


Bitonic st-Orderings for Upward Planar Graphs: Splits and Bends in the Variable Embedding Scenario

Item Type	Article
Authors	Angelini, Patrizio;Bekos, Michael A.;Förster, Henry;Gronemann, Martin
Citation	Angelini, Patrizio, Michael A. Bekos, Henry Förster, and Martin Gronemann. Bitonic st-Orderings for Upward Planar Graphs: Splits and Bends in the Variable Embedding Scenario. <i>Algorithmica</i> 85, 2667–2692 (2023).
DOI	https://doi.org/10.1007/s00453-023-01111-5
Publisher	Springer Nature
Rights	Attribution 4.0 International
Download date	2025-03-17 15:38:52
Item License	http://creativecommons.org/licenses/by/4.0/
Link to Item	https://hdl.handle.net/20.500.14490/325



Bitonic st -Orderings for Upward Planar Graphs: Splits and Bends in the Variable Embedding Scenario

Patrizio Angelini¹ · Michael A. Bekos² · Henry Förster³  · Martin Gronemann⁴

Received: 1 December 2021 / Accepted: 19 February 2023 / Published online: 11 March 2023
© The Author(s) 2023

Abstract

Bitonic st -orderings for st -planar graphs were introduced as a method to cope with several graph drawing problems. Notably, they have been used to obtain the best-known upper bound on the number of bends for upward planar polyline drawings with at most one bend per edge in polynomial area. For an st -planar graph that does not admit a bitonic st -ordering, one may split certain edges such that for the resulting graph such an ordering exists. Since each split is interpreted as a bend, one is usually interested in splitting as few edges as possible. While this optimization problem admits a linear-time algorithm in the fixed embedding setting, it remains open in the variable embedding setting. We close this gap in the literature by providing a linear-time algorithm that optimizes over all embeddings of the input st -planar graph. The best-known lower bound on the number of required splits of an st -planar graph with n vertices is $n - 3$. However, it is possible to compute a bitonic st -ordering without any split for the st -planar graph obtained by reversing the orientation of all edges. In terms of upward planar polyline drawings in polynomial area, the former translates into $n - 3$ bends, while the latter into no bends. We show that this idea cannot always be exploited by describing an st -planar graph that needs at least $n - 5$ splits in both orientations. We provide analogous bounds for graphs with small degree. Finally, we further investigate the relationship between splits in bitonic st -orderings and bends in upward planar polyline drawings with polynomial area, by providing bounds on the number of bends in such drawings.

Patrizio Angelini, Michael A. Bekos, Henry Förster and Martin Gronemann contributed equally to this work.

This work combines preliminary results from a paper presented at WG 2020 and a poster presented at GD 2021.

✉ Henry Förster
foersth@informatik.uni-tuebingen.de

Extended author information available on the last page of the article

Keywords Upward planar graphs · Bitonic st -orderings · Planar polyline drawings · Bend minimization

1 Introduction

Incremental drawing algorithms have a long history in the field of Graph Drawing. The central result of de Fraysseix, Pach and Pollack [1], who showed that every planar graph admits a planar straight-line drawing within quadratic area, marks the beginning of this line of research. In their seminal paper, they introduced the concept of *canonical ordering*, an ordering of the vertices that is used to drive their incremental drawing algorithm. In each step, one vertex at a time is placed, while it is ensured that certain invariants are satisfied. Another important result with respect to canonical orderings is by Kant [2]. While the original ordering is only defined for maximal planar graphs, he generalizes this concept to triconnected planar graphs. However, Kant's ordering is no longer a vertex ordering, instead it is an ordered partition of vertices. Later on, Harel and Sardas [3] show how one can further extend canonical orderings to the biconnected case.

Another type of vertex ordering that has its origins not in Graph Drawing, but finds its applications there [4, 5], is the so-called *st -ordering* [6]. However, *st -orderings* are not restricted to planar graphs, hence, the ordering is not related directly to the embedding of the underlying planar graph. This relation between a planar embedding and the ordering itself is established by the *bitonic st -orderings* [7], which have been used to solve various graph drawing problems, e.g., T-contact representations [7], L-drawings [8], finding universal slope sets [9]. Besides being a proper *st -ordering*, the definition takes the embedding into account and ensures that the vertex ordering has similar properties to a canonical ordering. Initially introduced for undirected graphs in [7], where it is shown that for every biconnected planar graph a bitonic *st -ordering* can be found in linear time, the concept has been extended to directed graphs [10].

The idea that led initially to the extension to directed graphs, namely the *st -planar graphs*, is rather simple. By slightly modifying the original algorithm of de Fraysseix, Pach and Pollack, one may use a bitonic *st -ordering* to obtain a planar straight-line drawing. Combined with the observation that a vertex is always drawn above its predecessors in the ordering, the resulting drawing is upward planar straight-line. However, not every *st -planar graph* admits such a bitonic *st -ordering*, but a full characterization is given in [10] that is based on the existence of so-called *forbidden configurations*. These configurations, however, can be eliminated by splitting certain edges in the graph, such that for the resulting graph one can then obtain the desired ordering. This technique is used to prove that every upward planar graph with n vertices, admits an upward planar polyline drawing with at most one bend per edge within $O(n^2)$ area. Moreover, the number of bends is at most $n - 3$, which is the best-known bound so far [10]. Hereby, each bend corresponds to a dummy vertex that has been introduced by splitting an edge. Note that in [10] an example is given that requires exactly $n - 3$ splits, which shows that this bound is tight.

In practice, one is interested in splitting as few edges as possible. In [10], a simple linear-time algorithm is described that finds the optimal set of edges to split. This

algorithm assumes the embedding of the underlying st -planar graph to be part of the input. Hence, it is only optimal in the fixed embedding scenario. Changing the embedding, however, may have a big impact on the required number of splits. Chaplick et al. [8] take a first step towards the variable embedding scenario by describing an algorithm based on SPQR-trees [11] to test whether an st -planar graph admits a bitonic st -ordering in any of its embeddings. In the positive case, their algorithm computes such an embedding and a corresponding bitonic st -ordering. Otherwise, one has to fall back to the fixed embedding algorithm for computing a set of edges to split. However, the number of edges to split may depend on the choice of the embedding, and a smaller set may be obtained from a different embedding.

1.1 Our Contribution

In this work, we first close the aforementioned gap in the literature by describing a linear-time algorithm to compute a smallest set of edges to split over all possible embeddings (see Theorem 1). Within the same time complexity, the algorithm outputs also a corresponding embedding.

Then, we turn our attention to upward planar polyline drawings obtained with the approach by Gronemann [10]. In this regard, Rettner [12] observed that an upward planar drawing can be obtained by reversing all edges of the graph, obtain an upward planar drawing for this reversed graph, and then mirror this drawing vertically. This idea stems from the observation that the example given in [10], which requires $n - 3$ splits, does not require any split at all when all edges have been reversed. In view of this property, one would naturally choose the orientation with the minimum number of splits. However, there exist limitations also in this approach, as there exist n -vertex st -planar graphs that require at least $\frac{3n}{4} - 3$ splits in each of the two orientations [12]. Still the question that arises is whether one of the two orientations always requires significantly less than $n - 3$ splits. We answer this question negatively by demonstrating n -vertex st -planar graphs that require $n - 5$ splits in each of the two orientations (see Theorem 3).

Note that the graphs supporting Theorem 3 are of maximum degree 6. On the other hand, if the maximum degree of the input st -planar graph is 3, then no split is required [13]. We show in Theorem 6 that $\frac{n}{2} - 2$ splits may be required for maximum degree 4 st -planar graphs. We finally prove that $\frac{n}{2}$ splits are sufficient even for degree-5 st -planar graphs (see Theorem 4).

Finally, we study lower bounds on the total number of bends in upward planar drawings under the polynomial-area requirement, independently of the required number of splits and of the allowed number of bends per edge. We show that $n - o(\log n)$ bends may be required for st -planar graphs of maximum degree 6 (see Theorem 8), while for maximum degree 4 the corresponding number of required bends is $\frac{n}{2} - o(\log n)$ (see Theorem 9). As a result, our findings imply that the upper bounds on the number of bends obtained by the approach by Gronemann are worst-case tight up to a logarithmic factor, even if more than one bend per edge is allowed.

Structure of the paper We give preliminaries in Sect. 2. Then, we devote Sect. 3 to describe our algorithm to minimize the number of splits in the variable embedding

setting. We complement our algorithm by providing upper and lower bounds on the number of splits in Sect. 4, and by discussing the relationship between bitonic embeddings and upward planarity in Sect. 5. Finally, we conclude with open problems in Sect. 6.

2 Preliminaries

Graph drawings and upward planarity A drawing Γ of a graph G maps the vertices of G to distinct points in the plane and the edges of G to simple Jordan arcs between their endpoints. Drawing Γ is *planar* if no two edges share an interior point. Planar drawings partition the plane into regions, called *faces*, whose boundaries consist of edges. The unbounded face is the *outer face*. An *embedding* is a class of drawings defining the sets of faces with the same boundaries.

A drawing Γ of a directed acyclic graph G is *upward planar* if for every edge (u, v) , vertex u lies below v in Γ and (u, v) is drawn as a y -monotone curve in Γ ; accordingly graph G is *upward planar* if it admits an upward planar drawing. A vertex v of a directed graph is a *source* (*sink*, resp.), if it only has outgoing (incoming, resp.) edges. A directed graph G is *st-planar* if it has a unique source s and a unique sink t such that there is an upward planar drawing of G , where s and t are incident to the outer face of it. In our definition, we assume that the edge (s, t) exists and is incident to the outer face. An upward planar embedding of an *st-planar* graph induces a left-to-right ordering of the incoming and outgoing edges of each vertex. We call the left-to-right ordered sequence of the neighbors of a vertex v connected with outgoing edges of v the *successor list* of v . Note that the faces of an *st-planar* graph have a unique source and a unique sink [14] connected by two paths. If one of these paths is a single edge, we call it *transitive*.

st-orderings An *st-ordering* of an *st-planar* graph is a linear ordering of its vertices with a prescribed vertex s being the first and a prescribed vertex t being the last vertex such that for every directed edge (u, v) , it holds that u precedes v [15]. Given an *st-ordering* of an embedded *st-planar* graph, the successor list of a vertex u is *monotonically increasing* (*decreasing*, resp.) if the outgoing neighbors of u appear in this successor list in the same (opposite, resp.) order as they appear in the *st-ordering*. Further, the successor list of u is *bitonic* if there exists an outgoing neighbor h of u , called *apex* of u , such that the successor list of u is monotonically increasing from the beginning up to h and monotonically decreasing from h up to the end. Note that a monotonically increasing (decreasing, resp.) successor list is bitonic with the rightmost (leftmost, resp.) outgoing neighbor being its apex. We call an embedding \mathcal{E} of an *st-planar* graph G *monotonic* (*bitonic*, resp.) if there exists an *st-ordering* of G such that the successor lists of all vertices defined by \mathcal{E} are monotonically increasing/decreasing (*bitonic*, resp.); we call the corresponding *st-ordering* *monotonic* (*bitonic*, resp.). Further, we say that an *st-planar* graph G is *monotonic* (*bitonic*, resp.) if G admits a *monotonic* (*bitonic*, resp.) embedding.

Forbidden configurations for bitonic st-orderings Consider an embedding and an *st-ordering* of an *st-planar* graph G . Let u be a vertex and h the outgoing neighbor of




Fig. 1 Forbidden configurations that prevent a bitonic successor list for u where **a** $v_{i+1} \neq v_j$ and **b** $v_{i+1} = v_j$, respectively

u with largest rank in the st -ordering. Note that h is the only possible apex for the successor list of u . Then, the successor list of u is not bitonic if and only if there exist two vertices v, w such that v precedes w in the st -ordering and v appears between w and h in the successor list. We call this configuration a *conflict*. It has been shown [10] that for a given embedding of an st -planar graph there exists an st -ordering without conflicts if and only if the embedding does not contain any *forbidden configuration*, where a forbidden configuration with *source* u is formed by two faces $f_1 = \langle u, v_{i+1}, \dots, v_i \rangle$ and $f_2 = \langle u, v_j, \dots, v_{j+1} \rangle$ such that the successor list of u contains $v_i, v_{i+1}, v_j, v_{j+1}$ in this order, with possibly $v_{i+1} = v_j$, and (v_{i+1}, \dots, v_i) and (v_j, \dots, v_{j+1}) are directed paths in G ; see Fig. 1. In order to obtain bitonic embeddings even in the presence of forbidden configurations, Gronemann [10] proposed to *split* at least one of the transitive edges (u, v_i) and (u, v_{j+1}) . More specifically, if we split edge (u, v_i) , we obtain two new edges (u, v'_i) and (v'_i, v_i) with dummy vertex v'_i . Note that v'_i then replaces v_i in the successor list of u in the obtained graph. Since there exists no directed path from v_{i+1} to v'_i , the forbidden configuration has been resolved.

Connectivity and SPQR-trees A graph is *connected* if for any pair of vertices there is a path connecting them. A graph is *k-connected* if the removal of any set of $k - 1$ vertices leaves it connected. A 2- or 3-connected graph is also referred to as *biconnected* or *triconnected*, respectively. Note that a triconnected planar graph has a unique embedding up to the choice of the outer face. Also note that st -planar graphs are always biconnected.

The *SPQR-tree* \mathcal{T} of an st -planar graph G is a labeled tree representing the decomposition of G into its triconnected components [11, 16]. Every triconnected component of G is associated with a node μ in \mathcal{T} . The two vertices separating the component associated with μ from the rest of the graph are called the *poles* s_μ and t_μ of μ . The *skeleton* of μ , denoted by $skel(\mu)$, is an st -planar graph where $s = s_\mu$ and $t = t_\mu$ whose edges are called *virtual edges*. In particular, there exists a virtual edge for every child ν of μ in \mathcal{T} plus a *parent virtual edge* (s_μ, t_μ) that corresponds to a virtual edge between s_μ and t_μ in the skeleton of its parent. A node $\mu \in \mathcal{T}$ can be of one of four different types:

- (i) *S-node*, if $skel(\mu)$ is composed of the parent virtual edge and a directed path of length at least 2 from s_μ to t_μ ;
- (ii) *P-node*, if $skel(\mu)$ is a bundle of at least three parallel edges from s_μ to t_μ ;

- (iii) *Q-node*, if $skel(\mu)$ consists of two parallel edges, one being the parent virtual edge and the other one being the corresponding edge in G ;
- (iv) *R-node*, if $skel(\mu)$ is a simple triconnected st -planar graph with $s = s_\mu$ and $t = t_\mu$.

The set of leaves of \mathcal{T} coincides with the set of Q-nodes, except for the Q-node ρ corresponding to edge (s, t) , which is selected as the root of \mathcal{T} . Also, neither two S-nodes, nor two P-nodes are adjacent in \mathcal{T} . The subtree \mathcal{T}_μ of \mathcal{T} rooted at μ induces a subgraph $pert(\mu)$ of G , called *pertinent*, which is described by \mathcal{T}_μ in the decomposition. In particular, $pert(\mu)$ is obtained from $skel(\mu)$ by recursively identifying each virtual edge with the corresponding parent virtual edge in the corresponding child node. We assume that the parent virtual edge of μ is not part of $pert(\mu)$. All embeddings of $pert(\mu)$ can be described by a permutation of the parallel virtual edges in each P-node in \mathcal{T}_μ and a flip of the skeleton of each R-node in \mathcal{T}_μ . SPQR-tree \mathcal{T} is unique, and can be computed in linear time [17].

3 Number of Splits in the Variable Embedding Setting

In this section, we present an algorithm that given an st -planar graph G computes a minimum-cardinality set of edges E' of G so that the graph G' obtained from G by splitting every edge in E' is bitonic. More precisely, our goal is to construct an embedding \mathcal{E} of G such that the embedding \mathcal{E}' of G' obtained from \mathcal{E} by splitting the edges in E' admits a bitonic ordering π' . At the end of the section we discuss an analogous (and simpler) algorithm that even guarantees G' to be monotonic albeit at the cost of more splits.

To compute \mathcal{E} , we adopt an SPQR-tree approach similar to the one by Chaplick et al. [8] to test whether an st -planar graph with fixed upward planar embedding is bitonic. In contrast, however, we do not explicitly augment the graph as Chaplick et al. [8] do. Instead, we specify the embedding \mathcal{E} and a labeling of the edges describing whether they will eventually be split.

Let \mathcal{T} be the SPQR-tree of G , rooted at the edge (s, t) . We associate each node μ of \mathcal{T} , with poles s_μ and t_μ , with two costs $c_b(\mu)$ and $c_m(\mu)$, and with two embeddings $\mathcal{E}_b(\mu)$ and $\mathcal{E}_m(\mu)$ of $pert(\mu)$, whose edges are labeled as `split` or `non-split`, such that the following invariants hold:

- I.1 $c_m(\mu)$ is the minimum number of splits to make $pert(\mu)$ bitonic, with the additional requirement that the successor list of s_μ is monotonically decreasing, and $\mathcal{E}_m(\mu)$ is an embedding of $pert(\mu)$ achieving this cost;
- I.2 $c_b(\mu)$ is the minimum number of splits to make $pert(\mu)$ bitonic, with no additional requirement, and $\mathcal{E}_b(\mu)$ is an embedding of $pert(\mu)$ achieving this cost;
- I.3 (a) An edge e in $\mathcal{E}_m(\mu)$ is labeled as `split` if and only if e contributes to $c_m(\mu)$,
(b) An edge e in $\mathcal{E}_b(\mu)$ is labeled as `split` if and only if e contributes to $c_b(\mu)$;
- I.4 If the edge (s_μ, t_μ) exists in $pert(\mu)$, then t_μ is the apex of s_μ in $\mathcal{E}_m(\mu)$ and the edge (s_μ, t_μ) is labeled as `non-split`.

Observe that, by definition, it holds that $c_b(\mu) \leq c_m(\mu)$.




Fig. 2 Bitonic embedding $\mathcal{E}_b(\mu)$ for a P-node μ . One child v_h uses its bitonic embedding $\mathcal{E}_b(v_h)$. Embedding $\mathcal{E}_m(v_\ell)$ of child v_ℓ appearing left of v_h is flipped

We perform a bottom-up traversal of \mathcal{T} and compute for each node μ the costs $c_b(\mu)$ and $c_m(\mu)$, the two embeddings $\mathcal{E}_b(\mu)$ and $\mathcal{E}_m(\mu)$ of $\text{pert}(\mu)$, and the labeling of their edges, so that Invariants 1.1–1.4 hold, assuming that they hold for all the children of μ . We distinguish cases based on the type of node μ .

Node μ is a Q-node that is a leaf of \mathcal{T} . We set both $\mathcal{E}_b(\mu)$ and $\mathcal{E}_m(\mu)$ to the unique embedding of $\text{pert}(\mu)$, which consists only of edge (s_μ, t_μ) . Since this embedding is monotonic, we set both costs $c_m(\mu)$ and $c_b(\mu)$ to 0, and we label (s_μ, t_μ) as non-split in both embeddings. Hence, Invariants 1.1–1.4 are satisfied.

Node μ is a P-node: Let v_1, \dots, v_k denote the children of μ . W.l.o.g., assume that if μ has a Q-node child then this child is v_1 . We construct both $\mathcal{E}_m(\mu)$ and $\mathcal{E}_b(\mu)$ by ordering the children of μ in clockwise order around s_μ from v_1 to v_k . Then, we choose embeddings and flips for the pertinent graphs of the children of μ in order to obtain $\mathcal{E}_m(\mu)$ and $\mathcal{E}_b(\mu)$, as follows.

In order to construct $\mathcal{E}_m(\mu)$, we choose the monotonic embedding $\mathcal{E}_m(v_i)$ for each child v_i and perform no flip. We set the monotonic cost $c_m(\mu)$ for μ to $\sum_{i=1}^k c_m(v_i)$, satisfying Invariant 1.1. The labeling of the edges in $\mathcal{E}_m(\mu)$ is inherited from the corresponding ones of $\mathcal{E}_m(v_1), \dots, \mathcal{E}_m(v_k)$, which ensures Invariants (a) and 1.4.

To specify $\mathcal{E}_b(\mu)$, we select one of the children of μ to contain the apex of s_μ , in such a way that the resulting bitonic cost $c_b(\mu)$ for μ is minimized. For this, we select the child v_h , with $1 \leq h \leq k$, such that the difference $c_m(v_h) - c_b(v_h)$ is maximum. If this difference is 0 for all children of μ , we set v_h to be v_1 . Then, we select the bitonic embedding $\mathcal{E}_b(v_h)$ for v_h and the monotonic embedding $\mathcal{E}_m(v_i)$ for each child $v_i \neq v_h$. Finally, we flip the embeddings $\mathcal{E}_m(v_1), \dots, \mathcal{E}_m(v_{h-1})$ of the pertinent graphs of v_1, \dots, v_{h-1} . Note that the flip of these embeddings results in a monotonically increasing successor list at s_μ for each of them, and hence guarantees that $\mathcal{E}_b(\mu)$ is bitonic, except when v_1 is a Q-node and $v_h \neq v_1$; see Fig. 2.

To guarantee that $\mathcal{E}_b(\mu)$ is bitonic also in this special case, edge (s_μ, t_μ) must be split; note that, by Invariant 1.4, edge $(s_\mu, t_\mu) = (s_{v_1}, t_{v_1})$ is labeled as non-split in the embedding $\mathcal{E}_m(v_1)$ of v_1 . So, to guarantee Invariant 1.2, we set $c_b(\mu) = c_m(\mu) - c_m(v_h) + c_b(v_h) + 1$ if v_1 is a Q-node and $v_h \neq v_1$, and $c_b(\mu) = c_m(\mu) - c_m(v_h) + c_b(v_h)$ otherwise.




Fig. 3 **a** An unavoidable conflict if v_j is selected as apex of u . **b–d** Different cases, that arise, when computing $c_b(u, j)$

To satisfy Invariant (b), we inherit the labeling of the edges in $\mathcal{E}_b(\mu)$ from the embeddings $\mathcal{E}_m(v_1), \dots, \mathcal{E}_m(v_{h-1}), \mathcal{E}_b(v_h), \mathcal{E}_m(v_{h+1}), \dots, \mathcal{E}_m(v_k)$. Further, in the special case in which v_1 is a Q-node and $v_h \neq v_1$, we label edge (s_μ, t_μ) as `split`.

Node μ is an S-node: Let v_1, \dots, v_k denote the children of μ , where $s_\mu = s_{v_1}$ and $t_\mu = t_{v_k}$. To compute $\mathcal{E}_m(\mu)$ and $\mathcal{E}_b(\mu)$, we use $\mathcal{E}_m(v_1)$ and $\mathcal{E}_b(v_1)$ for child v_1 , respectively, and the bitonic embeddings $\mathcal{E}_b(v_2), \dots, \mathcal{E}_b(v_k)$ for children v_2, \dots, v_k in both cases, without performing any flip. To guarantee Invariants I.1 and I.2, we set $c_m(\mu)$ and $c_b(\mu)$ to $c_m(v_1) + \sum_{i=2}^k c_b(v_i)$ and $\sum_{i=1}^k c_b(v_i)$, respectively. To guarantee Invariants (a) and (b), the labeling of the edges in $\mathcal{E}_b(\mu)$ and $\mathcal{E}_m(\mu)$ is inherited from the corresponding ones in the chosen embeddings of the children. Finally, Invariant I.4 is satisfied since edge (s_μ, t_μ) does not exist in $\text{pert}(\mu)$.

Node μ is an R-node: Since $\text{skel}(\mu)$ is triconnected, it has a unique embedding; we will construct both $\mathcal{E}_m(\mu)$ and $\mathcal{E}_b(\mu)$ based on such an embedding and by selecting for each child v of μ a suitable embedding of $\text{pert}(v)$ and a flip. Since each virtual edge is outgoing for only one of its end-vertices due to the fact that G is st -planar, we consider every vertex in $\text{skel}(\mu)$ independently, together with its outgoing virtual edges, similar to [10].

Let u be a vertex of $\text{skel}(\mu)$, and let $(u, v_1), \dots, (u, v_k)$ be the outgoing virtual edges of u , as they appear consecutively clockwise around u , and let v_1, \dots, v_k be the corresponding children of μ . If $u \neq s_\mu$, we can construct a bitonic successor list for u in both $\mathcal{E}_m(\mu)$ and $\mathcal{E}_b(\mu)$. Otherwise, we may need to perform different choices when constructing $\mathcal{E}_m(\mu)$ and $\mathcal{E}_b(\mu)$, to guarantee Invariants I.1 and I.2, respectively.

Suppose first that $u \neq s_\mu$. Similar to the P-node case, we determine a child v_h of μ , with $1 \leq h \leq k$, to contain the apex of u , in such a way to minimize the number of splits of the outgoing edges of u to make the successor list of u bitonic; we denote this number by $c_b(u)$. To determine v_h , we consider each child v_j , for $j = 1, \dots, k$, to be candidate for v_h , and compute the required number of splits for this choice, denoted by $c_b(u, j)$. We then obtain $c_b(u) = \min\{c_b(u, j) \mid j = 1, \dots, k\}$.

In contrast to the P-node case, we cannot conclude that $c_b(u, j) = c_b(v_j) + \sum_{i \neq j} c_m(v_i)$, since the choice of v_j and the structure of $\text{skel}(\mu)$ may result in new conflicts. Namely, consider a child v_i of μ with $i < j$ and assume that the edge (u, v_i) exists in $\text{pert}(\mu)$ and that there is a directed path in $\text{skel}(\mu)$ from v_{i+1} to v_i ; see e.g. Figure 3a. This implies that there is a directed path from a successor of u in $\text{pert}(v_{i+1})$ to v_i . Since $i < j$, this defines a forbidden configuration, and hence edge (u, v_i) must be split. On the other hand, by Invariant I.4 edge (u, v_i) is labeled as `non-split` in

$\mathcal{E}_m(v_i)$, and thus we have to account the cost of the split of edge (u, v_i) in the computation. Analogously, if $i > j$, a conflict may arise when there exists a directed path in $skel(\mu)$ from v_{i-1} to v_i . Denote by $c_s(u, j)$ the total number of these additional splits when v_j contains the apex of u . Thus,

$$c_b(u, j) = \sum_{i=1}^{j-1} c_m(v_i) + c_b(v_j) + \sum_{i=j+1}^k c_m(v_i) + c_s(u, j).$$

The computation of $c_b(u, j)$ for all $j = 1, \dots, k$ can be done in quadratic time with respect to the number of outgoing virtual edges of u . Next, we make use of ideas of the fixed-embedding algorithm [10] to achieve linear time. Namely, we first compute $c_b(u, 1)$ in linear time. Then, for each $j = 2, \dots, k$, we can compute $c_b(u, j)$ from $c_b(u, j - 1)$ in constant time as follows. Namely, for computing $c_b(u, j)$, we assume that we already computed $c_b(u, j - 1)$ and that the apex of u is to be contained in child v_j . In this transition, $pert(v_{j-1})$ changes its embedding from $\mathcal{E}_b(v_{j-1})$ to $\mathcal{E}_m(v_{j-1})$, while $pert(v_j)$ changes its embedding from $\mathcal{E}_m(v_j)$ to $\mathcal{E}_b(v_j)$; the pertinent graphs of the remaining children maintain their monotonic embeddings. We take this change into account by considering the corresponding difference $\delta_j(u) = c_b(v_j) - c_m(v_j) + c_m(v_{j-1}) - c_b(v_{j-1})$. In addition, we must also take into account the difference between $c_s(u, j - 1)$ and $c_s(u, j)$, whose computation can be done again by only considering the children v_{j-1} and v_j . More precisely, if (u, v_{j-1}) is an edge in $pert(\mu)$, then it did not need to be split when the apex of u was in v_{j-1} , but it has to be split when moving the apex to v_j , if there is a directed path from v_j to v_{j-1} ; see Fig. 3b. On the other hand, if (u, v_j) is an edge of $pert(\mu)$ and it had to be split when the apex of u was in v_{j-1} , i.e., there is a directed path from v_{j-1} to v_j , then edge (u, v_j) does not need to be split any longer when the apex is in v_j ; see Fig. 3c. Note that, if there is no directed path between v_{j-1} and v_j , then neither of the two cases occurs, and edges (u, v_{j-1}) and (u, v_j) (if they exist) do not need to be split; see Fig. 3d. In either case, we conclude that the difference between $c_s(u, j - 1)$ and $c_s(u, j)$ is at most 1. Depending on which of the three cases arises, we compute $c_b(u, j)$ as follows:

$$c_b(u, j) = c_b(u, j - 1) + \delta_j(u) + \begin{cases} 1 & \text{if } \exists \text{ directed path from } v_j \text{ to } v_{j-1}, \text{ and} \\ & (u, v_{j-1}) \text{ is an edge in } pert(\mu) \\ -1 & \text{if } \exists \text{ directed path from } v_{j-1} \text{ to } v_j, \text{ and} \\ & (u, v_j) \text{ is an edge in } pert(\mu) \\ 0 & \text{otherwise} \end{cases}$$

We remark that there is a directed path from v_j to v_{j-1} in G if and only if vertex v_{j-1} is the sink of the face shared with v_j and u in $skel(\mu)$. Similarly, a directed path from v_{j-1} to v_j in G exists if and only if vertex v_j is the sink of the face shared with v_{j-1} and u in $skel(\mu)$. Both properties can be checked in constant time as demonstrated in [10].

Once $c_b(u, j)$ has been computed for all $j = 1, \dots, k$, we choose v_h , with $1 \leq h \leq k$, so that $c_b(u, h)$ is minimum among all $c_b(u, j)$ and define $c_b(u) = c_b(u, h)$. Hence, in order to construct $\mathcal{E}_m(\mu)$ and $\mathcal{E}_b(\mu)$, we select the bitonic embedding $\mathcal{E}_b(v_h)$

for $\text{pert}(v_h)$ and the monotonic embedding $\mathcal{E}_m(v_i)$ for the pertinent graph $\text{pert}(v_i)$ for $i \in \{1, \dots, h-1, h+1, \dots, k\}$. We further flip the embeddings $\mathcal{E}_m(v_1), \dots, \mathcal{E}_m(v_{h-1})$, as in the P-node case. We inherit the labeling of the edges from the embeddings $\mathcal{E}_m(v_1), \dots, \mathcal{E}_m(v_{h-1}), \mathcal{E}_b(v_h), \mathcal{E}_m(v_{h+1}), \dots, \mathcal{E}_m(v_k)$, except for the edges that contribute to $c_s(u, h)$, which we label as `split`. We repeat the above operations for every vertex u of $\text{skel}(\mu)$ with $u \neq s_\mu$.

Consider now the case $u = s_\mu$. We distinguish two cases, based on which embedding of $\text{pert}(\mu)$ we are going to compute. Namely, for $\mathcal{E}_b(\mu)$ we perform the same operations as for any other vertex of $\text{skel}(\mu)$, since in this embedding we can have a bitonic successor list for s_μ . This guarantees Invariants I.2 and (b). In order to also guarantee Invariants I.1 and (a), we have to slightly adjust our approach. In particular, we have to obtain a monotonic successor list for s_μ in $\mathcal{E}_m(\mu)$. To achieve this, we first choose the monotonic embeddings for $\text{pert}(v_1), \dots, \text{pert}(v_k)$. Then, we have to choose whether v_1 or v_k contains the apex of s_μ . In order to perform this choice, we have to consider the conflicts that are created due to the presence of directed paths in $\text{skel}(\mu)$, as in the bitonic case. Thus, we compute $c_s(s_\mu, 1)$ and $c_s(s_\mu, k)$ and choose the minimum of the two. We choose the corresponding child to contain the apex of s_μ and label edges as `split` such that all conflicts are resolved and inherit the labeling of the remaining edges from embeddings $\mathcal{E}_m(v_1), \dots, \mathcal{E}_m(v_k)$. Note that if v_k contains the apex of s_μ , we also have to flip all the embeddings $\mathcal{E}_m(v_1), \dots, \mathcal{E}_m(v_k)$ and the resulting embedding of the entire R-node μ so to obtain a monotonically decreasing successor list for s_μ .

We account for the monotonic and the bitonic costs of node μ by summing up the corresponding costs of all vertices in V_μ , where V_μ denotes the vertex set of $\text{skel}(\mu)$. In particular, we can always choose the bitonic embedding for all the vertices different from s_μ . For s_μ on the other hand, we choose the corresponding embedding. This results in costs

$$c_b(\mu) = \sum_{u \in V_\mu} c_b(u)$$

and

$$c_m(\mu) = \sum_{\substack{u \in V_\mu \\ u \neq s_\mu}} c_b(u) + \underbrace{\sum_{i=1}^k c_m(v_i)}_{u=s_\mu} + \min\{c_s(s_\mu, 1), c_s(s_\mu, k)\}$$

for $\mathcal{E}_b(\mu)$ and $\mathcal{E}_m(\mu)$, respectively. We remark that Invariant I.4 is trivially satisfied, since edge (s_μ, t_μ) does not exist in $\text{pert}(\mu)$.

Node μ is a Q-node that is the root of \mathcal{T} : This case arises at the end of the traversal of \mathcal{T} . Since we seek to compute a bitonic embedding for G , we only have to compute $\mathcal{E}_b(\mu)$ and satisfy Invariants I.2 and (b) (i.e., Invariants I.1, (a) and I.4 can be safely neglected). Consider the unique child v of μ .

We first discuss the case where $c_b(v) < c_m(v)$. Here, assume for a contradiction, that the apex of $s_\mu = s_v$ in $\mathcal{E}_b(v)$ is incident to a face that contains both s_μ and t_μ .

Then either s_v has already a monotonic successor list in $\mathcal{E}_b(v)$ or v is a P-node and it is possible to obtain a monotonic successor list of s_v by only reordering and flipping the embeddings of the pertinent graphs of its children. But then $c_b(v) = c_m(v)$ holds, which contradicts our assumption. Hence s_μ is not incident to a face in $\mathcal{E}_b(v)$ that contains both s_μ and t_μ . Therefore, any possible embedding obtained from $\mathcal{E}_b(v)$ by adding the edge (s_μ, t_μ) violates the bitonicity of the successor list of s_μ . Thus, we label (s_μ, t_μ) as `split`, and inherit the labeling of the remaining edges from $\mathcal{E}_b(v)$.

Second, consider the case $c_b(v) = c_m(v)$. Here, we use the monotonic embedding $\mathcal{E}_m(v)$ and label the edges according to the labeling of $\mathcal{E}_m(v)$ while labeling (s_μ, t_μ) as `non-split`.

In both cases, edge (s_μ, t_μ) is embedded on the outer face of the embedding of $pert(v)$ such that t_μ is the leftmost successor of s_μ , which guarantees Invariant I.2. Invariant (b) is satisfied by the way we treat edge (s_μ, t_μ) and by inheriting the labeling of the remaining edges of $pert(\mu)$ from the chosen embedding of $pert(v)$.

We are now ready to prove the main theorem of this section.

Theorem 1 *Let $G = (V, E)$ be an st -planar graph with n vertices. There exists an $O(n)$ -time algorithm that computes an embedding and a set of edges $E' \subseteq E$ of minimum cardinality such that the graph G' obtained from G by splitting each edge in E' once is bitonic.*

Proof The correctness of our algorithm follows from the fact that at the end of the traversal of \mathcal{T} , Invariant I.2 is satisfied by the bitonic embedding $\mathcal{E}_b(\rho)$ of the root ρ of \mathcal{T} . Further, since the labeling of the edges of $\mathcal{E}_b(\rho)$ satisfies Invariant (b), we can set E' to be the set of edges that are labeled as `split` in $\mathcal{E}_b(\rho)$, which guarantees that E' is of minimum cardinality and that the graph G' obtained from G by splitting each edge in E' once is bitonic. Note that we can obtain an actual bitonic st -numbering for G' using the fixed-embedding algorithm [10] on the embedding of G' obtained from $\mathcal{E}_b(\rho)$ by splitting each edge of E' .

To complete the proof of the theorem, it remains to discuss the time complexity of our algorithm. The construction of the SPQR-tree can be done in $O(n)$ time [17]. Then, at each step of the algorithm, we consider a node μ of \mathcal{T} and we perform a set of operations in time linear to the size of $skel(\mu)$. This is clear for the Q-, S-, and P-node cases. In the R-node case, this follows from our analysis and the fact that the fixed-embedding algorithm [10] is linear in the size of the input embedding. Since the sum of the sizes of the skeletons over all the nodes of \mathcal{T} is $O(n)$ [18], the time complexity of the algorithm follows. □

We conclude this section with two remarks.

Remark 1 Our algorithm can be adjusted so that the resulting graph G' is monotonic. To achieve that, in the S- and R-node cases, we apply to all vertices of $skel(\mu)$ the same procedure as we applied to s_μ when computing $\mathcal{E}_m(\mu)$. In this way, we guarantee that the successor lists of all vertices are in fact monotonic.

Remark 2 Every series–parallel graph, oriented consistently with the series–parallel structure, is monotonic. This is because, in the absence of R-nodes, there is no need to split when computing a monotonic embedding.




Fig. 4 a Graph G_1 in \mathcal{G} . b Construction of graph G_k in \mathcal{G}

4 Bounds on the Number of Splits

In this section, we provide upper and lower bounds on the number of splits that are required for making either an n -vertex st -planar graph G or its reversed graph \tilde{G} bitonic. First, recall that there is an upper bound of $n - 3$ splits [10] and a lower bound of $\frac{3}{4}n - 3$ splits [12]. In the following, we will improve the lower bound to $n - 5$ hence showing that the upper bound is tight up to a small additive constant. Second, we investigate the number of splits that are required for graphs of bounded degree motivated by the fact that our lower bound construction has maximum degree 6.

We start by describing the family \mathcal{G} of graphs for the lower bound for general graphs.

Definition 1 (Graph family \mathcal{G}) For every integer $k \geq 1$, \mathcal{G} contains a graph $G_k = (V_k, E_k)$ that is recursively defined as follows:

- For $k = 1$, we set $V_1 = \{s_0, s_1, t_0, t_1, t_2\}$ and $E_1 = \{(s_0, t_0), (s_0, t_1), (s_0, t_2), (s_1, t_0), (s_1, t_1), (s_1, t_2), (s_1, s_0), (t_0, t_1), (t_1, t_2)\}$; see Fig. 4a.
- For $k > 1$, graph G_k constructed from G_{k-1} so that $V_k = V_{k-1} \cup \{s_k, t_{k+1}\}$ and $E_k = E_{k-1} \cup \{(s_k, t_{k-1}), (s_k, t_k), (s_k, t_{k+1}), (s_{k-1}, t_{k+1}), (s_k, s_{k-1}), (t_k, t_{k+1})\}$; see Fig. 4b.

With the following lemma, we first establish some properties of the graphs in \mathcal{G} . As a side note, we also mention that G_k has pathwidth 3.

Lemma 2 Each graph G_k in graph family \mathcal{G} contains the directed Hamiltonian path $\langle s_k, s_{k-1}, \dots, s_0, t_0, t_1, \dots, t_{k+1} \rangle$ and its underlying undirected graph is a triangulation with maximum degree 6.

Proof The fact that G_k is a triangulation follows easily by construction. Moreover, vertex t_i is connected to $t_{i+1}, t_{i-1}, s_{i-2}, s_{i-1}, s_i, s_{i+1}$ (if these vertices exist), while vertex s_i is connected to $s_{i+1}, s_{i-1}, t_{i-1}, t_i, t_{i+1}$ and t_{i+2} (again, if these vertices exist). Hence, each vertex of G_k has maximum degree 6. The existence of the Hamiltonian path can




Fig. 5 a, b Illustration of forbidden configurations and edges that must be split in both orientations

be shown inductively. Namely, G_1 contains the Hamiltonian path $\langle s_1, s_0, t_0, t_1, t_2 \rangle$ by definition. Since by induction hypothesis G_{k-1} contains the Hamiltonian path $\langle s_{k-1}, \dots, t_k \rangle$ and since G_k contains edges (s_k, s_{k-1}) and (t_k, t_{k+1}) , it holds that G_k contains the Hamiltonian path $\langle s_k, s_{k-1}, \dots, t_k, t_{k+1} \rangle$. \square

We now prove a new lower bound on the number of splits required for turning either G_k or \tilde{G}_k into a bitonic st -planar graph.

Theorem 3 *Let $G_k = (V_k, E_k)$ with $k > 1$ be a graph in graph family \mathcal{G} and let n be its number of vertices, i.e., $n = 2k + 3$. For every set $E' \subset E_k$ with $|E'| < n - 5$, neither the graph G'_k obtained from G_k by splitting each edge in E' once nor the reversed graph \tilde{G}'_k of G'_k is bitonic.*

Proof By Lemma 2, graph G_k is a triangulated Hamiltonian st -planar graph. Thus, it has a unique st -ordering and a unique embedding up to the choice of the outer face, which simply allows us to embed edge (s_k, t_{k+1}) as the leftmost or as the rightmost edge of vertex s_k . Similar arguments apply when considering \tilde{G}_k .

In the following, we describe forbidden configurations that inevitably appear in both embeddings of G_k , and then count the number of edge splits that are required to eliminate them. In particular, for vertex s_i with $1 \leq i \leq k - 1$, each of the two faces $\langle s_i, t_{i-1}, s_{i-1} \rangle$ and $\langle s_i, t_i, t_{i-1} \rangle$ form a forbidden configuration with each of the two faces $\langle s_i, s_{i-1}, t_{i+1} \rangle$ and $\langle s_i, t_{i+1}, t_{i+2} \rangle$; see the blue and red colored faces in Fig. 5a, respectively. In order to eliminate these forbidden configurations, at least one of the two pairs of edges $(s_i, t_{i-1}), (s_i, t_i)$ and $(s_i, t_{i+1}), (s_i, t_{i+2})$ must be split; see the blue and red edges in Fig. 5a, respectively. Forbidden configurations involving s_k are avoidable by embedding (s_k, t_{k+1}) as the leftmost edge of s_k .

We now describe the forbidden configurations that appear in both embeddings of \tilde{G}_k . Namely, for vertex t_i with $2 \leq i \leq k - 1$, each of the two faces $\langle t_i, t_{i-1}, s_i \rangle$ and $\langle t_i, s_i, s_{i+1} \rangle$ form a forbidden configuration with each of the two faces $\langle t_i, s_{i-2}, t_{i-1} \rangle$ and $\langle t_i, s_{i-1}, s_{i-2} \rangle$; see the blue and red colored faces in Fig. 5b, respectively. Hence, at least one of the two pairs of edges $(t_i, s_i), (t_i, s_{i+1})$ and $(t_i, s_{i-2}), (t_i, s_{i-1})$ must

be split; see the blue and red edges in Fig. 5b, respectively. Note that forbidden configurations that involve vertex t_{k+1} can be avoided by embedding edge (t_{k+1}, s_k) as the leftmost edge of t_{k+1} . Moreover, for vertex t_1 (t_k , resp.), face $\langle t_1, s_0, t_0 \rangle$ ($\langle t_k, t_{k-1}, s_k \rangle$, resp.) forms forbidden configurations with both faces $\langle t_1, t_0, s_1 \rangle$ and $\langle t_1, s_1, s_2 \rangle$ ($\langle t_k, s_{k-2}, t_{k-1} \rangle$ and $\langle t_k, s_{k-1}, s_{k-2} \rangle$, resp.). Hence, for each of t_1 and t_k at least one more incident edge must be split (namely, splitting (t_1, s_0) and (t_k, s_k) , respectively, eliminates these forbidden configurations).

We conclude that for both G_k and \tilde{G}_k , a set of edges E' of cardinality at least $2(k-1) = n-5$ has to be split to eliminate all the forbidden configurations discussed above. Note that in \tilde{G}_1 there is already one unavoidable forbidden configuration (at vertex t_1), while this is not the case for G_1 . \square

4.1 Bounds for Graphs of Bounded Degree

In Theorem 3, we showed that $n - O(1)$ is a tight upper bound for the required number of splits even for graphs with maximum degree 6 (see also Lemma 2). On the other hand, if the input graph has maximum degree 3, then no split is required [13]. In the next two theorems, we focus on graphs with maximum degree 5. For this graph class, we improve the general upper bound of $n-3$ [10] on the number of splits to $n/2$. Then, we provide a corresponding lower bound of $n/2-2$, which holds even for graphs with maximum degree 4.

Theorem 4 *Let $G = (V, E)$ be an st -planar graph with n vertices and maximum degree 5. There exists a set of edges $E' \subseteq E$ with $|E'| \leq n/2$ such that either the graph G' obtained by splitting each edge in E' or the reversed graph \tilde{G}' of G' is bitonic. Moreover, set E' can be computed in $O(n)$ time.*

Proof Consider any upward planar embedding of G . Consider a vertex v of G that is the source of at least one forbidden configuration. In the following, we will prove the existence of a face that belongs to all forbidden configurations with source v (after possibly redrawing edge (s, t)). Suppose first that v does not coincide with s , which implies that v has at most four outgoing edges. Hence, there exist at most three internal faces with v as a source, which implies that there is one that belongs to all forbidden configurations with source v . Suppose now that v coincides with s . In this case, s can be the source of four internal faces and, therefore, there might exist two forbidden configurations with source s not sharing a face. If s is the source of two such forbidden configurations and if (s, t) is the leftmost (rightmost) outgoing edge of s , then we redraw it as rightmost (leftmost, respectively). After this redrawing, s is necessarily the source of three forbidden configurations that share a single face f , i.e., one of the previous two forbidden configurations is resolved. Now, all remaining forbidden configurations with source s share the face f that has a transitive edge which is not (s, t) .

In both cases, there is a face f that belongs to all forbidden configurations with source v . Therefore, by splitting the transitive edge of f , every forbidden configuration with source v is resolved. We associate this split with vertex v . Since v is the source of at least one forbidden configuration, it has at least three outgoing edges, and thus at most

two incoming edges. Therefore, v cannot be the source of a forbidden configuration in \tilde{G} as well. This allows us to conclude that each vertex contributes at most one split in either G or \tilde{G} .

For the time complexity, we first observe that an upward-planar embedding can be computed in linear time for st -planar graphs [1]. Then, for each face we decide whether it contains a transitive edge or not. If this is the case, we assign the corresponding face to its source and sink. Afterwards, we iterate over all vertices and check how many faces are assigned to it (both as a source and as a sink) and perform the splits as described above. This last step can be done in constant time for each vertex as at most four faces are assigned to each vertex in the previous step. \square

Next, we prove that the upper bound given in Theorem 4 is tight by providing a corresponding lower bound. First, we present the class of graphs used in our lower bound construction. For this purpose, we define the following family of auxiliary graphs:

Definition 2 (Auxiliary graph family $\hat{\mathcal{H}}$) For every integer $k \geq 1$, $\hat{\mathcal{H}}$ contains a graph $\hat{H}_k = (V_k, E_k)$ that is recursively defined as follows:

- For $k = 1$, graph \hat{H}_1 consists of six vertices that form a directed path $\langle b_{1,b}, r_{1,b}, g_{1,b}, b_{1,t}, g_{1,t}, r_{1,t} \rangle$. Additionally, it contains four edges $(b_{1,b}, b_{1,t})$, $(r_{1,b}, b_{1,t})$, $(r_{1,b}, r_{1,t})$ and $(g_{1,b}, g_{1,t})$; see Fig. 6a.
- For $k > 1$, graph \hat{H}_k contains \hat{H}_{k-1} as a subgraph and six additional vertices $b_{k,t}, g_{k,t}, r_{k,t}, b_{k,b}, r_{k,b}, g_{k,b}$ that form two directed paths $\langle r_{k-1,t}, b_{k,t}, g_{k,t}, r_{k,t} \rangle$ and $\langle b_{k,b}, r_{k,b}, g_{k,b}, b_{k-1,b} \rangle$ with $r_{k-1,t}$ and $b_{k-1,b}$ of \hat{H}_{k-1} , respectively. The recursive construction of \hat{H}_k is completed by adding the six edges $(b_{i-1,b}, b_{i,t})$, $(b_{i,b}, b_{i,t})$, $(g_{i,b}, g_{i-1,t})$, $(g_{i,b}, g_{i,t})$, $(r_{i,b}, r_{i-1,t})$ and $(r_{i,b}, r_{i,t})$; see Fig. 6b.

Based on $\hat{\mathcal{H}}$, we define the following graph family:

Definition 3 (Graph family \mathcal{H}) For every integer $k \geq 1$, \mathcal{H} contains a graph $H_k = (V_k, E_k)$ obtained from \hat{H}_k of $\hat{\mathcal{H}}$ by adding the two edges $(b_{k,b}, g_{k,t})$ and $(b_{k,b}, r_{k,t})$ to it; see the dashed edges in Fig. 6b.

In the following, we discuss some properties of each graph H_k of \mathcal{H} .

Observation 1 Let $H_k \in \mathcal{H}$. The edges of H_k can be partitioned into the following four paths (ignoring the edge orientations):

- $P = \langle b_{k,b}, r_{k,b}, g_{k,b}, \dots, b_{1,b}, r_{1,b}, g_{1,b}, b_{1,t}, g_{1,t}, r_{1,t}, \dots, b_{k,t}, g_{k,t}, r_{k,t} \rangle$,
- $P_b = \langle b_{1,t}, b_{1,b}, b_{2,t}, b_{2,b}, \dots, b_{k,t}, b_{k,b} \rangle$,
- $P_g = \langle g_{1,b}, g_{1,t}, g_{2,b}, g_{2,t}, \dots, g_{k,b}, g_{k,t}, b_{k,b} \rangle$, and
- $P_r = \langle b_{1,t}, r_{1,b}, r_{1,t}, r_{2,b}, r_{2,t}, \dots, r_{k,b}, r_{k,t}, b_{k,b} \rangle$.

Moreover, P is a directed Hamiltonian path.

Lemma 5 The underlying undirected graph of each $H_k \in \mathcal{H}$ is triconnected and has maximum degree 4.

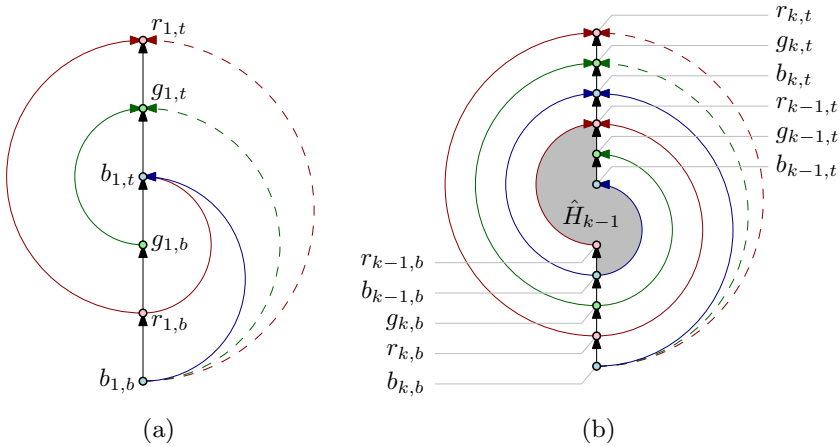


Fig. 6 **a** Graphs \hat{H}_1 in $\hat{\mathcal{H}}$ and H_1 in \mathcal{H} . **b** Construction of graph \hat{H}_k in $\hat{\mathcal{H}}$ and H_k in \mathcal{H} . In both subfigures, dashed edges are part of H_k but not of \hat{H}_k

Proof By Observation 1, the edges of H_k can be partitioned into the four paths P , P_b , P_g and P_r (ignoring the edge orientations). Each vertex $b_{i,j}$ of H_k is on path P_b is also connected to one vertex $r_{k,\ell}$ via P or P_r and to one vertex $g_{x,y}$ via P or P_g for some values of i, j, k, ℓ, x and y . Analogous properties hold for vertices $g_{i,j}$ and $r_{i,j}$ of H_k . Thus, for each pair of vertices, there exists three disjoint paths such that each follows a distinct subpath of P_b , P_g and P_r possibly except for the very first and the very last edge which may also belong to P .

For the maximum degree, observe that every vertex except for $b_{1,t}$ and $b_{k,b}$ belongs to exactly two of the four paths partitioning the edges of H_k . Vertex $b_{1,t}$ is an internal vertex of P and an extremal vertex of both P_b and P_r . Finally, $b_{k,b}$ is an extremal vertex in all four paths. □

We are now ready to prove that the upper bound given in Theorem 4 is tight up to a small additive constant.

Theorem 6 *Let $H_k = (V_k, E_k)$ be a graph in graph family \mathcal{H} , which has maximum degree 4, and let n be its number of vertices, i.e., $n = 6k$. For every set $E' \subset E_k$ with $|E'| < n/2 - 2$, neither the graph H'_k obtained from H_k by splitting each edge in E' once nor the reversed graph of H'_k is bitonic.*

Proof Since H_k is triconnected and Hamiltonian by Lemma 5 and Observation 1, it admits a unique upward planar embedding (see Fig. 6a, b) up to a flip and the choice of the outer face which can either be $\langle b_{k,b}, r_{k,b}, r_{k,t} \rangle$ as shown in Fig. 6b or $\langle r_{k,t}, g_{k,t}, b_{k,b} \rangle$ depending on whether $\langle b_{k,b}, r_{k,t} \rangle$ is drawn on the right or on the left side, respectively. We emphasize that the forbidden configurations that we describe next occur in each of the four possible upward planar embeddings.

First consider H_k . We find the following forbidden configurations:

- For $2 \leq i \leq k$, vertex $r_{i,b}$ is source of the forbidden configuration formed by the faces $\langle r_{i,b}, r_{i,t}, g_{i,t}, g_{i,b} \rangle$ and $\langle r_{i,b}, g_{i,b}, g_{i-1,t}, r_{i-1,t} \rangle$; see Fig. 7a.

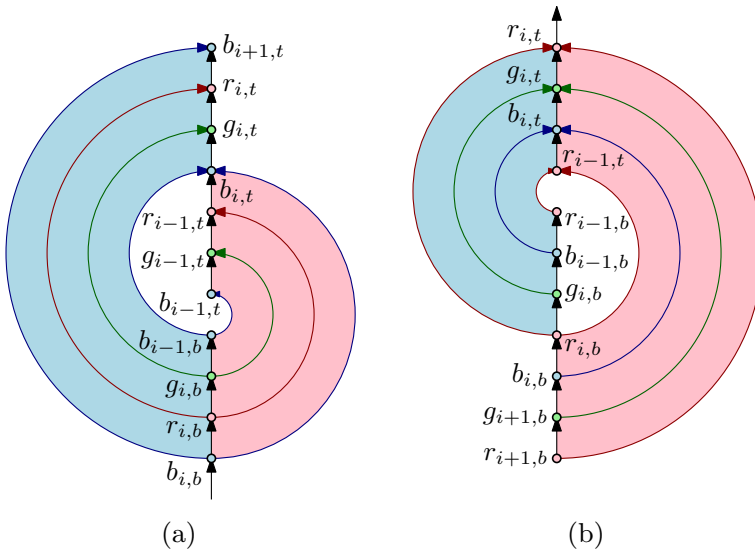


Fig. 7 a, b Illustration of forbidden configurations and edges that must be split in both orientations

- For $2 \leq i \leq k$, vertex $g_{i,b}$ is source of the forbidden configuration formed by the faces $\langle g_{i,b}, g_{i,t}, b_{i,t}, b_{i-1,b} \rangle$ and $\langle g_{i,b}, b_{i-1,b}, b_{i-1,t}, g_{i-1,t} \rangle$; see Fig. 7a.
- For $2 \leq i \leq k - 1$, vertex $b_{i,b}$ is source of the forbidden configuration formed by the faces $\langle b_{i,b}, b_{i+1,t}, r_{i,t}, r_{i,b} \rangle$ and $\langle b_{i,b}, r_{i,b}, r_{i-1,t}, b_{i,t} \rangle$; see Fig. 7a.
- Vertex $r_{1,b}$ is source of a forbidden configuration formed by faces $\langle r_{1,b}, r_{1,t}, g_{1,t}, g_{1,b} \rangle$ and $\langle r_{1,b}, g_{1,b}, b_{1,t} \rangle$; see Fig. 6a.
- Vertex $b_{1,b}$ is source of a forbidden configuration formed by faces $\langle b_{1,b}, b_{2,t}, r_{1,t}, r_{1,b} \rangle$ and $\langle b_{1,b}, r_{1,b}, b_{1,t} \rangle$; see Fig. 6a.

Thus, there are $3k - 2$ forbidden configurations in any upward planar embedding of H_k that each require at least one split to obtain a bitonic subdivision.

Next consider the reversed graph \tilde{H}_k . In each of the four upward planar embeddings of graph \tilde{H}_k , we find the following forbidden configurations:

- For $1 \leq i \leq k - 1$, vertex $r_{i,t}$ is source of the forbidden configuration formed by the faces $\langle r_{i,t}, g_{i,t}, g_{i,b}, r_{i,b} \rangle$ and $\langle r_{i,t}, r_{i+1,b}, g_{i+1,b}, g_{i,t} \rangle$; see Fig. 7b.
- For $2 \leq i \leq k - 1$, vertex $g_{i,t}$ is source of the forbidden configuration formed by the faces $\langle g_{i,t}, b_{i,t}, b_{i-1,b}, g_{i,b} \rangle$ and $\langle g_{i,t}, g_{i+1,b}, b_{i,b}, b_{i,t} \rangle$; see Fig. 7b.
- For $2 \leq i \leq k$, vertex $b_{i,t}$ is source of the forbidden configuration formed by the faces $\langle b_{i,t}, r_{i-1,t}, r_{i-1,b}, b_{i-1,b} \rangle$ and $\langle b_{i,t}, b_{i,b}, r_{i,b}, r_{i-1,t} \rangle$; see Fig. 7b.
- Vertex $g_{1,t}$ is source of a forbidden configuration formed by faces $\langle g_{1,t}, b_{1,t}, g_{1,b} \rangle$ and $\langle g_{1,t}, g_{2,b}, b_{1,b}, b_{1,t} \rangle$; see Fig. 6a.
- Vertex $g_{k,t}$ is source of a forbidden configuration formed by faces $\langle g_{k,t}, b_{k,t}, b_{k-1,b}, g_{k,b} \rangle$ and $\langle g_{k,t}, b_{k,b}, b_{k,t} \rangle$; see Fig. 6b.

We conclude that there are also $3k - 2$ forbidden configurations in any upward planar embedding of \tilde{H}_k that each require at least one split to obtain a bitonic subdivision.

Since H_k has $6k$ vertices, $3k - 2 = n/2 - 2$ and the statement of the theorem follows. \square

5 Relationship to Upward Planarity

For an n -vertex st -planar graph that can be made bitonic by splitting a subset E' of its edges, Gronemann suggested an approach in [10] to construct an upward planar polyline drawing of it in $O(n^2)$ area, where each edge of E' has one bend while the remaining edges are drawn as straight lines. This approach has been used to prove that every n -vertex st -planar graph admits an upward planar drawing in $O(n^2)$ area with at most one bend per edge and at most $n - 3$ bends in total [10], and that every n -vertex st -planar graph of maximum degree 3 can be drawn upward planar without any bends in $O(n^2)$ area [13]. A notable consequence of Theorem 4 is that the bound on the total number of bends can be reduced from $n - 3$ to $n/2$ for n -vertex st -planar graphs of maximum degree 5.

Theorems 3 and 6, on the other hand, show the limitations of this approach. In particular, when constructing upward planar drawings with the approach by Gronemann [10], $n - 5$ bends may be required for graphs of maximum degree 6 (by Theorem 3), and $n/2 - 2$ bends may be required for graphs of maximum degree 4 (by Theorem 6). Note that these limitations are tailored to the adopted approach. As a matter of fact, Gronemann [7] further observed that there exist st -planar graphs which require a linear number of splits to become bitonic, and at the same time admit bend-less upward planar drawings in even linear area.

In this section, we investigate whether these limitations are caused by the specific drawing technique or are already imposed by the nature of the upward planarity problem. To this end, we study lower bounds on the total number of bends in upward planar drawings under the polynomial-area requirement, independently of the required number of splits and of the allowed number of bends per edge. Interestingly, our findings imply that the upper bounds on the number of bends obtained by the approach by Gronemann are worst-case almost tight, even if more than one bend per edge is allowed.

Central in our lower bound studies is the following structure, the so-called coil; for an illustration refer to Fig. 8a:

Definition 4 (k-coil) A k -coil $\xi = \langle v_0, v_1, \dots, v_k \rangle$ with $k \geq 2$ in an upward planar embedding of an st -planar graph G is an embedded subgraph of G so that

- (i) either $(v_i, v_{i-1}), (v_i, v_{i+1}) \in E(\xi)$, i.e., v_i is source in ξ , or $(v_{i-1}, v_i), (v_{i+1}, v_i) \in E(\xi)$, i.e., v_i is sink in ξ , for $1 \leq i \leq k - 1$,
- (ii) there is a directed st -path P_ξ in G that passes through all vertices of ξ so that v_i follows v_{i+2} along P_ξ if v_i is sink in ξ or precedes v_{i+2} along P_ξ if v_i is source in ξ , and,
- (iii) for $1 \leq i \leq k - 1$, the edges $\{v_i, v_{i-1}\}, \{v_i, n_i^*\}, \{v_i, v_{i+1}\}$ appear in this order either consistently clockwise or consistently counter-clockwise around v_i where n_i^* is the predecessor of v_i on P_ξ if v_i is source in ξ or the successor of v_i on P_ξ if v_i is sink in ξ .




Fig. 8 **a** A 6-coil $\langle v_0, \dots, v_6 \rangle$. **b** The region $R(\xi_i)$

Note that by Property [(i)] of Definition 4 it follows that if v_i is sink in ξ , then v_{i+1} is source in ξ , and vice-versa. We also remark that a similar concept has already been used for area lower bounds in [19] and [20]. In particular, Frati [20] proved that a k -coil requires $\Omega(2^k)$ area in any upward planar straight-line drawing. In the following, we generalize this result by showing that superpolynomial area is required for a coil unless roughly half of its edges have at least one bend each.

Lemma 7 *Let G be an st -planar graph with a fixed upward planar embedding containing a k -coil $\xi = \langle v_0, v_1, \dots, v_k \rangle$. In any polyline upward planar drawing of G , ξ is drawn in $\omega(\text{poly}(k)) = \omega(2^{\log(k)})$ area unless $k/2 - O(\log(k))$ edges of ξ have at least one bend.*

Proof Assume for a contradiction that G admits an upward st -planar drawing Γ in which the k -coil ξ is drawn in $O(\text{poly}(k))$ area such that $k/2 - \omega(\log(k))$ edges of ξ are bent.

Consider a 2-coil $\xi_i = \langle v_{i-1}, v_i, v_{i+1} \rangle$ that is part of ξ . We say that ξ_i is a *V-shape* if and only if v_i is a source of ξ . Similarly, we say that ξ_i is a *Λ -shape* if and only if v_i is a sink of ξ . By definition of k -coil, the number of 2-coils in ξ is $k - 1$; more precisely, ξ_1, \dots, ξ_{k-1} are the 2-coils in ξ , where ξ_i is a V-shape if and only if ξ_{i+1} is a Λ -shape. Furthermore, we say that ξ_i is *eliminated* in Γ if either edge (v_i, v_{i-1}) or edge (v_i, v_{i+1}) is bent in Γ . Conversely, we say that a bend b on one of these two edges *eliminates* ξ_i . We call a V-shape ξ_i *valid* if and only if ξ_i is not eliminated and the next non-eliminated 2-coil ξ_j along ξ with $j > i$ is a Λ -shape. Similarly, we call a Λ -shape ξ_i *valid* if and only if it is not eliminated and the next non-eliminated 2-coil ξ_j along ξ with $j > i$ is a V-shape.




Fig. 9 Illustration for the proof of Lemma 7

Let $\xi' = \langle v_j, v_{j+1}, \dots, v_{j+k'} \rangle$ be a k' -subcoil of ξ , so that all k' edges of ξ' are bent in Γ . Then, the bends along ξ' eliminate all 2-coils $\xi_j, \dots, \xi_{j+k'}$ of ξ' , which are in total $k' + 1$. In addition, if k' is even, ξ_{j-1} and $\xi_{j+k'+1}$ are either both V-shapes or both Λ -shapes. Thus, the bends along ξ' make $k' + 1$ of the 2-coils in ξ not valid, if k' is odd, or $k' + 2$ of the 2-coils in ξ , if k' is even. In either case, the bends along ξ' make at most $2k'$ of the 2-coils in ξ not valid, namely if $k' \in \{1, 2\}$. Recall now that in Γ only $k/2 - \omega(\log(k))$ edges of ξ are bent. Hence, by the previous analysis, at most $k - 2\omega(\log(k))$ 2-coils in ξ are not valid. As ξ contains $k - 1$ 2-coils, we conclude that ξ contains $c = \omega(\log(k))$ valid 2-coils ξ_1^*, \dots, ξ_c^* so that if $\xi_i^* = \xi_j$ and $\xi_{i+1}^* = \xi_{j'}$ it holds that $j < j'$, for $1 \leq i < c$. In particular, if ξ_i^* is a Λ -shape, then ξ_{i+1}^* is a V-shape and vice versa.

We first observe, that all of v_{i+2}, \dots, v_k are located in the bounded region $R(\xi_i)$ delimited by edge $\{v_i, v_{i+1}\}$, the horizontal h_{i+1} through v_{i+1} and the part of edge $\{v_i, v_{i-1}\}$ between v_i and the crossing with h_{i+1} ; see Fig. 8b (in this and the following figures $\{v_i, v_{i-1}\}, \{v_i, v_i^*\}, \{v_i, v_{i+1}\}$ appear in this order clockwise around v_i ; the counter-clockwise case is symmetric). We now show, that the region $R(\xi_i^*)$ has at least four times as much area as $R(\xi_{i+1}^*)$ for all $1 \leq i < c$. We assume w.l.o.g. that $\xi_i^* = \langle \beta_1, \beta_2, \beta_3 \rangle$ is a V-shape while $\xi_{i-1}^* = \langle \lambda_1, \lambda_2, \lambda_3 \rangle$ is a Λ -shape. Further, we assume w.l.o.g. that β_2 and λ_2 have the same x -coordinate, otherwise this property may be obtained by shearing the drawing horizontally. Note that Euclidean area is invariant under shear mapping [21][Thm.9-2].

Consider the following:

- horizontals y_2 and y_3 through vertices λ_2 and λ_3 , resp.,
- rays r_1 and r_3 from λ_2 through λ_1 and λ_3 , resp.,
- point λ'_1 on the intersection of y_3 and r_1 ,
- verticals x_1, x_2 and x_3 through λ'_1, λ_2 and λ_3 , resp.,
- the line p_1 through λ_3 which is parallel to edge (β_2, β_1) , and,
- the line p_3 through λ'_1 which is parallel to edge (β_2, β_3) .

We consider the following three cases:

Case 1: Both r_1 and r_3 cross (β_2, β_1) . For an illustration refer to Fig. 9a. The area of $R(\xi_{i+1}^*)$ is at most half the area of the axis-aligned rectangle R spanned by edge (λ_3, λ_2) . Moreover, $R(\xi_i^*)$ contains the right triangle T bounded by y_2, x_2 and p_1 . Since R is an inscribed rectangle of T , R has at most half the area of T . Thus, we conclude that the area of $R(\xi_i^*)$ is at least four times as large as the area of $R(\xi_{i+1}^*)$.

Case 2: Both r_1 and r_3 cross (β_2, β_3) . This case is symmetric to Case 1; here R is spanned by the segment (λ_2, λ'_1) .

Case 3: r_1 crosses (β_2, β_1) and r_3 crosses (β_2, β_3) . For an illustration refer to Fig. 9b. The area of $R(\xi_{i+1}^*)$ is at most half the area of the axis-aligned rectangle R bounded by x_1, x_3, y_2 and y_3 . Moreover, $R(\xi_i^*)$ contains the triangle T bounded by y_2, p_1 and p_3 . Since R is an inscribed rectangle of T , R has at most half the area of T . Thus, we conclude that the area of $R(\xi_i^*)$ is at least four times as large as the area of $R(\xi_{i+1}^*)$.

Recall that by assumption, $k/2 - \omega(\log(k))$ edges of ξ are bent. As we have proven, there is the sequence ξ_1^*, \dots, ξ_c^* of 2-coils with $c = \omega(\log(k))$ as claimed above. Finally, we showed that the area of $R(\xi_i)$ is at least four times as large as the area of $R(\xi_{i+1}^*)$ for $1 \leq i < c$.

We now conclude that $R(\xi_1^*)$ has at least $4^{c-2} \cdot \Omega(1) = 4^{\omega(\log(k))}$ area, which is superpolynomial in k assuming that $R(\xi_{c-1}^*)$ has area $\Omega(1)$. This leads to a contradiction.

Hence, it remains to show that $R(\xi_{c-1}^*)$ has area $\Omega(1)$. Assume w.l.o.g. that $\lambda_1, \lambda_2, \lambda_3$ are the vertices of ξ_c^* and that $\beta_1, \beta_2, \beta_3$ are the vertices of ξ_{c-1}^* ; as shown in Fig. 9. Observe that $R(\xi_{c-1}^*)$ contains a triangle T bounded by β_2, β_3 , and λ_2 . By Pick’s theorem [22, 23], T , and hence also $R(\xi_{c-1}^*)$, has area at least $1/2$. This concludes the proof. □

We now shift our attention back to the two graph families introduced in Sect. 4. As we shall see in the next two theorems, members of these graph families contain coils which require many bends in any polynomial area upward planar drawing. More precisely, the induced number of required bends almost matches the upper bound obtained via splitting technique for obtaining a bitonic subdivision:

Theorem 8 *Let $G_k \in \mathcal{G}$ and let n be its number of vertices. Graph G_k does not admit an upward planar drawing with $n - o(\log(n))$ bends within polynomial area.*

Proof Recall the definition of the graph family \mathcal{G} described in Definition 1. In particular, we may assume $G_k \in \mathcal{G}$ to be upward planar embedded as it is a triangulated graph with a Hamiltonian directed path. Observe that in G_k

- the path $\xi_h = \langle s_k, t_k, s_{k-1}, t_{k-1}, \dots, s_1, t_1, s_0 \rangle$ is a $(2k - 1)$ -coil, and
- $\xi_1 = \langle s_k, t_{k-1}, s_{k-3}, t_{k-4}, s_{k-6}, \dots \rangle$, $\xi_2 = \langle s_{k-1}, t_{k-2}, s_{k-4}, t_{k-5}, s_{k-7}, \dots \rangle$ and $\xi_3 = \langle s_{k-2}, t_{k-3}, s_{k-5}, t_{k-6}, s_{k-8}, \dots \rangle$ are three $(2k/3 - O(1))$ -coils.

By Lemma 7, every polynomial area upward drawing Γ of G_k has at least $\frac{2k}{2} - o(\log(2k)) = k - o(\log(k))$ bent edges along ξ_h and at least $\frac{2k/3}{2} - o(\log(2k/3)) = \frac{k}{3} - o(\log(k))$ bent edges along each of ξ_1, ξ_2 and ξ_3 . Since ξ_h, ξ_1, ξ_2 and ξ_3 are pairwise edge-disjoint, Γ has at least $2k - o(\log(k)) = n - o(\log(n))$ bent edges in total. □

Theorem 9 *For infinitely many values of $n \in \mathbb{N}$, there exists an n -vertex st -planar graph with maximum degree 4 that does not admit any upward planar drawing with $\frac{n}{2} - o(\log(n))$ bends within polynomial area.*

Proof Consider the graph H_k from the proof of Theorem 6. Observe that the paths P_r , P_b and P_g are disjoint $(2k + 1)$ -, $(2k - 1)$, and $2k$ -coils, resp. By Lemma 7, every polynomial area upward drawing Γ of H_k has at least $\frac{2k}{2} - o(\log(2k)) = k - o(\log(k))$ bent edges along each of P_r , P_b and P_g . Thus, Γ has at least $3k - o(\log(k)) = \frac{n}{2} - o(\log(n))$ bent edges in total. \square

Note that Theorem 9 improves upon the result of Di Battista et al. [19] in two ways: First, we show that the superpolynomial area is still required if we allow an almost linear number of bends. Second, our area lower bound construction has maximum degree 4 which closes the gap towards the maximum degree 3 graphs, which are known to always admit a straight-line upward drawing in polynomial area.

This result also strengthens the observation on the relationship between the number of splits in bitonic st -orderings and the number of bends in polynomial-area upward planar drawings. Consequently, we may now ask whether graphs that require a certain number of splits in any bitonic st -ordering also require a certain number of bends in a polynomial-area upward planar drawing. In the following theorem, we answer this question negatively even for graphs of maximum degree 4.

Theorem 10 *There exist infinitely many n -vertex st -planar graphs $G = (V, E)$ with maximum degree 4 so that*

- (i) *for every set $E' \subset E$ with $|E'| < n/4 - 5/2$, neither the graph G' obtained from G by splitting each edge in E' once nor the reversed graph of G' is bitonic and*
- (ii) *there is a straight-line upward planar drawing of G within quadratic area.*

Proof Our construction for the n -vertex (such that $n \equiv 2(\text{mod } 4)$) st -planar graph G consists of source s and sink t and subgraphs G_1 and G_2 with sources s_1 and s_2 , respectively, and sinks t_1 and t_2 , respectively, which are st -planar except for lacking the st -edge; refer also Fig. 10. We first describe G_1 .

Let $k = n/4 - 3/2$. Then, G_1 contains the three directed paths

- $\langle s_1 = v_1, v_2, \dots, v_k \rangle$,
- $\langle v_k, \ell_{(k-1)/2}, \dots, \ell_2, \ell_1 \rangle$ and
- $\langle v_k, r_{(k-1)/2}, \dots, r_2, r_1 \rangle$.

In addition, there is a vertex t_1 and the following edges:

- (ℓ_1, t_1) ,
- (r_1, t_1) ,
- $(v_i, \ell_{\lceil i/2 \rceil})$ for $i \in \{1, \dots, k - 1\}$ and
- $(v_i, \ell_{\lfloor i/2 \rfloor})$ for $i \in \{1, \dots, k - 1\}$.

Now, we observe that vertex v_i for $i \in \{1, \dots, k - 1\}$ is source of a forbidden configuration, namely,




Fig. 10 Illustration for the proof of Theorem 10

- if i is odd, consisting of face $(v_i, \ell_{(i+1)/2}, v_{i+1})$ with transitive edge $(v_i, \ell_{(i+1)/2})$ and face $(v_i, r_{(i+1)/2}, v_{i+1})$ with transitive edge $(v_i, r_{(i+1)/2})$, or,
- if i is even, consisting of face $(v_i, \ell_{i/2}, \ell_{i/2+1}, v_{i+1})$ with transitive edge $(v_i, \ell_{i/2})$ and face $(v_i, r_{i/2}, r_{i/2+1}, v_{i+1})$ with transitive edge $(v_i, r_{i/2})$.

Hence, G_1 requires at least $k - 1 = n/4 - 5/2$ splits in G . In addition, G_1 admits an upward planar drawing in quadratic area (as shown in Fig. 10), where

- v_i is located at $(0, i)$ for $i \in \{1, \dots, k\}$,
- ℓ_i is located at $(-(k + 1)/2 + i, k + i)$ for $i \in \{1, \dots, (k - 1)/2\}$,
- r_i is located at $((k + 1)/2 - i, k + i)$ for $i \in \{1, \dots, (k - 1)/2\}$ and
- t_1 is located at $(0, 3k/2 + 1/2)$.

Now the construction of G is completed as follows. G_2 is isomorphic to \tilde{G}_1 . Further, vertices s and t are incident to edges (s, t) , (s, s_1) , (s, s_2) , (t_1, t) and (t_2, t) .

Since G_2 is isomorphic to \tilde{G}_1 , it requires at least $k - 1 = n/4 - 5/2$ splits in \tilde{G} and admits an upward planar drawing isomorphic to the one described above for G_1 . We can combine the described drawings of G_1 and G_2 to a drawing of G as follows:

- G_1 is drawn as described above,
- vertex s is placed at $((k + 1)/2, 0)$, and,
- the drawing of G_2 is isomorphic (rotated by $\pi/2$) to the drawing of G_1 and vertex s_2 of G_2 (which is isomorphic to t_1) is located at $(k + 1, 1)$.

With the above construction, we obtain a drawing in area $O(k) \times O(k) = \mathcal{O}(n^2)$. In addition, we have already seen that G_1 requires at least $n/4 - 5/2$ splits in G , while G_2 requires at least $n/4 - 5/2$ splits in \tilde{G} . This concludes the proof. \square




Fig. 11 Illustration of a slight modification of a lower bound construction in [12]

Similarly to the previous theorem, the graph shown in Fig. 11 that is inspired by a lower bound construction of Rettner [12] requires only linear area for an upward planar drawing but $n/2$ splits in each orientation. Thus, we conclude that the number of bends required in a polynomial area upward planar drawing is not an upper bound for the number of splits in a bitonic st -ordering while the reverse relation holds.

6 Conclusions and Open Problems

In this work, we proposed a linear-time algorithm to minimize the number of splits over all embeddings to make a given st -planar graph bitonic. We then provided bounds on the number of required splits that are tight up to an additive constant in the worst case. Finally, we studied the relationship between the required number of such splits and the number of bends in polynomial-area upward planar drawings.

We conclude with some open problems raised by our work.

- (i) An experimental evaluation of our algorithm would allow to estimate the required number of splits in practice. In addition, our investigation of the relationship to upward planar drawings in Sect. 5 suggests that an actual implementation of our algorithm would be a viable tool in practical applications.
- (ii) In view of Remark 2 and Theorem 4, it is worth investigating other meaningful subclasses of st -planar graphs that admit improved upper bounds on the required number of splits; note that our lower bound examples in Theorems 3 and 6 already impose strong restrictions. Namely, they are Hamiltonian have pathwidth 3 and 4, respectively, and the latter construction has maximum degree 4.
- (iii) Another possible direction would be to study whether some of the results for st -planar graphs translate to general upward planar graphs. Note that the definition of bitonic st -orderings is based on st -planar graphs, hence, it would have to be extended to support general upward planar graphs. Recall that a directed

acyclic graph is upward planar if and only if it can be augmented to an st -planar graph [24].

Acknowledgements The authors would like to thank Michael Kaufmann and Antonios Symvonis for fruitful discussions.

Funding Open Access funding enabled and organized by Projekt DEAL. No funding was received for conducting this study.

Declarations

Conflict of interest 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. de Fraysseix, H., Pach, J., Pollack, R.: How to draw a planar graph on a grid. *Combinatorica* **10**(1), 41–51 (1990). <https://doi.org/10.1007/BF02122694>
2. Kant, G.: Drawing planar graphs using the canonical ordering. *Algorithmica* **16**(1), 4–32 (1996). <https://doi.org/10.1007/BF02086606>
3. Harel, D., Sardas, M.: An algorithm for straight-line drawing of planar graphs. *Algorithmica* **20**(2), 119–135 (1998). <https://doi.org/10.1007/PL00009189>
4. Biedl, T.C., Kant, G.: A better heuristic for orthogonal graph drawings. *Comput Geom.* **9**(3), 159–180 (1998). [https://doi.org/10.1016/S0925-7721\(97\)00026-6](https://doi.org/10.1016/S0925-7721(97)00026-6)
5. Rosenstiehl, P., Tarjan, R.E.: Rectilinear planar layouts and bipolar orientations of planar graphs. *Discrete Comput. Geom.* **1**, 343–353 (1986). <https://doi.org/10.1007/BF02187706>
6. Even, S., Tarjan, R.E.: Computing an st -Numbering. *Theor. Comput. Sci.* **2**(3), 339–344 (1976). [https://doi.org/10.1016/0304-3975\(76\)90086-4](https://doi.org/10.1016/0304-3975(76)90086-4)
7. Gronemann, M.: Bitonic st -orderings of biconnected planar graphs. In: Duncan, C.A., Symvonis, A. (eds.) *Graph Drawing*, vol. 8871 of LNCS, pp. 162–173. Springer, Cham (2014)
8. Chaplick, S., Chimani, M., Cornelsen, S., Da Lozzo, G., Nöllenburg, M., Patrignani, M., et al.: Planar L -drawings of directed graphs. In: Frati, F., Ma, K. (eds.) *Graph Drawing and Network Visualization*, vol. 10692 of LNCS, pp. 465–478. Springer, Cham (2017)
9. Angelini, P., Bekos, M.A., Liotta, G., Montecchiani, F.: Universal slope sets for 1-bend planar drawings. *Algorithmica* **81**(6), 2527–2556 (2019). <https://doi.org/10.1007/s00453-018-00542-9>
10. Gronemann, M.: Bitonic st -orderings for upward planar graphs. In: Hu, Y., Nöllenburg, M. (eds.) *Graph Drawing and Network Visualization*, vol. 9801 of LNCS, pp. 222–235. Springer, Cham (2016)
11. Di Battista, G., Tamassia, R.: On-line planarity testing. *SIAM J. Comput.* **25**(5), 956–997 (1996). <https://doi.org/10.1137/S0097539794280736>
12. Rettner, C.: Flipped Bitonic st -Orderings of Upward Plane Graphs [Bachelor's Thesis]. University of Würzburg (2019)
13. Bekos, M.A., Di Giacomo, E., Didimo, W., Liotta, G., Montecchiani, F.: Universal slope sets for upward planar drawings. In: Biedl, T.C., Kerren, A. (eds.) *Graph Drawing and Network Visualization*, vol. 11282 of LNCS, pp. 77–91. Springer, Cham (2018)

14. Tamassia, R., Tollis, I.G.: A unified approach to visibility representations of planar graphs. *Discrete Comput. Geom.* **1**, 321–341 (1986). <https://doi.org/10.1007/BF02187705>
15. Otten, R.H.J.M., van Wijk, J.G.: Graph representations in interactive layout design. In: *Proceedings of the IEEE International Symposium on Circuits and Systems*, pp. 914–918 (1978)
16. Di Battista, G., Tamassia, R.: On-line maintenance of triconnected components with SPQR-trees. *Algorithmica* **15**(4), 302–318 (1996). <https://doi.org/10.1007/BF01961541>
17. Gutwenger, C., Mutzel, P.: A linear time implementation of SPQR-trees. In: Marks, J. (ed.) *Graph Drawing*, vol. 1984 of LNCS, pp. 77–90. Springer, Cham (2000)
18. Bertolazzi, P., Di Battista, G., Didimo, W.: Computing orthogonal drawings with the minimum number of bends. *IEEE Trans. Comput.* **49**(8), 826–840 (2000). <https://doi.org/10.1109/12.868028>
19. Di Battista, G., Tamassia, R., Tollis, I.G.: Area requirement and symmetry display of planar upward drawings. *Discrete Comput. Geom.* **7**, 381–401 (1992). <https://doi.org/10.1007/BF02187850>
20. Frati, F.: On minimum area planar upward drawings of directed trees and other families of directed acyclic graphs. *Int. J. Comput. Geom. Appl.* **18**(3), 251–271 (2008). <https://doi.org/10.1142/S021819590800260X>
21. Lang, S., Murrow, G.: Volumes. In: *Geometry—A High School Course*, 2nd edn., pp. 261–294. Springer (1997)
22. Aigner, M., Ziegler, G.M.: Three applications of Euler’s formula: Pick’s theorem. In *Proofs from THE BOOK*, 6th edn., pp. 93–94. Springer (2018)
23. Pick, G.: Geometrisches zur Zahlenlehre. *Sonderabdr. Naturw.-medizin. Verein f. Böhmen “Lotos”* Nr. 8, 9 S. 8° (1899)
24. Di Battista, G., Tamassia, R.: Algorithms for plane representations of acyclic digraphs. *Theor. Comput. Sci.* **61**, 175–198 (1988). [https://doi.org/10.1016/0304-3975\(88\)90123-5](https://doi.org/10.1016/0304-3975(88)90123-5)

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

Authors and Affiliations

Patrizio Angelini¹ · Michael A. Bekos² · Henry Förster³  · Martin Gronemann⁴

Patrizio Angelini
pangelini@johncabot.edu

Michael A. Bekos
bekos@uoi.gr

Martin Gronemann
m Gronemann@ac.tuwien.ac.at

¹ Department of Mathematics, Natural, and Applied Sciences, John Cabot University, Via della Lungara 233, 00165 Rome, Lazio, Italy

² Department of Mathematics, University of Ioannina, University Campus, 45110 Ioannina, Greece

³ Department of Computer Science, University of Tübingen, Sand 13, 72076 Tübingen, Baden-Württemberg, Germany

⁴ Algorithms and Complexity Group, TU Wien, Favoritenstraße 9-11, 1040 Vienna, Vienna, Austria