

Book Embeddings and Point-Set Embeddings of Series-Parallel Digraphs*

Emilio Di Giacomo¹, Walter Didimo¹,
Giuseppe Liotta¹, and Stephen K. Wismath²

¹ Dipartimento di Ingegneria Elettronica e dell'Informazione,
Università degli Studi di Perugia, Perugia, Italy.

{digiacomo, didimo, liotta}@diei.unipg.it

² Department of Mathematics and Computer Science,
University of Lethbridge, Lethbridge, Alberta, Canada.
wismath@cs.uleth.ca

Abstract. An optimal $O(n)$ -time algorithm to compute an upward two-page book embedding of a series-parallel digraph with n vertices is presented. A previous algorithm of Alzohairi and Rival [1] runs in $O(n^3)$ time and assumes that the input series-parallel digraph does not have transitive edges. One consequence of our result is that series-parallel (undirected) graphs are necessarily sub-hamiltonian. This extends a previous result by Chung, Leighton, and Rosenberg [5] who proved sub-hamiltonicity for a subset of planar series-parallel graphs. Also, this paper investigates the problem of mapping planar digraphs onto a given set of points in the plane, so that the edges are drawn upward planar. This problem is called the upward point-set embedding problem. The equivalence between the problem of computing an upward two-page book embedding and an upward point-set embedding with at most one bend per edge on any given set of points is proved. An $O(n \log n)$ -time algorithm for computing an upward point-set embedding with at most one bend per edge on any given set of points for planar series-parallel digraphs is presented.

1 Introduction

A *book* consists of a line in Euclidean 3D space called the *spine* and some number of half-planes, called *pages*, having the spine as a boundary. A *book embedding* of a graph G is a linear ordering of the vertices along the spine of a book and an assignment of edges to pages such that edges assigned to the same page can be drawn as simple Jordan curves on that page without intersecting each other. The problem of computing a book embedding of a graph that uses the minimum number of pages has been the subject of intense research activities. (see e.g., [2,5,17]). In their seminal paper [2], Bernhart and Kainen prove that a planar

* Research supported in part by the Natural Sciences and Engineering Research Council of Canada, and by “Progetto Giovani Ricercatori: Algoritmi per la Visualizzazione di Grafi di Grandi Dimensioni” of the University of Perugia, Italy.

graph has a book embedding on two pages if and only if it is sub-hamiltonian. Yannakakis [17] shows that all planar graphs have a book embedding on four pages and that four pages are necessary for some planar graphs.

Also *upward book embeddings* of directed acyclic graphs and of posets have been widely investigated (see e.g., [1,11,12,13,15]). An upward book embedding of a digraph G is a book embedding of G such that the ordering of the vertices along the spine is a topological ordering of G . In contrast to the result by Yannakakis [17] concerning the book embeddability of undirected planar graphs, the minimum number of pages required by an upward book embedding of an upward planar digraph is not known [1,15]. Heath, Pemmaraju, and Trenk [13] show that oriented trees have a two-page upward book embedding. Alzohairi and Rival [1] show that a planar series-parallel ordered set has a two-page upward book embedding and that such an embedding can be computed in $O(n^3)$ time if the poset is a series-parallel planar lattice with n elements (note that the covering graph of a planar ordered set is an upward planar digraph).

In this paper we study upward book embeddings and find relations between book-embeddability and another relevant graph drawing question, the *point-set embeddability* problem. Before giving an overview of our main results, we briefly digress on this second topic to better clarify our contributions.

Let P be a set of n distinct points in the plane and let G be a planar graph with n vertices. A *point-set embedding* of G on P is a planar drawing of G such that each vertex is mapped to a distinct point of P and each edge is drawn as a polygonal chain connecting its endpoints (note that the mapping of the vertices to the points is not specified as part of the input). The problem of computing a point-set embedding of a graph while minimizing the number of bends per edge has been widely investigated (see e.g. [4,3,10,14]). Since outerplanar graphs are the largest class of graphs admitting a straight-line point-set embedding on any set of points [10], Kaufmann and Wiese [14] investigate the point-set embeddability problem for planar graphs allowing few bends per edge.

They show that if a graph has a hamiltonian cycle with an edge on the external face then it admits a point-set embedding with at most one bend per edge on any set of points on the plane and present an $O(n \log n)$ time algorithm to compute it. Also, they show that for a planar graph G with n vertices and for any set of n points on the plane, a point-set embedding of G with at most three bends per edge can be computed in $O(n \log n)$ time and that the number of bends per edge can be reduced to be at most two at the expense of increasing the time complexity to $O(n^2)$. The research of Kaufmann and Wiese is extended in this paper to *upward point-set embeddings*. An upward point-set embeddings of an upward planar digraph G on a set P of points in the plane is a point-set embedding of G onto P such that the edges are drawn monotonically increasing in a common direction. An overview of our results follows.

- An optimal $O(n)$ time algorithm to compute a two-page upward book embedding of a planar series-parallel digraph with n vertices is presented. The best previously known algorithm runs in $O(n^3)$ time and assumes that the input graph does not have transitive edges [1].

- The equivalence between the problem of computing a two-page upward book embedding and an upward point-set embedding with at most one bend per edge on any given set of points is proved.
- The investigation is extended to the undirected book-embeddability and the point set embeddability problem. One consequence of our result is that series-parallel (undirected) graphs are necessarily sub-hamiltonian.

2 Preliminaries

2.1 Series-Parallel Digraphs and Their Decomposition Tree

Let u and v be two vertices of a digraph. We use the notation $u \rightarrow v$ to say that there is a directed path from u to v in the digraph.

A *two terminal series-parallel digraph* (also called *TTSP-digraph* in the following) is a planar digraph recursively defined as follows [6,16]:

- An edge joining two vertices is a TTSP-digraph.
- Let G' and G'' be two TTSP-digraphs; the digraph obtained by identifying the sink of G' with the source of G'' (*Series Composition*) is also a TTSP-digraph.
- Let G' and G'' be two TTSP-digraphs; the digraph obtained by identifying the source of G' with the source of G'' and the sink of G' with the sink of G'' (*Parallel Composition*) is also a TTSP-digraph.

A TTSP-digraph has one source and one sink that are called its *poles*. The undirected underlying graph of a TTSP-digraph is called a *two terminal series-parallel graph* or *TTSP-graph* for short.

We recall the definition of the *SPQR-tree* of an *st-graph* G , also called the *decomposition tree* of G [7]. A *split pair* of G is either a separation pair or a pair of adjacent vertices. A *split component* of a split pair $\{u, v\}$ is either an edge (u, v) or a maximal subgraph C of G such that C is an uv -graph and $\{u, v\}$ is not a split pair of C . A *maximal split pair* $\{u, v\}$ of G is such that there is no other split pair $\{u', v'\}$ in G such that $\{u, v\}$ is contained in a split component of $\{u', v'\}$. The *SPQR-tree* T of an *st-graph* G is a rooted ordered tree whose nodes are of four types: *S*, *P*, *Q*, and *R*. Each node μ of T has an associated planar *st-graph* (possibly with multiple edges), called the *skeleton* of μ and denoted by $skeleton(\mu)$. Also it is associated with an edge of the skeleton of the parent ν of μ , called the *virtual edge* of μ in $skeleton(\nu)$. Tree T is recursively defined as follows (see Figure 3(b)):

Trivial case: If G consists of a single edge from s to t , then T consists of a single *Q*-node whose skeleton is G itself.

Series case: If G is not biconnected, let c_1, \dots, c_{k-1} ($k \geq 2$) be the cutvertices of G . Since G is planarly biconnectible, each cutvertex c_i is contained in exactly two blocks G_i and G_{i+1} such that s is in G_1 and t is in G_k . The root of T is an *S*-node μ . Graph $skeleton(\mu)$ consists of the chain e_1, \dots, e_k , where edge e_i goes from c_{i-1} to c_i , $c_0 = s$, and $c_k = t$.

Parallel case: If s and t are a split pair for G with split components G_1, \dots, G_k ($k \geq 2$), the root of T is a P -node μ . Graph $skeleton(\mu)$ consists of k parallel edges from s to t , denoted e_1, \dots, e_k .

Rigid case: If none of the above cases applies, let $\{s_1, t_1\}, \dots, \{s_k, t_k\}$ be the maximal split pairs of G ($k \geq 1$), and for $i = 1, \dots, k$, let G_i be the union of all the split components of $\{s_i, t_i\}$. The root of T is an R -node μ . Graph $skeleton(\mu)$ is obtained from G by replacing each subgraph G_i with edge e_i from s_i to t_i .

Note that for a TTSP-graph, the Rigid case never applies and thus the decomposition tree of a TTSP-graph does not contain any R -node.

2.2 Two Page Book Embeddings

Let G be a graph. Without loss of generality, we assume that in each drawing of a two-page book embedding of G , the spine is always vertical, and we call the half-plane to the left and to the right of the spine, the *left page* and the *right page* respectively. Also, we denote by L_G the linear ordering of the vertices of G on the spine, and, for each edge e of G , we denote by d_e the page (left or right) to which e is assigned.

The following theorem gives a simple characterization of two-page upward book embeddings of planar digraphs, that will be useful in the following.

Theorem 1. *Let G be a planar digraph. The linear ordering L_G and the assignment d_e for every edge e of G define a two-page upward book embedding of G if and only if the following two properties hold:*

UBE1: *If (u, v) is a directed edge of G then u precedes v in L_G .*

UBE2: *There are no two edges $e_1 = (u, v)$ and $e_2 = (w, z)$ (with u, v, w, z all distinct) for which $d_{e_1} = d_{e_2}$ and u, w, v, z occur in this order in L_G .*

Proof. Given an upward book embedding of G , the ordering L_G of the vertices of G along the spine is a topological ordering of the vertices of G , hence Property UBE1 is verified. Also, suppose for a contradiction that Property UBE2 is not satisfied. This implies that there exist two edges $e_1 = (u, v)$ and $e_2 = (w, z)$ (with u, v, w, z all distinct) for which $d_{e_1} = d_{e_2}$ and u, w, v, z occur in this order in L_G . In this case, however, e_1 and e_2 can not be drawn on the same page without intersecting each other (see Figure 1(a)), which is a contradiction. Hence, also Property UBE2 is verified.

Suppose that L_G and the assignment d_e (for every edge e of G) are such that Properties UBE1 and UBE2 hold. Property UBE1 implies that L_G defines a topological ordering of the vertices on the spine. Also, consider any two edges $e_1 = (u, v)$ and $e_2 = (w, z)$ of G . From Property UBE1, u precedes v , and w precedes z in L_G .

- If $d_{e_1} \neq d_{e_2}$, that is, the two edges are assigned to different pages, then it is always possible to draw e_1 and e_2 without intersecting each other (see for example Figure 1(b)).

- If $d_{e_1} = d_{e_2}$, from Property UBE2, only the following sub-cases are possible (we do not examine symmetric sub-cases): (i) Edges e_1 and e_2 have no common vertex, and u, v, w, z occur in this order in L_G . Edge e_2 can be drawn completely above e_1 (see Figure 1(c)). (ii) Edges e_1 and e_2 have no common vertex, and u, w, z, v occur in this order in L_G . Edge e_2 can be drawn nested into e_1 (see Figure 1(d)). (iii) Edges e_1 and e_2 have one common vertex. The two edges can be drawn either nesting one edge into another (see Figure 1(e)) or putting one edge completely above the other (see Figure 1(f)).




Fig. 1. Illustration of the cases of Theorem 1

Given a two-page book embedding of G , a drawing of this book embedding can be easily computed mapping the vertices of G at consecutive points on a vertical line (the spine of the book), according to the ordering L_G , and drawing each edge e as a semi-circle between its end-vertices on page d_e with a diameter chosen as the distance between the end-vertices of e (see Figure 3(d)).

3 Two-Page Upward Book Embeddings of Series-Parallel Digraphs in Linear Time

In this section we describe a linear time algorithm for computing a two-page upward book embedding of a TTSP-digraph G . We can assume that edge (s, t) exists in G . If it does not exist, we temporarily add such an edge and then remove it at the end of the algorithm. Finally, since s and t are on the external face, edge (s, t) can always be embedded on the external face.

Algorithm SP-Book Embedder

Input: A TTSP-digraph G with edge (s, t) on the external face.

Output: A two-page upward book embedding of G , defined by a linear ordering L_G and by a page assignment d_e for each edge e of G .

Step 1: Compute the $SPQR$ -tree T of G . Since G is a TTSP-digraph, T has no R nodes.

- For each node μ of T denote by G_μ the pertinent digraph of μ , and let L_μ be a linear list used to define a linear ordering of the vertices of G_μ .
- Define the following *merge* operation \circ between two ordered vertex lists $L_\mu = \langle u_1, u_2, \dots, u_k \rangle$ and $L_\nu = \langle v_1, v_2, \dots, v_h \rangle$:
 - $L_\mu \circ L_\nu = \langle u_1, u_2, \dots, u_k, v_1, v_2, \dots, v_h \rangle$, if $u_k \neq v_1$.
 - $L_\mu \circ L_\nu = \langle u_1, u_2, \dots, u_k, v_2, \dots, v_h \rangle$, if $u_k = v_1$.

Step 2: Denote by $\{\mu_1, \mu_2, \dots, \mu_h\}$ the nodes of T ordered according to a post-order visit.

Step 3: For $i = 1, \dots, h$ do:

Case 1 μ_i is a *Q*-node (see Figure 2(a)). Let $e = (u, v)$ be the directed edge represented by μ_i .

- Set $L_{\mu_i} = \langle u, v \rangle$.
- Set $d_e = \textit{left}$.

Case 2 μ_i is an *S*-node (see Figure 2(b)). Let $(u = u_1, u_2), (u_2, u_3), \dots, (u_k, u_{k+1} = v)$ be the edges of the skeleton of μ_i where $u \rightarrow v$ in G , and let $\nu_j, (j = 1, \dots, k)$ be the child of μ_i in T corresponding to edge (u_j, u_{j+1}) . Node ν_k will be called the *top child* of μ_i in the following.

- Set $L_{\mu_i} = L_{\nu_1} \circ L_{\nu_2} \circ \dots \circ L_{\nu_k}$.
- If ν_k is a *Q*-node representing an edge e_k of G , then set $d_{e_k} = \textit{right}$.

Case 3 μ_i is a *P*-node (see Figure 2(c)). Let u, v be the poles of the skeleton of μ_i , such that $u \rightarrow v$. Let $\nu_1, \nu_2, \dots, \nu_k$ be the children of μ in T and assume that if the edge (u, v) exists in G_{μ_i} , the corresponding *Q*-node is ν_k . Also, denote by $L_{\nu_j}^-$ the list L_{ν_j} without the first vertex u and the last vertex v , ($j = 1, \dots, k$).

- Set $L_{\mu_i} = \langle u \rangle \circ L_{\nu_k}^- \circ L_{\nu_{k-1}}^- \circ \dots \circ L_{\nu_1}^- \circ \langle v \rangle$.
- If ν_k is a *Q*-node representing an edge e_k and if μ_i is the top child of an *S*-node in T , set $d_{e_k} = \textit{right}$.

Step 4: Set $L_G = L_{\mu_h}$.




Fig. 2. Illustration of the cases of Algorithm *SP-Book Embedder*

In order to prove the correctness of Algorithm *SP-Book Embedder*, we need the following preliminary lemma.

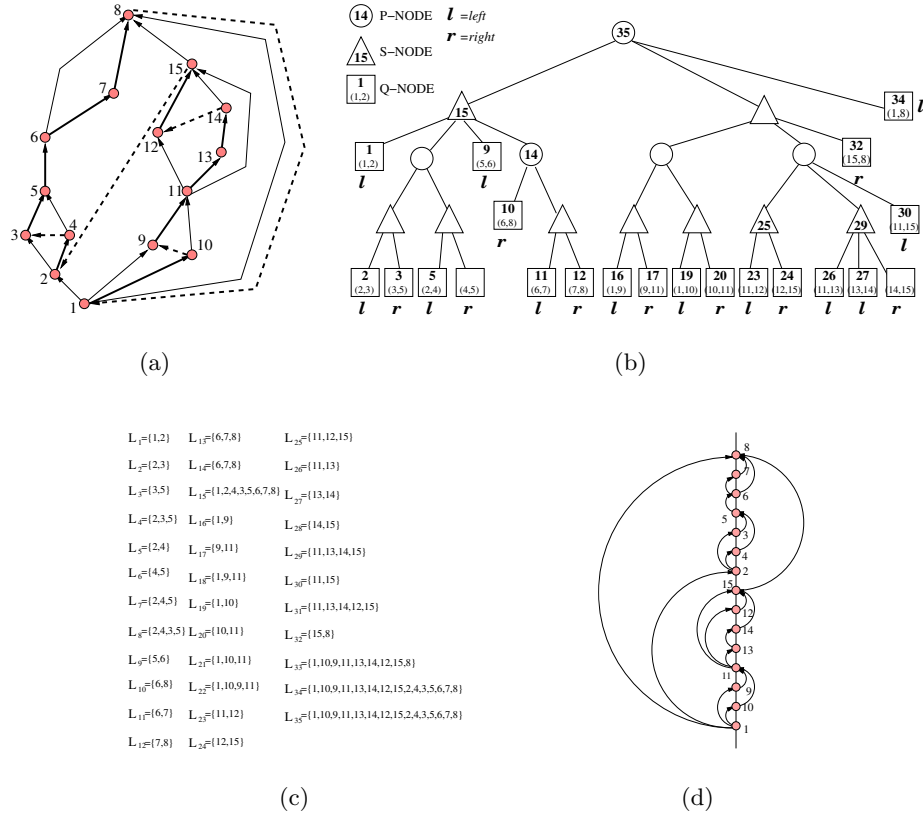


Fig. 3. (a) A TTSP-digraph G . The dashed edges are dummy edges that form an augmented hamiltonian cycle along with the bold edges. (b) The $SPQR$ -tree of G . For each node the number given by the post order visit is shown. For each Q -node the edge associated and the page assigned are also shown. (c) The lists L_μ for each node μ . (d) The final book embedding.

Lemma 1. *Let G be a TTSP-digraph, let T be its $SPQR$ -tree, and let μ be an S -node of T . Let G_μ be the pertinent digraph of μ and let u and v be the poles of μ such that $u \rightarrow v$. Algorithm **SP-Book Embedder** assigns the edges of G_μ to the pages in such a way that each edge incident on u is assigned to the left page and each edge incident on v is assigned to the right page.*

Proof. We prove the statement by induction on the height h_μ of the subtree $T_\mu \subset T$ rooted at μ (note that h_μ is always an odd number). If $h_\mu = 1$, then all children ν_1, \dots, ν_k of μ are Q -nodes. Let e_1, \dots, e_k be the edges associated with ν_1, \dots, ν_k , respectively; vertex u is the starting vertex of e_1 and v is the end vertex of e_k . Algorithm **SP-Book Embedder** sets $d_{e_i} = left$ for $1 \leq i \leq (k - 1)$

(Case 1 of Step 3) and $d_{e_k} = \text{right}$ (Case 2 of Step 3). Hence, the lemma holds when $h_\mu = 1$.

Let us assume that the lemma is true for $h_\mu = k \geq 1$ and consider the case $h_\mu = k + 2$. The children ν_1, \dots, ν_k of μ are either P -nodes or Q -nodes. In particular, in order to verify that the lemma is true, we only need to verify the assignment on the edges incident on u and v , that is, the edges of the pertinent digraphs of ν_1 and ν_k .

- Consider node ν_1 : If it is a Q -node representing an edge e , then the algorithm sets $d_e = \text{left}$ (Case 1 of Step 3); if ν_1 is a P -node then the children of ν_1 are all S -nodes except at most one which is a Q -node. For each child which is an S -node, the statement is true by the inductive hypothesis. For the edge e associated with the only child (if present) which is Q -node the algorithm sets $d_e = \text{left}$ (Case 1 of Step 3).
- Consider node ν_k : If it is a Q -node representing an edge e , then the algorithm sets $d_e = \text{right}$, since it is the top child of μ (Case 2 of Step 3); if ν_k is a P -node then the children of ν_k are all S -nodes except at most one, which is a Q -node. For each child which is an S -node the statement is true by the inductive hypothesis. For the edge e associated with the only child (if present) which is a Q -node the algorithm sets $d_e = \text{right}$, since ν_k is the top child of μ (Case 3 of Step 3).

Theorem 2. *Let G be a TTSP-digraph with n vertices. There exists an algorithm that computes a two-page upward book embedding of G in $O(n)$ time and space.*

Proof. Let T be the SPQR-tree of G , and denote by G_μ the pertinent digraph of a node μ of T . We prove by induction that, after the i -th iteration of Step 3 has been executed, a two-page upward book embedding of G_{μ_i} is correctly given by list L_{μ_i} and by variable d_e , for each edge e of G_{μ_i} . Namely, based on Theorem 1 we prove that Properties UBE1 and UBE2 hold for G_{μ_i} . We adopt the same notation used for describing Algorithm SP-Book Embedder.

Since Step 3 iterates on the nodes of T in post-order, μ_1 is a leaf of T , that is, a Q -node. Hence G_{μ_1} is a single edge $e = (u, v)$, and Case 1 applies. In this case the proof is trivial. Suppose the statement is true for $i \geq 1$; we prove that it is true for $i + 1$. There are three possible cases:

1. μ_{i+1} is a Q -node. In this case $G_{\mu_{i+1}}$ is a single edge and the proof is trivial.
2. μ_{i+1} is an S -node. Let $\nu_1 \dots \nu_k$ be the children of μ_{i+1} . For each edge $e = (w, z)$ of $G_{\mu_{i+1}}$, w and z belong to the same list L_{ν_j} for some $1 \leq j \leq k$. Since Property UBE1 is verified for L_{ν_j} and since the merge operation keeps unchanged the ordering of the vertices in L_{ν_j} , Property UBE1 holds for $L_{\mu_{i+1}}$. Let $e_1 = (w_1, z_1)$ and $e_2 = (w_2, z_2)$ be two edges of $G_{\mu_{i+1}}$ such that w_1, z_1, w_2, z_2 are all different. If w_1, z_1, w_2, z_2 belong to the same list L_{ν_j} for some $1 \leq j \leq k$ then Property UBE2 follows from the inductive hypothesis. Otherwise, $w_1, z_1 \in L_{\nu_{j_1}}$ and $w_2, z_2 \in L_{\nu_{j_2}}$ with $j_1 \neq j_2$. Assume, without

loss of generality, that $j_1 < j_2$. Since $L_{\mu_{i+1}} = L_{\nu_1} \circ L_{\nu_2} \circ \dots \circ L_{\nu_k}$, all the vertices in $L_{\nu_{j_1}}$ precede all the vertices in $L_{\nu_{j_2}}$. Thus Property UBE2 holds.

3. μ_{i+1} is a *P-node*. Let $\nu_1 \dots \nu_k$ be the children of μ_{i+1} , and let u and v be the poles of μ_{i+1} where $u \rightarrow v$.

We prove first that Property UBE1 is satisfied. Let $e = (w, z)$ be a directed edge of $G_{\mu_{i+1}}$. There are three cases to consider:

- if $w = u$ then w is the first vertex in $L_{\mu_{i+1}}$ and so Property UBE1 holds;
- if $z = v$ then z is the last vertex in $L_{\mu_{i+1}}$ and so Property UBE1 holds;
- If $w \neq u$ and $z \neq v$ then $w, z \in L_{\nu_j}^-$ for some $1 \leq j \leq k$. Since the merge operation keeps unchanged the ordering of the vertices in $L_{\nu_j}^-$ Property UBE1 follows from the inductive hypothesis.

We now show that Property UBE2 is true. Let $e_1 = (w_1, z_1)$ and $e_2 = (w_2, z_2)$ be two directed edges of $G_{\mu_{i+1}}$ such that w_1, z_1, w_2, z_2 are all different. Also in this case we have three cases:

- $w_j \neq u$ $j = 1, 2$ and $z_j \neq v$ $j = 1, 2$. In this case $w_1, z_1 \in L_{\nu_{j_1}}^-$ for some $1 \leq j_1 \leq k$ and $w_2, z_2 \in L_{\nu_{j_2}}^-$ for some $1 \leq j_2 \leq k$.
 - If $j_1 = j_2 = j$ then the ordering of w_1, z_1, w_2, z_2 in $L_{\mu_{i+1}}$ is the same as in $L_{\nu_j}^-$ and Property UBE2 follows from the inductive hypothesis.
 - If $j_1 \neq j_2$ (assume without loss of generality $j_1 < j_2$). Since $L_{\mu_{i+1}} = \langle u \rangle \circ L_{\nu_k}^- \circ \dots \circ L_{\nu_1}^- \langle v \rangle$ all the vertices in $L_{\nu_{j_2}}^-$ precedes all the vertices in $L_{\nu_{j_1}}^-$. Thus the ordering of w_1, z_1, w_2, z_2 in $L_{\mu_{i+1}}$ is w_2, z_2, w_1, z_1 , and Property UBE2 holds.
- One of w_1, z_1, w_2, z_2 is equal to u or v (without loss of generality assume that $w_1 = u$, other cases are similar). Let $w_1, z_1 \in L_{\nu_{j_1}}^-$ and $w_2, z_2 \in L_{\nu_{j_2}}^-$.
 - If $j_1 = j_2 = j$ then the ordering of w_1, z_1, w_2, z_2 in $L_{\mu_{i+1}}$ is the same as in $L_{\nu_j}^-$ since the merge operation keeps unchanged the ordering of the vertices $z_1, w_2, z_2 \in L_{\nu_j}^-$ and the vertex $w_1 = u$ which precedes all z_1, w_2, z_2 in $L_{\nu_j}^-$ also precedes all z_1, w_2, z_2 in $L_{\mu_{i+1}}$. Thus Property UBE2 follows from the inductive hypothesis.
 - If $j_1 < j_2$ then all the vertices in $L_{\nu_{j_2}}^-$ precede all the vertices in $L_{\nu_{j_1}}^-$, and $w_1 = u$ precedes all other vertices in $L_{\mu_{i+1}}$. Thus the ordering of w_1, z_1, w_2, z_2 in $L_{\mu_{i+1}}$ is $w_1 = u, w_2, z_2, z_1$, and Property UBE2 holds.
 - If $j_1 > j_2$ then all the vertices in $L_{\nu_{j_1}}^-$ precede all the vertices in $L_{\nu_{j_2}}^-$, and $w_1 = u$ precedes all other vertices in $L_{\mu_{i+1}}$. Thus the ordering of w_1, z_1, w_2, z_2 in $L_{\mu_{i+1}}$ is $w_1 = u, z_1, w_2, z_2$, and Property UBE2 holds.
- Two of w_1, z_1, w_2, z_2 are equal to u or v .
 - If $w_1 = u$ and $z_1 = v$ (or symmetrically $w_2 = u$ and $z_2 = v$) $w_1 = u$ ($w_2 = u$) is the first vertex in $L_{\mu_{i+1}}$, and $z_1 = v$ ($z_2 = v$) is the last vertex in $L_{\mu_{i+1}}$. Thus the ordering of w_1, z_1, w_2, z_2 in $L_{\mu_{i+1}}$ is $w_1 = u, w_2, z_2, z_1 = v$ (symmetrically $w_2 = u, w_1, z_1, z_2 = v$), and Property UBE2 holds.
 - If $w_1 = u$ and $z_2 = v$ (or symmetrically $w_2 = u$ and $z_1 = v$) then one among e_1 and e_2 is incident on u and the other is incident on v . By Lemma 1 they are on different pages and Property UBE2 holds.

Proofs about time and space complexity are omitted for reasons of space.

4 Hamiltonicity, Book and Point-Set Embeddability

In their investigation of book-embeddings, Bernhart and Kainen [2] show that an undirected planar graph admits a book embedding on two pages if and only if it is sub-hamiltonian, *i.e.* it can be augmented with edges to create a hamiltonian cycle while still remaining planar. In a subsequent paper, Chung, Leighton, and Rosenberg [5] show that a family of planar graphs that they call series-parallel graphs have a book embedding on two pages and therefore are sub-hamiltonian. However, it must be noted that the definition of Chung, Leighton, and Rosenberg appears somewhat restrictive if compared to the one commonly used for series-parallel graphs by the vast majority of researchers (see e.g. [6,16]). In particular, in [5] the series composition is the usual one, but the parallel composition is defined as follows: Let G_1, \dots, G_n be series-parallel graphs with source vertices s_1, \dots, s_n and sink vertices t_1, \dots, t_n ; the parallel composition of G_1, \dots, G_n is obtained by taking a new source s and a new sink t and by adding edges between each s_i and s and between each t_i and t .

Clearly the family of series-parallel graphs defined by Chung, Leighton, and Rosenberg is a proper subset of the class of TTSP-graphs. For example, a three-cycle is a TTSP-graph, but is not a series-parallel graph as defined by Chung *et al.* Furthermore, it is not obvious how to extend the proof of sub-hamiltonicity given in [5] to all TTSP-graphs. On the other hand, we observe that, since each TTSP undirected graph can be *st*-oriented in linear time with s and t on the external face [8], Algorithm **SP-Book Embedder** can also be used for computing a two-page book embedding of a two terminal series-parallel graph. This observation, together with the result by Bernhart and Kainen [2] leads to the following. We recall that the problem of determining if a graph is sub-hamiltonian is NP-Complete [5]

Theorem 3. *Let G be a TTSP-graph with n vertices. Then G admits a two-page book embedding and therefore it is sub-hamiltonian. Also there exists an $O(n)$ -time algorithm that computes an augmented hamiltonian cycle of G .*

In a recent paper, Ganley and Heath [9] conjecture that every k -tree has a book-embedding on exactly k pages. Since every 2-tree is a TTSP-graph, a consequence of Theorem 3 is to prove their conjecture for $k = 2$.

Corollary 1. *Every 2-tree has a book embedding on two pages.*

As a second application of the results in Section 3, we consider the upward point-set embeddability problem, that is the problem of computing a point-set embedding of an upward planar digraph on a given set of points by minimizing the number of bends along the edges and by drawing all edges monotonically increasing in a common direction. We remark that the upward point-set embeddability problem is, in general, different from the non-upward problem. For example, let G be the upward planar graph of Figure 4. Graph G does not admit an upward point set embedding if the points of P are collinear as in Figure 4(b). Conversely, Figure 4(c) shows a non upward point-set embedding of G on P with at most one bend per edge.




Fig. 4. (a) An upward planar graph G . (b) There is no upward point-set embeddings of G on P . (c) A (non-upward) point-set embedding of G on a set of points P on the same line.

Theorem 4. *An upward planar digraph G with n vertices has an upward point-set embedding with at most one bend per edge on every set of n points if and only if G has a two-page upward book embedding. Also, from a two-page upward book embedding of G it is possible to construct in $O(n \log n)$ time an upward point-set embedding of G with at most one bend per edge on every set of n points.*

Proof. If G has an upward point-set embedding with at most one bend per edge on every set of n points then it has an upward point-set embedding on consecutive points of a straight line l . In this upward point-set embedding the edges of G are drawn either on the half-plane to the left of l (assuming l to be a vertical line), or on the half-plane to the right of l . The ordering of the vertices of G along l and the assignment of the edges of G to the left or right half-plane define the two-page upward book embedding of G .

Consider now a two-page upward book embedding of G and let $v_i, i = 1, \dots, n$ be the vertices of G in the order they appear along the spine. Since G has a two-page upward book embedding, G is sub-hamiltonian, and so we can add a certain number of edges in order to create a hamiltonian cycle C . Let $C = (v_1, v_2)(v_2, v_3) \dots (v_{n-1}, v_n)(v_n, v_1)$, *i.e.*, the ordering of the vertices in C is set to be the ordering of the vertices along the spine. In order to compute an upward point-set embedding of G on P with at most one bend per edge, we apply the algorithm described by Kaufmann and Wiese [14]. Namely, given a planar graph G' with a hamiltonian cycle $(v_1, v_2, \dots, v_n, v_1)$, the algorithm in [14] computes a point-set embedding of G' on P with at most one bend per edge such that v_i is mapped to p_i ($i = 1, \dots, n$). About time complexity, we just observe that the algorithm in [14] takes $O(n \log n)$ time and that the operations of ordering the points of P according to their y -coordinates and possibly rotating the plane can be performed in $O(n \log n)$ time.

The equivalence shown above is also valid for the non-upward case. Thus combining Theorem 4, Theorem 3 and, Theorem 2, we have the following.

Corollary 2. *Let G be a TTSP-(di)graph with n vertices. There exists an $O(n \log n)$ -time algorithm that computes an (upward) point-set embedding of G with at most one bend per edge on every set of n distinct points on the plane.*

References

1. M. Alzohairi and I. Rival. Series-parallel ordered sets havepagenumber two. In S. North, editor, *Graph Drawing (Proc. GD '96)*, volume 1190 of *Lecture Notes Comput. Sci.*, pages 11–24. Springer-Verlag, 1996.
2. F. Bernhart and P. C. Kainen. The book thickness of a graph. *J. Combin. Theory, Ser. B* 27:320–331, 1979.
3. P. Bose. On embedding an outer-planar graph on a point set. In G. Di Battista, editor, *Graph Drawing (Proc. GD '97)*, volume 1353 of *Lecture Notes Comput. Sci.*, pages 25–36. Springer-Verlag, 1997.
4. P. Bose, M. McAllister, and J. Snoeyink. Optimal algorithms to embed trees in a point set. *Journal of Graph Algorithms and Applications*, 2(1):1–15, 1997.
5. F. R. K. Chung, F. T. Leighton, and A. Rosenberg. Embedding graphs in books: A layout problem with applications to VLSI design. *SIAM Journal on Algebraic and Discrete Methods*, 8:33–58, 1987.
6. G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing*. Prentice Hall, Upper Saddle River, NJ, 1999.
7. G. Di Battista and R. Tamassia. On-line planarity testing. *SIAM Journal on Computing*, 25(5):956–997, October 1996.
8. S. Even and R. E. Tarjan. Computing an st-numbering. *Theoret. Comput. Sci.*, 2:339–344, 1976.
9. J. Ganley and L. Heath. Thepagenumber of k -trees is $O(k)$. *Discrete Applied Mathematics*, 109:215–221, 2001.
10. P. Gritzmann, B. Mohar, J. Pach, and R. Pollack. Embedding a planar triangulation with vertices at specified points. *Amer. Math. Monthly*, 98(2):165–166, 1991.
11. L. S. Heath and S. V. Pemmaraju. Stack and queue layouts of posets. *SIAM Journal on Discrete Mathematics*, 10:599–625, 1997.
12. L. S. Heath and S. V. Pemmaraju. Stack and queue layouts of directed acyclic graphs: Part II. *SIAM Journal on Computing*, 28:1588–1626, 1999.
13. L. S. Heath, S. V. Pemmaraju, and A. Trenk. Stack and queue layouts of directed acyclic graphs: Part I. *SIAM Journal on Computing*, 28:1510–1539, 1999.
14. M. Kaufmann and R. Wiese. Embedding vertices at points: Few bends suffice for planar graphs. *Journal of Graph Algorithms and Applications*, 6(1):115–129, 2002.
15. R. Nowakowski and A. Parker. Ordered sets,pagenumbers and planarity. *Order*, 6:209–218, 1989.
16. J. Valdes, R. E. Tarjan, and E. Lawler. The recognition of series-parallel digraphs. *SIAM J. Comput.*, 11(2):298–313, 1982.
17. M. Yannakakis. Embedding planar graphs in four pages. *Journal of Computer and System Sciences*, 38:36–67, 1989.