdivision vertices introduced while augmenting H to H'). We refer to the union of these vertices of G as the *tripod-vertices* of c , namely, the tripod-vertices of c are the vertices of G contained in the tripods of G that correspond to the vertices of c . By Invariant (T.1), v_τ precedes the vertices of c_s^1, c_s^2, c_t^1 and c_t^2 in $QL(H')$. Also by Invariant (T.2), we may assume that the vertices of c_s^1 (c_t^1) precede those of c_s^2 (c_t^2 , respectively). Additionally, Lemma 3 ensures that the vertices of c_s^2 precede those of c_t^1 .

Since an edge $(v_\tau, v_{\tau'})$ exists in H , if τ is a parent of τ' , or vice versa, by Lemma 4, for each vertex $v_{\tau'}$ of H that is a neighbor of v_τ in one of c_s^1, c_s^2, c_t^1 and c_t^2 , it follows that τ is a parent of τ' . By Property (P.6), there is a vertical path of τ in G , say p , such that no tripod-vertex of c_s^2 is adjacent to it in G . Similarly, there is a vertical path of τ in G , say q , such that no tripod-vertex of c_t^2 is adjacent to it in G . Note that p and q might be the same vertical path of τ .

We now describe the order of the three vertical paths of τ . We only specify the first one; the other two can be arbitrarily ordered:

- (i) if the tripod-vertices of c_s^1 and c_s^2 are connected to all three vertical paths of τ in G , then p is the first vertical path of τ ;
- (ii) if the tripod-vertices of c_t^1 and c_t^2 are connected to all three vertical paths of τ in G , then q is the first vertical path of τ ;
- (iii) otherwise, any vertical path of τ can be first.

Before proving formally that Cases (i) and (ii) cannot apply simultaneously (which implies that the order of the vertical paths of τ is well-defined), we state two important implications of choosing the first vertical path of τ .

- (P.11) By assuming that all vertices of c_s^1 precede those of c_s^2 , if tripod-vertices of c_s^1 and c_s^2 are connected to all three vertical paths of τ , then tripod-vertices of c_s^2 are not connected to the first vertical path of τ .
- (P.12) Also, under the assumption that all vertices of c_t^1 precede those of c_t^2 , if tripod-vertices of c_t^1 and c_t^2 are connected to all three vertical paths of τ , then tripod-vertices of c_t^2 are not connected to the first vertical path of τ .

We now prove that Cases (i) and (ii) in the selection of the first vertical path of τ cannot apply simultaneously. Assume for a contradiction that the tripod-vertices of c_s^1 and c_s^2 are connected to all three vertical paths of τ , and also that the tripod-vertices of c_t^1 and c_t^2 are connected to all three vertical paths of τ . Since v_τ is the top vertex of the two triangular faces f_t^1 and f_t^2 that contain c_t^1 and c_t^2 in their interior in $\Gamma(H')$, it follows that the edges connecting v_τ to c_t^1 and c_t^2 appear consecutively around v_τ in the cyclic




Fig. 8 Illustrations for: **a–b** the fact that Cases (i) and (ii) in the selection of the first vertical path of τ cannot apply simultaneously, and **c** the proof of Lemma 5

order of the binding edges between levels L_l and L_{l+1} of H' ; refer to the blue edges of Fig. 8a. Also, since v_τ is the middle vertex of the two triangular faces f_s^1 and f_s^2 that contain c_s^1 and c_s^2 in their interior in $\Gamma(H'_l)$, we can assume that the edges connecting v_τ to c_s^1 immediately precede the edges connecting v_τ to c_t^1 and c_t^2 , while the edges connecting v_τ to c_s^2 immediately follow them; refer to the green edges of Fig. 8a. Note that since we have assumed that the vertices of c_s^1 precede the vertices of c_s^2 , our assumption on the order of the edges around v_τ is a property of the algorithm by Alam et al. [10]. Now, if the tripod-vertices of c_s^1 and c_s^2 are connected to all three vertical paths of τ , and the same holds for the tripod-vertices of c_t^1 and c_t^2 , then Property (P.10) implies that in G the edges connecting vertices of τ to tripod-vertices of c_t^1 and c_t^2 alternate with edges connecting vertices of τ to tripod-vertices of c_s^1 and c_s^2 ; see e.g. Fig. 8b. Hence, the cyclic order of the edges around v_τ in H' is not the same as the cyclic order of the edges incident to vertices of τ along a clockwise traversal of τ in G , contradicting Property (P.9). Hence, Case (i) and (ii) cannot apply simultaneously, as we initially claimed.

5 Reducing the Bound

To reduce the upper bound on the queue number of planar graphs, we turn our attention to the analysis of the required number of queues for the intra-bag inter-layer edges (E.2), as well as for the inter-bag edges (that is, either intra-layer (E.3) or inter-layer (E.4)–(E.5)); refer to Section 3.3.

For intra-bag inter-layer edges (E.2), the original algorithm by Dujmović et al. [5] uses three queues, since the layered width ℓ is 3; see blue edges in Fig. 4a. We prove that no three intra-bag inter-layer edges form a 3-rainbow, implying that the upper bound on the queue number of planar graphs can be improved from 49 to 48. Note that this first improvement does not require any of the modifications presented in Sect. 4.

Lemma 5 *In the queue layout computed by the algorithm by Dujmović et al. [5], no three intra-bag inter-layer edges of G form a 3-rainbow.*

Proof Assume to the contrary that there exist three such edges (u_1, v_1) , (u_2, v_2) and (u_3, v_3) forming a 3-rainbow in $QL(G)$ such that $u_1 <_G u_2 <_G u_3 <_G v_3 <_G v_2 <_G$

v_1 . By Property (P.1) these edges belong to the same bag A of the H -partition, while by Property (P.2) their endpoints belong to two consecutive layers V_i and V_{i+1} of \mathcal{L} , for some $i \geq 0$. The underlying linear order implies that vertices u_1, u_2 and u_3 belong to V_i , while v_1, v_2 and v_3 belong to V_{i+1} . The order of $A \cap V_i$ and $A \cap V_{i+1}$ is $u_1 <_G u_2 <_G u_3$ and $v_3 <_G v_2 <_G v_1$; see Fig. 8c. Let p_1, p_2 and p_3 be the first, second and third vertical paths of tripod τ forming A . Then, $(u_1, v_3) \in p_1, (u_2, v_2) \in p_2$ and $(u_3, v_1) \in p_3$. However, (u_1, v_1) and (u_3, v_3) contradict Property (P.8). Hence, there is no 3-rainbow formed by intra-bag inter-layer edges of G . \square

For inter-bag edges (E.3)–(E.5), the algorithm by Dujmović et al. [5] uses $3 \cdot 15$ queues, since the layered width ℓ is 3 and the best known upper bound on the queue number k of planar 3-trees is 5; see green, purple and orange edges in Fig. 4b. We exploit Property (P.3) to prove that $3 \cdot 13$ queues suffice for inter-bag edges. This further improves the upper bound on the queue number of planar graphs from 48 to 42. Note that this proof is more involved than the previous one and requires all modifications presented in Sect. 4.

Lemma 6 *In the queue layout computed by our modification of the algorithm by Dujmović et al. [5], the inter-bag edges of G do not form a 40-rainbow.*

Proof To prove the statement, it suffices to show that the inter-bag edges of G that belong to each of (E.3), (E.4) and (E.5) form at most a 13-rainbow. We focus on the edges of (E.3), that is, on the intra-layer inter-bag edges of G ; a similar argument applies for the other two types of edges, that is, forward and backward inter-layer inter-bag edges. We partition the edges of (E.3) into five parts E_3^0, \dots, E_3^4 , such that for $i = 0, \dots, 4$, part E_3^i contains the following type-(E.3) edges of G :

$$E_3^i = \{(u, v) \in (E.3) : u \in A_x, v \in A_y, (x, y) \in Q_i\}.$$

By Property (P.3), and since $\ell = 3$, the edges of E_3^i can form an r -rainbow with $r \leq 3$ (this property actually implies the initial bound of $3 \cdot 5 = 15$ queues for all intra-layer inter-bag edges of G by Dujmović et al. [5]). We next prove that neither the edges of E_3^2 nor the edges of E_3^3 form a 3-rainbow, which yields the desired reduction from 15 to 13 on the size of the maximum rainbow formed by all edges of (E.3). Assume that this is not true, say for the edges of E_3^2 ; a symmetric argument applies for the edges of E_3^3 . Let $(u_1, v_1), (u_2, v_2)$ and (u_3, v_3) be three edge of E_3^2 that form a 3-rainbow, such that $u_1 <_G u_2 <_G u_3 <_G v_3 <_G v_2 <_G v_1$. Since $(u_1, v_1), (u_2, v_2)$ and (u_3, v_3) are of type (E.3), that is, intra-layer edges, their endpoints u_1, u_2, u_3, v_1, v_2 and v_3 all belong to the same layer of the BFS-layering \mathcal{L} of G ⁵. Assuming that $u_i \in A_{x_i}$ and $v_i \in A_{y_i}$, it follows that (x_i, y_i) is an edge of H assigned to Q_2 and $x_i <_H y_i$, for $i = 1, 2, 3$.

⁵ Note that if we had assumed that $(u_1, v_1), (u_2, v_2)$ and (u_3, v_3) belonged to one of (E.4) or (E.5), that is, inter-layer edges, vertices u_1, u_2 and u_3 would all belong to one layer, while vertices v_1, v_2 and v_3 would belong to the next layer. In what follows, we will only use the fact that vertices u_1, u_2 and u_3 belong to the same layer and that vertices v_1, v_2 and v_3 belong to the same layer, without discriminating whether the two layers are the same or not.

By Property (P.3), either $x_1 = x_2 = x_3$ or $y_1 = y_2 = y_3$ holds. First, we prove that $x_1 = x_2 = x_3$ always holds. Assume for a contradiction that, without loss of generality, $x_1 \neq x_2$. Then, $y = y_1 = y_2 = y_3$ holds. Since the edges of \mathcal{Q}_2 are binding in the peeling-into-levels decomposition of H' and since $x_1 \prec_H x_2 \prec_H y$, vertices x_1 and x_2 belong to level L_j of H' , while y belongs to level L_{j+1} , for some $0 \leq j < \lambda$. By Property (T.4), y has only one neighbor on level L_j such that the edge connecting y to this neighbor is assigned to \mathcal{Q}_2 . This contradicts the fact that both edges (x_1, y) and (x_2, y) are assigned to \mathcal{Q}_2 , since $x_1 \neq x_2$ and x_1 and x_2 belong to L_j . Hence $x = x_1 = x_2 = x_3$ holds.

Now, since the edges of \mathcal{Q}_2 are binding edges in the peeling-into-levels decomposition of H' , and since, for $i = 1, 2, 3$, $x \prec_H y_i$, it follows that vertex x belongs to level L_j of H' , while vertex y_i belongs to level L_{j+1} , for some $0 \leq j < \lambda$. Let τ be the tripod of bag A_x and let τ_1, τ_2, τ_3 be the tripods of A_{y_1}, A_{y_2} and A_{y_3} , respectively. Note that vertices y_1, y_2 and y_3 are not necessarily distinct, that is, the corresponding tripods τ_1, τ_2 and τ_3 are not necessarily pairwise different. However, since $u_1 \prec_G u_2 \prec_G u_3$, and since u_1, u_2 and u_3 belong to the same layer in the BFS-layering \mathcal{L} of G , we can conclude that u_1, u_2 and u_3 belong to p_1, p_2 and p_3 , respectively, where p_1, p_2 and p_3 are the first, second and third vertical paths of τ . From the fact that, for $i = 1, 2, 3$, the edge (x, y_i) belongs to H , we have that τ is a parent tripod of τ_i or vice versa. Since x belongs to level L_j of H' and vertices y_1, y_2 and y_3 belong to level L_{j+1} , Lemma 4 implies that τ is a parent tripod of τ_1, τ_2 and τ_3 . By Lemma 3, there exist at most two connected components c_s^1 and c_s^2 of H'_{j+1} , such that the edges connecting x to vertices of c_s^1 and c_s^2 have been assigned to \mathcal{Q}_2 . Without loss of generality, by Invariant (T.2) we assume that all vertices of c_s^1 (if any) precede those of c_s^2 (if any).

We first argue that not all vertices y_1, y_2 and y_3 belong to only one of c_s^1 or c_s^2 , which further implies that both components c_s^1 and c_s^2 exist. Assume to the contrary that y_1, y_2 and y_3 belong to c_s^1 ; a similar argument applies for c_s^2 . By Property (P.6), tripod-vertices of component c_s^1 are not adjacent to the vertices of at least one vertical path p_i of τ in G , where $1 \leq i \leq 3$. In this case, however, we obtain a contradiction to the fact that v_i , which is a tripod-vertex of c_s^1 , is connect to u_i that belongs to p_i , i.e., edge (u_i, v_i) cannot exist in G .

Hence, not all vertices y_1, y_2 and y_3 belong to only one of c_s^1 or c_s^2 , as we initially claimed. In particular, since $y_3 \preceq_H y_2 \preceq_H y_1$ and since all vertices of c_s^1 precede those of c_s^2 , it follows that y_3 belongs to c_s^1 , while y_1 belongs to c_s^2 . On the other hand, vertices u_1, u_2 and u_3 belong to p_1, p_2 and p_3 , respectively, which implies that the vertices of c_s^1 and c_s^2 are connected to all vertical paths of τ . Thus, Property (P.11) applies, which states that the tripod-vertices of c_s^2 (including v_1) are not connected to the first vertical path p_1 of τ . However, this contradicts the fact that (u_1, v_1) is an edge of G , with $u_1 \in p_1$. Hence, the edges of E_3^2 cannot form a 3-rainbow. Using Property (P.12) instead of Property (P.11), we can symmetrically prove that the edges of E_3^3 cannot form a 3-rainbow. This completes the proof of the lemma. \square

We are now ready to state the main theorem of this section by combining Lemmas 5 and 6.

Theorem 7 *Every planar graph has queue number at most 42.*

6 Conclusions

In this work, we improved the upper bound on the queue number of planar graphs by refining the three techniques involved in the original algorithm [5]. We believe that our approach has the potential to further reduce the upper bound by at least 3 more queues (i.e., from 42 to 39). However, more elaborate arguments that exploit planarity in more depth are required, and several details have to be worked out. Towards narrowing the gap with the lower bound of 4, determining the exact queue number of planar 3-trees becomes critical; an improvement of the current upper bound of 5 (to meet the lower bound of 4) implies a direct improvement on the upper bound on the queue number of general planar graphs. On the other hand, to obtain a better understanding of the general problem, it is also reasonable to further examine subclasses of planar graphs, such as bipartite planar graphs or planar graphs with bounded degree (e.g., cubic planar graphs).

Funding Open access funding provided by HEAL-Link Greece. No funding was received for conducting this study.

Declarations

Competing Interests The authors have no competing interests to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Díaz, J., Petit, J., Serna, M.J.: A survey of graph layout problems. *ACM Comput. Surv.* **34**(3), 313–356 (2002). <https://doi.org/10.1145/568522.568523>
2. Barth, D., Pellegrini, F., Raspaud, A., Roman, J.: On bandwidth, cutwidth, and quotient graphs. *RAIRO Theor. Inform. Appl.* **29**(6), 487–508 (1995). <https://doi.org/10.1051/ita/1995290604871>
3. Chinn, P.Z., Chvatalova, J., Dewdney, A.K., Gibbs, N.E.: The bandwidth problem for graphs and matrices - a survey. *J. Graph Theory.* **6**(3), 223–254 (1982). <https://doi.org/10.1002/jgt.3190060302>
4. Horton, S.B., Parker, R.G., Borie, R.B.: On minimum cuts and the linear arrangement problem. *Discret. Appl. Math.* **103**(1–3), 127–139 (2000). [https://doi.org/10.1016/S0166-218X\(00\)00173-6](https://doi.org/10.1016/S0166-218X(00)00173-6)
5. Dujmovic, V., Joret, G., Micek, P., Morin, P., Ueckerdt, T., Wood, D.R.: Planar Graphs Have Bounded Queue-Number. *J. ACM.* **67**(4):22:1–22:38 (2020)
6. Heath, L.S., Leighton, F.T., Rosenberg, A.L.: Comparing queues and stacks as mechanisms for laying out graphs. *SIAM J. Discrete Math.* **5**(3), 398–412 (1992). <https://doi.org/10.1137/0405031>
7. Heath, L.S., Rosenberg, A.L.: Laying out graphs using queues. *SIAM J. Comput.* **21**(5), 927–958 (1992). <https://doi.org/10.1137/0221055>
8. Wiechert, V.: On the queue-number of graphs with bounded tree-width. *Electr. J. Comb.* **24**(1), P1.65 (2017)

9. Bernhart, F., Kainen, P.C.: The book thickness of a graph. *J. Comb. Theory Ser. B.* **27**(3), 320–331 (1979). [https://doi.org/10.1016/0095-8956\(79\)90021-2](https://doi.org/10.1016/0095-8956(79)90021-2)
10. Alam, J.M., Bekos, M.A., Gronemann, M., Kaufmann, M., Pupyrev, S.: Queue layouts of planar 3-trees. *Algorithmica* **82**(9), 2564–2585 (2020). <https://doi.org/10.1007/s00453-020-00697-4>
11. Bannister, M.J., Devanny, W.E., Dujmovic, V., Eppstein, D., Wood, D.R.: Track layouts, layered path decompositions, and leveled planarity. *Algorithmica* **81**(4), 1561–1583 (2019). <https://doi.org/10.1007/s00453-018-0487-5>
12. Bekos, M.A., Förster, H., Gronemann, M., Mchedlidze, T., Montecchiani, F., Raftopoulou, C.N., et al.: Planar graphs of bounded degree have bounded queue number. *SIAM J. Comput.* **48**(5), 1487–1502 (2019). <https://doi.org/10.1137/19M125340X>
13. Dujmović, V., Morin, P., Wood, D.R.: Layout of graphs with bounded tree-width. *SIAM J. Comput.* **34**(3), 553–579 (2005). <https://doi.org/10.1137/S0097539702416141>
14. Wood, D.R.: Queue layouts, tree-width, and three-dimensional graph drawing. In: Agrawal M, Seth A, editors. *FST TCS*. vol. 2556 of LNCS. Springer pp. 348–359 (2002)
15. Dujmović, V., Wood, D.R.: On linear layouts of graphs. *Discrete Math. Theor. Computer Sci.* **6**(2), 339–358 (2004)
16. Di Battista, G., Frati, F., Pach, J.: On the queue number of planar graphs. *SIAM J. Comput.* **42**(6), 2243–2285 (2013). <https://doi.org/10.1137/130908051>
17. Dujmović, V., Frati, F.: Stack and queue layouts via layered separators. *J. Graph Algorithms Appl.* **22**(1), 89–99 (2018). <https://doi.org/10.7155/jgaa.00454>
18. Dujmović, V., Joret, G., Micek, P., Morin, P., Ueckerdt, T., Wood, D.R.: Planar graphs have bounded queue-number. In: Zuckerman, D. (ed.) *FOCS*, pp. 862–875. IEEE Computer Society (2019)
19. Bekos, M.A., Kaufmann, M., Klute, F., Pupyrev, S., Raftopoulou, C.N., Ueckerdt, T.: Four pages are indeed necessary for planar graphs. *J. Comput. Geom.* **11**(1), 332–353 (2020)
20. Yannakakis, M.: Embedding planar graphs in four pages. *J. Comput. Syst. Sci.* **38**(1), 36–67 (1989). [https://doi.org/10.1016/0022-0000\(89\)90032-9](https://doi.org/10.1016/0022-0000(89)90032-9)
21. Dujmović, V., Pór, A., Wood, D.R.: Track layouts of graphs. *Discrete Math. Theor. Computer Sci.* **6**(2), 497–522 (2004)
22. Heath, L.S.: Embedding planar graphs in seven pages. In: *FOCS*. IEEE Computer Society pp. 74–83 (1984)
23. Kratochvíl, J., Vaner, M.: A note on planar partial 3-trees (2012) *CoRR* [arXiv:1210.8113](https://arxiv.org/abs/1210.8113)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.