Edge-Orders

Lena Schlipf[1] and Jens M. Schmidt[*][2]

1   LG Theoretische Informatik, FernUniversität in Hagen, Hagen, Germany
    lena.schlipf@fernuni-hagen.de
2   Institut für Mathematik, TU Ilmenau, Ilmenau, Germany
    jens.schmidt@tu-ilmenau.de

──── **Abstract** ────

Canonical orderings and their relatives such as st-numberings have been used as a key tool in algorithmic graph theory for the last decades. Recently, a unifying link behind all these orders has been shown that links them to well-known graph decompositions into parts that have a prescribed vertex-connectivity.

Despite extensive interest in canonical orderings, no analogue of this unifying concept is known for edge-connectivity. In this paper, we establish such a concept named *edge-orders* and show how to compute *(1,1)-edge-orders* of 2-edge-connected graphs as well as *(2,1)-edge-orders* of 3-edge-connected graphs in linear time, respectively. While the former can be seen as the edge-variants of *st-numberings*, the latter are the edge-variants of *Mondshein sequences* and *non-separating ear decompositions*. The methods that we use for obtaining such edge-orders differ considerably in almost all details from the ones used for their vertex-counterparts, as different graph-theoretic constructions are used in the inductive proof and standard reductions from edge- to vertex-connectivity are bound to fail.

As a first application, we consider the famous *Edge-Independent Spanning Tree Conjecture*, which asserts that every k-edge-connected graph contains k rooted spanning trees that are pairwise edge-independent. We illustrate the impact of the above edge-orders by deducing algorithms that construct 2- and 3-edge independent spanning trees of 2- and 3-edge-connected graphs, the latter of which improves the best known running time from $O(n^2)$ to linear time.

1   Introduction

Canonical orderings serve as a fundamental tool in various fields of algorithmic graph theory, see [2, 26] for a wealth of over 30 applications. Under this name, canonical orderings were published in 1988 for maximal planar graphs [8] and soon after generalized to 3-connected planar graphs [14]. Interestingly, it turned out only recently [26] that the well-known *non-separating ear decompositions* [6] are in fact strict generalizations of canonical orderings to arbitrary 3-connected graphs, and that this generalization was, independently, already known as *(2,1)-sequences* [19] in 1971 long before canonical orderings were even proposed (anticipating many of their later planar features).

Mondshein [19] characterized (2,1)-sequences, or *(2,1)-orders*, as we will call them, by decomposing a graph into 2-connected and connected parts. Indeed, the unifying link above

■ **Table 1** *Left*: (k, l)-orders of $(k + l)$-connected graphs known so far and the best-known running times for constructing them. *Right*: (k, l)-edge-orders of $(k + l)$-edge-connected graphs (this paper).

$k \backslash l$	1	2		$k \backslash l$	1
1	st-numbering [9] $O(m)$			1	st-edge-numbering [1] $O(m)$ *(+in this paper)*
2	Mondshein sequence [25] $O(m)$	Chain decomposition [7] $O(n^2 m)$; if planar [21] $O(m)$		2	(2,1)-edge-order $O(m)$ *(in this paper)*
3	(3,1)-order for triangulations [4] $O(m)$	5-canonical decomposition for triangulations [20] $O(m)$		3	
4				4	

allows to describe any canonical ordering of a graph $G = (V, E)$ as a total order on V such that for certain i, the first i vertices induce a 2-connected graph and the remaining vertices induce a connected graph in G [26] (and hence, does not use any reference to planarity). The general concept behind canonical orderings is thus connectivity, with all of its implications for planarity, instead of planarity itself.

Several publications [20, 7, 4] extended this approach to (k, l)-*orders* with $(k, l) \neq (2, 1)$. Such (k, l)-orders may be described canonically as total orders on V such that for certain i, the first i vertices induce a k-connected graph and the remaining vertices induce a l-connected graph (a related description for planar triangulations is given in [4]). We note that this is not a definition, as "certain i" has to be quantified for every particular (k, l). This is usually done in dependence of a graph decomposition, which tend to become more complex, as k or l grow: e.g. for $(2, 1)$-orders, "certain i" is quantified by taking every vertex i that completes an ear with the predecessors of i in a fixed open ear decomposition of G.

Several relatives of $(2,1)$-orders fit into the context of (k, l)-orders: The well-known *st-numberings* and *st-orientations* are actually $(1,1)$-orders of 2-connected graphs, where i ranges over all vertices, the *chain decompositions* of [7] are $(2,2)$-orders of 4-connected graphs, and more orders on restricted graph classes such as planar graphs and triangulations are known (see Table 1 left).

The purpose of this paper is to extend this unifying view further to (k, l)-*edge-orders*, each of which can be described as a total order on E such that for certain i, the first i edges induce a k-edge-connected graph and the remaining edges induce a l-edge-connected graph. Despite the many known and heavily used vertex-orders above, these natural edge-variants do not seem to be well-studied. In fact, we are only aware of one technical report by Annexstein et al. [1], which deals with $(1,1)$-edge-orders (under the name *st-edge-orderings*). For the $(1, 1)$-edge-order we present, i ranges over all edges except st; for the $(2, 1)$-edge-order, i ranges over all edges that complete an ear with the predecessors of i in a fixed ear decomposition of G.

We show a simple algorithm how a $(1,1)$-edge-order can be computed and prove that it has running time $O(m)$. Our main contribution is then an algorithm that computes a $(2,1)$-edge-order of a 3-edge-connected graph in time $O(m)$ (see Table 1 right), of which the corresponding result for the vertex-counterpart took over 40 years.

Just like $(2,1)$-orders, which immediately led to improvements on the best-known running time for five applications [5, 26], $(2,1)$-edge-orders seem to be an important and useful tool for many graph algorithms. We give an application of them, which is related to the edge-independent spanning tree conjecture [13]: By using a $(2,1)$-edge-order, we show that three edge-independent spanning trees of 3-edge-connected graphs can be computed in time $O(m)$, improving the best-known running time $O(n^2)$ by Gopalan et al. [11].

We also considered the *3-edge-partition problem*, but surprisingly did not find an easy reduction to $(2,1)$-edge-orders. However, we note that this problem can be solved in linear time using existing algorithms: A 3-edge-partition can be computed by two linear-time reductions, first to the *vertex-subset tripartitioning problem* [28, Theorem 2b], and then [27] to the problem of computing a non-separating ear decomposition. It is also possible to find an alternative simple and direct linear-time reduction along the lines of [26, Application 5].

After giving preliminary facts on ear decompositions, we explain the linear-time algorithms for computing (1,1)- and (2,1)-edge-orders in Sections 3–5. Section 6 then shows algorithms for computing two and three edge-independent spanning trees.

1.1 Vertex-connectivity vs. edge-connectivity

In many cases, the vertex-variant of a connectivity problem is more challenging than its edge-variant, as the latter may be reduced to the former by taking its line-graph or by using the reduction from k-edge- to k-vertex-connectivity of Galil and Italiano [10]. From a top-level perspective, our $(2,1)$-edge-order algorithm follows the proof outline of its vertex-counterpart in [26]. Thus, it needs to be motivated that there is no obvious linear-time reduction to [26] that produces the results of this paper (of course there is a non-obvious reduction that just takes the algorithm of this paper and does not invoke [26] at all).

Clearly, a reduction to line-graphs is not possible, as this may involve a quadratic blow-up in the graph size and thus in the running time. Using the reduction of Galil-Italiano, we can reduce a 3-edge-connected graph G to a 3-connected graph G', and then compute a $(2,1)$-order of G' in linear time using [26]. However, it can be shown that there is no obvious way of transforming the $(2,1)$-order of G' back to a $(2,1)$-edge-order of G.

Another hint that such a reduction might be elusive is given by our application to edge-independent spanning trees. Despite extensive research, it is still not known how to reduce these to vertex-independent spanning trees (which may in turn be computed from a $(2,1)$-order [26]), not even for the corresponding existence results. In fact, an attempt trying to prove this turned out to be false [12]. If there was a reduction to $(2,1)$-orders, it would directly imply a reduction to vertex-independent spanning trees.

Hence, there is no obvious way of producing our results using old ones. Indeed, the different parts of our proof require substantially new ideas and non-trivial formalizations in comparison to [26]: Mader-sequences differ from the (BG)-sequences used in [26] (and, although they are not too far apart, it took a 27-page paper to show that the former can be computed in linear time as well [18]), the notions of non-separateness and $\overline{G_i}$ differ considerably, and, here, we need last-values in addition to just birth-values.

2 Preliminaries

We use standard graph-theoretic terminology and consider only graphs that are finite and undirected, but may contain parallel edges and self-loops. In particular, cycles may have length one or two. A separator of size one is called a *cut-vertex*. The *2-connected components* of a graph are its inclusion-wise maximal connected subgraphs having no cut-vertex. For $k \geq 1$, let a graph G be *k-edge-connected* if $n := |V| \geq 2$ and G has no edge-cut of size less than k.

▶ **Definition 1** ([15, 29]). An *ear decomposition* of a graph $G = (V, E)$ is a sequence $(P_0, P_1, \ldots, P_k)$ of subgraphs of G that partition E such that (i) P_0 is a cycle that is no self-loop and (ii) every P_i, $1 \leq i \leq k$, is either a path that intersects $P_0 \cup \cdots \cup P_{i-1}$ in its

endpoints or a cycle that intersects $P_0 \cup \cdots \cup P_{i-1}$ in a unique vertex q_i (which we call *endpoint* as well). Each P_i is called an *ear*. An ear is *short* if it is an edge and *long* otherwise.

▶ **Theorem 2** ([22]). *A graph is 2-edge-connected if and only if it has an ear decomposition.*

According to Whitney [29], every ear decomposition has exactly $m - n + 1$ ears ($m := |E|$). For any i, let $G_i = (V_i, E_i) := P_0 \cup \cdots \cup P_i$ and $\overline{E_i} := E - E_i$. We denote the subgraph of G that is induced by $\overline{E_i}$ as $\overline{G_i} = (\overline{V_i}, \overline{E_i})$. Clearly, $\overline{G_j} \subset \overline{G_i}$ for every $i < j$. We note that this definition of $\overline{G_i}$ differs from the definition $\overline{G_i} := G - V_i$ that was used for (2,1)-vertex-orders [26], due to the weaker edge-connectivity assumption.

For any ear P_i, let $inner(P_i) := V(P_i) - G_{i-1}$ be the set of *inner vertices* of P_i (for P_0, every vertex is an inner vertex). Hence, for a cycle $P_i \neq P_0$, $inner(P_i) = V(P_i) - q_i$. Every vertex of G is an inner vertex of exactly one long ear, which implies that, in an ear decomposition, the inner vertex sets of the long ears partition V.

▶ **Definition 3.** Let $D = (P_0, P_1, \ldots, P_{m-n})$ be an ear decomposition of G. For an edge e, let $birth_D(e)$ be the index i such that P_i contains e. For a vertex v, let $birth_D(v)$ be the index i such that P_i contains v as inner vertex and let $last_D(v)$ be the maximal index $birth(vw)$ over all neighbors w of v. Whenever D is clear from the context, we will omit the subscript D.

Thus, $P_{last(v)}$ is the last ear that contains v and, seen from another perspective, the first ear P_i such that $\overline{G_i}$ does not contain v. Clearly, a vertex v is contained in $\overline{G_i}$ if and only if $last(v) > i$.

3 The (1,1)-edge-order

Although (1,1)-edge-orders can be seen as edge-counterparts of st-numberings, they do not seem to be well-known. Let two edges be *neighbors* if they share a common vertex. Annexstein et al. gave essentially the following definition.

▶ **Definition 4** ([1]). Let $G = (V, E)$ be a graph with an edge st that is not a self-loop. A *(1,1)-edge-order through st* of G is a total order $<$ on the edge set $E - st$ such that $m \geq 2$,
- every edge e, except for one incident to s, has a neighbor e' with $e' < e$ and
- every edge e, except for one incident to t, has a neighbor e' with $e < e'$.

Hence, the two exceptional edges incident to s and t must be, respectively, the minimal and maximal edge of $E - st$ with respect to $<$. Clearly, if G has a (1,1)-edge-order through st, G is 2-edge-connected, as neither st nor any other edge can be a bridge of G (note that this requires $m \geq 2$). The converse statement was shown in [1, Prop. 4] using a special type of ear decompositions based on breadth-first-search (however, without giving details of the linear-time algorithm). Here, we aim for a simple and direct (unlike, e.g., reducing to (1,1)-orders via line-graphs) exposition of the underlying idea and show that *any* ear decomposition can be transformed to a (1,1)-edge-order in linear time.

We will use the *incremental list order-maintenance problem*, which maintains a total order subject to the operations of (i) *inserting* an element after a given element and (ii) *comparing* two distinct given elements by returning the one that is smaller in the order. Bender et al. [3] show a simple solution for an even more general problem with amortized constant time per operation; we will call this the *order data structure*.

▶ **Lemma 5.** *Let G be a 2-edge-connected graph with an edge st that is not a self-loop. Then a (1,1)-edge-order through st can be computed in time $O(m)$.*

Proof. We compute an ear decomposition D of G such that $st \in P_0$. This can be done in linear time by any text-book-algorithm; see [24] for a simple one. Let $<_0$ be the total order that orders the edges in $P_0 - st$ consecutively from s to t. Thus, every edge has a smaller and a larger neighbor, except for st and the two exceptional edges incident to s and t. Clearly, $<_0$ is a (1,1)-edge-order through st of the 2-edge-connected graph G_0. We extend $<_{i-1}$ iteratively to a (1,1)-edge-order $<_i$ of G_i by adding the next ear P_i of D; then $<_{m-n}$ gives the claim.

The order itself is stored in the *order data structure*. For every vertex x in G_{i-1}, let $min(x)$ be the smaller of its two incident edges in $P_{birth(x)}$ with respect to $<_{i-1}$ (for later arguments, define $max(x)$ analogously as the larger such edge); clearly, $min(x)$ and $max(x)$ can be computed in constant time while adding P_j. When adding the ear P_i with (not necessarily distinct) endpoints x and y, let e be the smallest edge in $\{min(x), min(y)\}$ with respect to $<_{i-1}$ (this needs amortized constant time by using at most one comparison of the data structure). Consider all edges of P_i in consecutive order starting with a neighbor of e. We obtain $<_i$ from $<_{i-1}$ by inserting these edges as one consecutive block immediately after the edge e (if P_i is a cycle with endpoint s the edges are insert in front of the other edges); this takes amortized time proportional to the length of P_i. Then the first edge of P_i has a smaller neighbor in $<_i$ while the last has a larger neighbor in $<_i$ (for cycles $P_i \neq P_0$, this exploits that q_i has another incident edge in G_{i-1} or the exceptional edge incident to s (or t) might change), which implies that $<_i$ is a (1,1)-edge-order.   ◀

This (special) (1,1)-edge-order will allow for a very easy computation of two edge-independent spanning trees in Section 6 and serve as a building block for the computation of three such trees. If one wants to keep the root-paths in two edge-independent spanning trees short, a different (1,1)-edge-order [1] may be computed by maintaining $min(x)$ as the incident edge of x that is minimal in G_i in the above algorithm (this can be done efficiently by updating $min(x)$ whenever an ear with endpoint x is added). However, the latter order cannot be used for three edge-independent spanning trees.

4   The (2,1)-edge-order

We define (2,1)-orders as special ear decompositions.

▶ **Definition 6.** Let G be a graph with distinct edges rt and ru ($t = u$ is possible). A *(2,1)-edge-order* through rt and avoiding ru (see Figure 1) is an ear decomposition D of G such that

1. $rt \in P_0$,
2. $P_{m-n} = ru$, and   ▷ i.e., the last ear is the short ear ru
3. for every $0 \leq i < m - n$, $\overline{G_i}$ contains $inner(P_i)$ and, if P_i is short, at least one endpoint of P_i.

Property 6.2 implies that $\overline{G_i}$ contains the vertices r and u for every $0 \leq i < m - n$. We call Property 6.3 the *non-separateness* of D. The non-separateness of D states that every inner vertex of a long ear P_i has an incident edge in G that is in $\overline{G_i}$, and that every short ear P_i (seen as edge) has a neighbor in $\overline{G_i}$. The name refers to the following helpful property.

▶ **Lemma 7.** *Let D be a (2,1)-edge-order. Then, for every $0 \leq i < m - n$, $\overline{G_i}$ is connected.*

Proof. Consider any $i < m - n$ and let e be any edge in $\overline{G_i}$. By Property 6.2, $r \in \overline{G_i}$. We show that $\overline{G_i}$ contains a path from one of the endpoints of e to r. This gives the claim, as $\overline{G_i}$ is an edge-induced graph and therefore does not contain isolated vertices.

Figure 1 A (2,1)-edge-order of a 3-edge connected graph.

Let P_j be the unique ear that contains e. If P_j is short, $P_j = e$ and e has a neighbor in $\overline{G_j}$ due to the non-separateness of D. If P_j is long, at least one endpoint of e must be an inner vertex of P_j and e has a neighbor in $\overline{G_j}$ for the same reason. Hence, in both cases we find a neighbor that is contained in an ear P_k with $k > j$. By applying induction on the indices of these ears, we find a path that starts with an endpoint of e and ends with the only edge left in $\overline{G_{m-n-1}}$, namely ru. ◀

Next, we show that the existence of a (2,1)-edge-order proves the graph to be 3-edge-connected.

▶ **Lemma 8.** *If G has a (2,1)-edge-order, G is 3-edge-connected.*

Proof. Let D be a (2,1)-edge-order through rt and avoiding ru. Consider any vertex v of G. By transitivity of edge-connectivity, it suffices to show that G contains three edge-disjoint paths between v and r. Let P_i be the ear that contains v as inner vertex. In particular $i < m - n$, as P_i is long. Then G_i has an ear decomposition and, due to Theorem 2, contains two edge-disjoint paths between v and r. By Properties 6.2+3, $\overline{G_i}$ contains v and r. According to Lemma 7, $\overline{G_i}$ is connected. Thus, $\overline{G_i}$ contains a third path between v and r, which is edge-disjoint from the first two, as G_i and $\overline{G_i}$ are edge-disjoint. ◀

Let G have a (2,1)-edge-order. Then Lemma 8 implies $\delta(G) \geq 3$. This in turn gives that, for every vertex v, $P_{last(v)}$ is not the first ear that contains v, which implies that $P_{last(v)}$ must have v as endpoint. In particular, if vw is an edge and $last(v) = last(w) = birth(vw)$, $P_{birth(vw)}$ is the short ear vw and, according to the non-separateness of D, we have $i = m - n$, which implies $vw = ru$.

▶ **Lemma 9.** *For any vertex v, $P_{last(v)}$ has v as an endpoint. For any edge vw satisfying $last(v) = last(w) = birth(vw)$, $vw = ru$.*

The converse of Lemma 8 is also true: If G is 3-edge-connected, G has a (2,1)-edge-order. This gives a full characterization of 3-edge-connected graphs; however, proving the latter direction is more involved than Lemma 8. In the next section, we will prove the stronger statement that such a (2,1)-edge-order does not only exist but can actually be computed efficiently.

5   Computing a (2,1)-edge-order

At the heart of our algorithm is the following classical construction of 3-edge-connected graphs due to Mader.

(a) vertex-vertex-addition: $v = w$ is allowed. **(b)** edge-vertex-addition: $w \in \{a, b\}$ is allowed. **(c)** edge-edge-addition: $a, b \in \{c, d\}$ is allowed.

Figure 2 Mader-operations.

Figure 3 A (2,1)-edge-order of K_2^3 through rt and avoiding ru.

▶ **Definition 10.** The following operations on graphs are called *Mader-operations* (see Figure 2).
(a) *vertex-vertex-addition*: Add an edge between the not necessarily distinct vertices v and w (possibly a parallel edge or, if $v = w$, a self-loop).
(b) *edge-vertex-addition*: Subdivide an edge ab with a vertex v and add the edge vw for a vertex w.
(c) *edge-edge-addition*: Subdivide two distinct edges ab and cd with vertices v and w, respectively, and add the edge vw.

The edge vw is called the *added edge* of the Mader-operation. Let K_2^3 be the graph that consists of exactly two vertices and three parallel edges.

▶ **Theorem 11** ([16]). *A graph G is 3-edge-connected if and only if G can be constructed from K_2^3 using Mader-operations.*

According to Theorem 11, applying Mader-operations on 3-edge-connected graphs preserves 3-edge-connectivity. We will call a sequence of Mader-operations that constructs a 3-edge-connected graph a *Mader-sequence*. It has been shown that a Mader-sequence can be computed efficiently.

▶ **Theorem 12** ([18, Thm. 4]). *A Mader-sequence of a 3-edge-connected graph can be computed in time $O(n + m)$.*

Our algorithm for computing a (2,1)-edge-order works as follows. Assume we want a (2,1)-edge-order of G through $r\bar{t}$ and avoiding $r\bar{u}$. We first compute a suitable Mader-sequence of G using Theorem 12 and start with a (2,1)-edge-order of its first graph K_2^3. This (2,1)-edge-order is easy to find (see Figure 3). The crucial part of the algorithm is then to iteratively modify the given (2,1)-edge-order to a (2,1)-edge-order of the next graph in the sequence efficiently.

There are several technical difficulties to master. First, the edges $r\bar{t}$ and $r\bar{u}$ may be contained in different 2-connected components A' and B' (implying that r is a cut-vertex). As this would raise problems in the computation of the initial K_2^3 later, we perform in such a case the following reduction in advance. Let $\overline{A}$ be the connected component of $G \setminus \{r\}$

containing $\bar{t}$, $A := G[V(\overline{A}) \cup \{r\}]$ and $B := G \setminus V(\overline{A})$ (note that r may still be a cut-vertex of B). Since r is a cut-vertex of G, A and B are still 3-edge-connected. We compute a (2,1)-edge-order D_A of A avoiding $r\bar{t}$ through an arbitrary edge $ru_A \in A' \setminus \{r\bar{t}\}$, and a (2,1)-edge-order D_B of B avoiding an arbitrary edge $rt_B \in B' \setminus \{r\overline{u}\}$ through $r\overline{u}$. Then concatenating D_A with D_B gives a (2,1)-edge-order of G. Hence, we assume from now on that $r\bar{t}$ and $r\overline{u}$ are in the same 2-connected component of the input graph G.

Second, the edge $r\bar{t}$ (and analogously $r\overline{u}$) of G is not necessarily contained in the previous graph of the Mader-sequence, as it may have been created by a Mader-operation that subdivided a previous edge rt with the new vertex $\bar{t}$ (a more general view on this dynamics follows from the bijection between the graphs H of the Mader-sequence and H-subdivisions that are contained in G as subgraphs [18, Thm.+Cor. 1]; we refer to [23, Sections 2.3 and 4] for details of this bijection). In such cases, we take t as *replacement* vertex for $\bar{t}$ (and likewise u for $\overline{u}$) in the previous graph, and iterate this procedure to obtain replacement vertices for t and u in the graph before that previous graph, and so forth. This way, the replacement vertices t and u in any graph of the Mader-sequence containing r are neighbors of r.

Now a special Mader-sequence is used to harness the dynamics of the vertices r, t and u: Choose a DFS-tree of G with root r such that $r\bar{t}$ and $r\overline{u}$ are backedges (this is possible, since r has degree at least three) and compute a Mader-sequence of this DFS-tree that contains these two edges in its initial K_2^3 (this is possible, since $r\bar{t}$ and $r\overline{u}$ are in the same 2-connected component of G). This way the K_2^3 consists of the two vertices r and $t = u$ by the construction of [18, p. 6], and thus all graphs in the Mader-sequence contain r (and t and u are always neighbors of r). The vertices $\bar{t}$ and $\overline{u}$ are not present in this initial K_2^3 unless they are identical to $t = u$ (they are however contained in the two paths from r to $t = u$ of the K_2^3-subdivision the bijection maps to). For every graph in the Mader-sequence, we will compute a (2,1)-edge-order through rt and avoiding ru using the previous (2,1)-edge-order (which depends on the previous and possibly different replacement vertices); then the choice of t and u ensures that the final (2,1)-edge-order of G is indeed through $r\bar{t}$ and avoids $r\overline{u}$, as desired.

Thus, consider a graph G of the above Mader-sequence for which we know a (2,1)-edge-order D and let G' be the next graph in that sequence. Then G' is only one Mader-operation away and we aim for an efficient modification of D into a (2,1)-edge-order D' of G'. We will prove that there is always a modification that is local in the sense that the only ears that are modified are "near" the added edge of the Mader-operation.

▶ **Lemma 13.** *Let $D = (P_0, P_1, \ldots, P_{m-n})$ be a (2,1)-edge-order of a 3-edge-connected graph G through rt and avoiding ru for replacement vertices t and u. Let G' be obtained from G by applying one Mader-operation Γ and let t' and u' be the replacement vertices of G'. Then a (2,1)-edge-order D' of G' through rt' avoiding ru' can be computed from D using only constantly many amortized constant-time modifications.*

Lemma 13 is our main technical contribution and we split its proof into the following three sections. First, we introduce the operations *leg*, *belly* and *head* in order to combine several cases that can be handled similarly for the different types of Γ. Second, we show how to modify D to D' and, third, we discuss computational issues.

For all three sections, let vw be the added edge of Γ such that v subdivides the edge $ab \in E(G)$ and w subdivides $cd \in E(G)$ (if applicable). Thus, the vertex t' in G' is either t, v or w, and the vertex u' in G' is either u, v or w (hence, $t'r$ and ru' will never be self-loops). In all three sections, *birth* and *last* will always refer to D, unless stated otherwise.

Let $P_i \neq P_0$ be an ear with a given orientation and let x be a vertex in P_i. If P_i is a path, we define $P_i[, x]$ and $P_i[x,]$ as the *maximal subpaths* of P_i that end and start at x,

Figure 4 The result of operation *leg* (dashed lines), black vertices are in $G_{birth(ab)-1}$.

respectively; if P_i is a cycle, we take the same definition with the additional restriction that $P_i[,x]$ starts at q_i and $P_i[x,]$ ends at q_i. Occasionally, the orientation of P_i will not matter; if none is given, an arbitrary orientation can be taken. For paths A and B, let $A + B$ be the *concatenation* of A and B.

5.1 Legs, bellies and heads

While the operations *leg* and *belly* are inspired by the ones in [26], the operation *head* is new. All three operations will show for some special cases how D can be modified to a (2,1)-edge-order D'. A complete description for all cases (using these operations) will be given in the next section.

5.1.1 Legs

Let Γ be either an edge-vertex-addition such that $ab \neq ru$ and $last(w) < birth(ab)$ or an edge-edge-addition such that $ab \neq ru$ and $birth(cd) < birth(ab)$. If $P_{birth(ab)}$ is long, at least one of a and b is an inner vertex, say w.l.o.g. b. Otherwise, $P_{birth(ab)} = ab$ is short and, as D is non-separating, at least one of a and b, say w.l.o.g. b, has an incident edge in $\overline{G_{birth(ab)}}$ (note that this requires $ab \neq ru$). In both cases, orient $P_{birth(ab)}$ from a to b. The operation *leg* constructs D' from D by replacing the ear $P_{birth(ab)}$ of D by the two consecutive ears $P_{birth(ab)}[,a] + av + vw$ and $vb + P_{birth(ab)}[b,]$ in that order and, if Γ is an edge-edge-addition, additionally subdividing the edge cd in $P_{birth(cd)}$ with w (see Figure 4). Note that this definition is well-defined also for cycles $P_{birth(ab)}$, including self-loops.

We omit the proof that D' is a (2,1)-edge-order through rt' avoiding ru'.

5.1.2 Bellies

Let Γ be either an edge-vertex-addition such that $last(w) = birth(ab)$ and $w \notin \{a, b\}$ or an edge-edge-addition such that $birth(cd) = birth(ab)$ (note that $c, d \in \{a, b\}$ is allowed.) Consider the shortest path in $P_{birth(ab)}$ from an endpoint to one of the vertices $\{a, b\}$, say w.l.o.g. b, such that w is contained in this path. We orient $P_{birth(ab)}$ from a to b. $P_{birth(ab)}$ is a long ear with b as inner vertex. If Γ is an edge-edge-addition, one of the vertices $\{c, d\}$, say w.l.o.g. c, is contained in $P_{birth(ab)}[,w]$.

If $birth(ab) > 0$, the operation *belly* constructs D' from D by replacing the ear $P_{birth(ab)}$ of D by the two consecutive ears $P_{birth(ab)}[,a] + av + vw + P_{birth(ab)}[w,]$ and $vb + P_{birth(ab)}[b,w]$ in that order (if edge-vertex-addition) and by the two consecutive ears $P_{birth(ab)}[,a] + av + vw + wd + P_{birth(ab)}[d,]$ and $vb + P_{birth(ab)}[b,c] + cw$ (if edge-edge-addition), see Figure 5. Note that this definition is well-defined also if $P_{birth(ab)}$ is a cycle. If $birth(ab) = 0$, the vertices v and w cut P_0 in two distinct paths $P_{0,1}$ and $P_{0,2}$ having endpoints v and w. Let $P_{0,1}$ be the path containing r. Then the operation *belly* constructs D' from D by replacing the ear $P_{birth(ab)}$ of D by the two consecutive ears $P_{0,1} + vw$ and $P_{0,2}$ in this order. If $rt \in \{ab, cd\}$, then either $v = t'$ or $w = t'$, respectively.

We omit the proof that D' is a (2,1)-edge-order through rt' avoiding ru'.

Figure 5 The result of the operation *belly* (dashed lines).

Figure 6 The dashed lines show the result of the operation *head*.

5.1.3 Heads

Let Γ be an edge-vertex-addition such that $w \in \{a, b\}$, $last(a) = birth(ab)$ and, if $ab = ru$, then $r \neq a$. W.l.o.g. let $w = a$. Then a is an endpoint of $P_{birth(ab)}$ ($P_{birth(ab)}$ cannot be a self-loop, as $last(a) = birth(ab)$). We orient $P_{birth(ab)}$ from a to b. The operation *head* constructs D' from D by replacing the ear $P_{birth(ab)}$ of D by the two consecutive ears $av + va$ and $vb + P_{birth(ab)}[b,]$ in that order (see Figure 6). Note that this definition is well-defined also for cycles $P_{birth(ab)}$.

We omit the proof that D' is a (2,1)-edge-order through rt' avoiding ru'.

5.2 Modifying D to D'

We will now show how to obtain a (2,1)-edge-order D' through rt' avoiding ru' from D. By symmetry, assume w.l.o.g. that $birth(ab) \geq birth(cd)$. Note that applying the operations *belly*, *leg* and *head* preserves all properties of a $(2,1)$-edge-order. Recall that, for every subdivision the Mader-sequences does on rt or ru, respectively, the subdividing vertex is t' or u', as explained after Figure 3. We have the following case distinctions:

1. **Γ is a vertex-vertex-addition.** (See Figure 2a.)
   (a) vw is a self-loop at v ($v = w$): Obtain D' from D by adding the new short ear vv directly after the ear $P_{last(v)-1}$. This ensures that the new ear is non-separating.
   (b) $v \neq w$ and $vw \neq \{rt, ru\}$: If $last(v) \leq last(w)$, D' is obtained from D by adding the new short ear vw directly after the ear $P_{last(w)-1}$, ensuring that the new ear is non-separating. If $last(v) > last(w)$, the new short ear vw is added directly after the ear $P_{last(v)-1}$.
   (c) $vw = rt$ (the added edge is a parallel edge): the Mader-sequence gives us the information whether rt is rt' or the new added edge is rt'. If $rt = rt'$ then add the new edge immediately after the ear $P_{last(t)-1}$. Otherwise obtain D' from D by replacing rt with rt' in P_0 and adding the old edge rt as an short ear immediately after the ear $P_{last(t)-1}$.
   (d) $vw = ru$ (the added edge is a parallel edge): the Mader-sequence gives us the information whether ru is ru' or the new added edge is ru'. Depending on this information, obtain D' from D by either adding the new edge directly before or directly after the last ear of D.
5. **Γ is an edge-vertex-addition.** (See Figure 2b.)
   (a) $birth(ab) < last(w)$: Obtain D' from D by adding the new short ear vw directly after the ear $P_{last(w)-1}$ and subdivide the ear $P_{birth(ab)}$ with v. This operation is also well-defined when $P_{birth(ab)}$ is a cycle or self-loop. Also, the new ear is non-separating and, since v is incident to w, the ear $P_{birth(ab)}$ remains non-separating.

    (b) $last(w) < birth(ab)$ and $ab \neq ru$: Apply *leg*

    (c) $birth(ab) = last(w)$ and $w \notin \{a, b\}$: Apply *belly*.

    (d) $birth(ab) = last(w)$ and $w \in \{a, b\}$; if $ab = ru$, then $r \neq w$: Apply *head*.

    (e) $ab = ru$ and if $birth(ab) = last(w)$ and $w \in \{a, b\}$ then $r = w$: Obtain D' from D by replacing the ear ru by the two consecutive ears $wv + vu$ and rv.

6. **Γ is an edge-edge-addition.** (See Figure 2c.)

    (a) $birth(ab) = birth(cd)$: Apply *belly*.

    (b) $birth(ab) > birth(cd)$ and $ab \neq ru$: Apply *leg*.

    (c) $ab = ru$: Let w.l.o.g. $r = a$. Obtain D' from D by replacing the last ear of D by the two consecutive ears $bv + vw$ and rv in this order.

In all cases, D' is clearly an ear decomposition. Properties 6.1–3 are satisfied due to the given case distinction and the mentioned properties. Hence, D' is a $(2,1)$-edge-order through rt' avoiding ru'.

There are several subtleties in sorting out the computational complexity of this approach, mostly raised by the question how fast we can compute one of the above cases in which we are in. The proof of the linear runtime is omitted due to space constraints.

▶ **Theorem 14.** *Given edges tr and ru of a 3-edge-connected graph G, a (2,1)-edge-order D of G through tr and avoiding ru can be computed in time $O(m)$.*

The proposed algorithms for $(1,1)$-edge-orders and $(2,1)$-edge-orders (as well as the computation of edge-independent spanning trees in the next section) are *certifying* in the sense of [17]: For $(1,1)$-edge-orders through st, it suffices to check that every edge $e \neq st$ has indeed a smaller and larger neighboring edge. For $(2,1)$-edge-orders, it suffices to check in linear time that D is an ear decomposition of G and that D satisfies Properties 6.1–3.
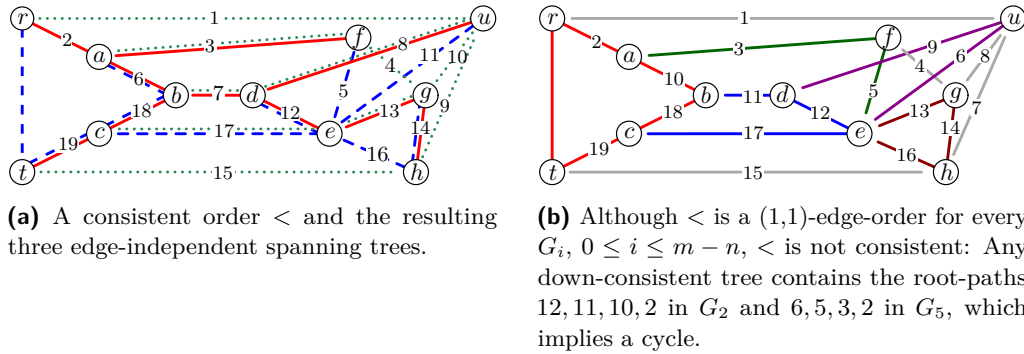
6   Edge-Independent Spanning Trees

Let k spanning trees of a graph be *edge-independent* if they all have the same root vertex r and, for every vertex $x \neq r$, the paths from x to r in the k spanning trees are edge disjoint. The following conjecture was stated 1988 by Itai and Rodeh.

▶ **Conjecture** (Edge-Independent Spanning Tree Conjecture [13]). *Every k-edge-connected graph contains k edge-independent spanning trees.*

The conjecture has been proven constructively for $k \leq 2$ [13] and $k = 3$ [11] with running times $O(m)$ and $O(n^2)$, respectively, for computing the corresponding edge-independent spanning trees. For every $k \geq 4$, the conjecture is open. We first give a short description of an algorithm for $k = 2$ and then show the first linear-time algorithm for $k = 3$.

For $k = 2$, compute the $(1,1)$-edge-order $<$ through tr using Lemma 5. The first tree T_1 consists of the edges $min(x)$ for all vertices $x \neq r$ (as defined in Lemma 5), while the second tree T_2 consists of tr and the edges $max(x)$ for all vertices $x \notin \{r, t\}$. Then T_1 and T_2 are spanning, as no edge can be taken twice, and edge-independent, as, from every vertex x, the path of smaller edges to r obtained by iteratively applying $min()$ must be edge-disjoint from the path of larger edges to r.

For $k = 3$, choose any vertex r and two distinct edges tr and ru in the 3-edge-connected graph G. Compute a $(2,1)$-edge-order D through tr and avoiding ru in time $O(m)$ using Theorem 14. For every vertex $x \in V$, the idea is now to find *two* edge-disjoint paths from x to r in $G_{birth(x)}$ (after all, $G_{birth(x)}$ is 2-edge-connected and thus contains a $(1,1)$-edge-order)

(a) A consistent order $<$ and the resulting three edge-independent spanning trees.

(b) Although $<$ is a (1,1)-edge-order for every G_i, $0 \le i \le m - n$, $<$ is not consistent: Any down-consistent tree contains the root-paths $12, 11, 10, 2$ in G_2 and $6, 5, 3, 2$ in G_5, which implies a cycle.

Figure 7 (1,1)-edge-orders that are consistent and not consistent to the (2,1)-edge-order of Figure 1.

and a *third* path from x to r in $\overline{G_{birth(x)}}$ using the non-separateness of D. The subtle part is to make this idea precise: We have to construct the first tree T_1 in such a consistent way that the paths of smaller edges from x to r for *all* vertices $x \in V$ are contained in T_1 (and the same for T_2 and paths of larger edges).

For a (1,1)-edge-order $<$ through tr of G, let a spanning tree $T_1 \subseteq G$ be *down-consistent* to a given (2,1)-edge-order through tr if (a) every path in T_1 to r is strictly decreasing in $<$ and (b) for every $0 \le i \le m - n$, $T_1 \cap G_i$ is a spanning tree of G_i (analogously, *up-consistent* spanning trees T_2 of $G - r$ are defined by strictly increasing paths to t). Now let a (1,1)-edge-order be *consistent* to a given (2,1)-edge-order D' if G contains r-rooted spanning trees T_1 and T_2 that are down- and up-consistent to D', respectively. By the very same argument as used for $k = 2$, T_1 and $T_2 + tr$ are edge-independent and, in addition, do not use any edge of $\overline{G_{birth(x)}}$ for any $x \in V$.

In fact, the special (1,1)-edge-order that is computed by Lemma 5 is consistent to D: There, the trees T_1 and T_2 consist of the edges $min(x)$ and $max(x)$ for $x \in V$, which makes T_1 down-consistent and $T_2 + tr$ up-consistent to D (see Figure 7a). We note that a simpler definition of consistent as used for the vertex-variant [6], i.e., as *orders that remain (1,1)-edge-orders for all subgraphs G_i, $0 \le i \le m - n$*, does not suffice here (see Figure 7b).

It remains to construct the third edge-independent spanning tree. For every edge $e \ne ru$ of G, we compute a pointer to an arbitrary neighboring edge e' in $\overline{G_{birth(e)}}$. This edge e' exists, as D is non-separating, and satisfies $birth(e') > birth(e)$. Similarly, for every vertex $x \in V - r - u$, we compute a pointer to an incident edge e' of x with $birth(e') > birth(x)$. Both computations take linear total time by comparing $birth$ values. The third edge-independent spanning tree is then the union of ur and the u-rooted spanning tree of $G - r$ that interprets the pointers as parent edges. Hence, we obtain the following theorem.

▶ **Theorem 15.** *Given the two edges rt and ru of a 3-edge-connected graph G, three edge-independent spanning trees of G rooted at r (such that two of them contain rt and ru as unique root edges, respectively) can be computed in time $O(m)$.*

Similarly as for the more general (2,1)-edge-orders, one could be interested why the reduction from k-edge- to k-vertex-connectivity by Galil and Italiano [10] does not give edge-independent spanning trees from their vertex-counterparts; we omit the argument due to space constraints.

References

1 Fred Annexstein, Ken Berman, and Ram Swaminathan. Independent spanning trees with small stretch factors. Technical Report 96-13, DIMACS, June 1996.

2 M. Badent, U. Brandes, and S. Cornelsen. More canonical ordering. *Journal of Graph Algorithms and Applications*, 15(1):97–126, 2011.

3 M. A. Bender, R. Cole, E. D. Demaine, M. Farach-Colton, and J. Zito. Two simplified algorithms for maintaining order in a list. In *Proceedings of the 10th European Symposium on Algorithms (ESA'02)*, pages 152–164, 2002. `doi:10.1007/3-540-45749-6_17`.

4 T. Biedl and M. Derka. The (3,1)-ordering for 4-connected planar triangulations. *Journal of Graph Algorithms and Applications*, 20(2):347–362, 2016.

5 T. Biedl and J. M. Schmidt. Small-area orthogonal drawings of 3-connected graphs. In *Proceedings of the 23rd International Symposium on Graph Drawing (GD'15)*, pages 153–165, 2015. `doi:10.1007/978-3-319-27261-0_13`.

6 J. Cheriyan and S. N. Maheshwari. Finding nonseparating induced cycles and independent spanning trees in 3-connected graphs. *Journal of Algorithms*, 9(4):507–537, 1988.

7 S. Curran, O. Lee, and X. Yu. Chain decompositions of 4-connected graphs. *SIAM J. Discrete Math.*, 19(4):848–880, 2005. `doi:10.1137/S0895480103434592`.

8 H. de Fraysseix, J. Pach, and R. Pollack. Small sets supporting fary embeddings of planar graphs. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC'88)*, pages 426–433, 1988. `doi:10.1145/62212.62254`.

9 S. Even and R. E. Tarjan. Computing an st-Numbering. *Theor. Comput. Sci.*, 2(3):339–344, 1976.

10 Z. Galil and G. F. Italiano. Reducing edge connectivity to vertex connectivity. *SIGACT News*, 22(1):57–61, 1991. `doi:10.1145/122413.122416`.

11 A. Gopalan and S. Ramasubramanian. On constructing three edge independent spanning trees. Manuscript (see `http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.406.7119`), March 2011.

12 A. Gopalan and S. Ramasubramanian. A counterexample for the proof of implication conjecture on independent spanning trees. *Information Processing Letters*, 113(14-16):522–526, 2013. `doi:10.1016/j.ipl.2013.04.008`.

13 A. Itai and M. Rodeh. The multi-tree approach to reliability in distributed networks. *Information and Computation*, 79:43–59, 1988.

14 G. Kant. Drawing planar graphs using the lmc-ordering. In *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science (FOCS'92)*, pages 101–110, 1992. `doi:10.1109/SFCS.1992.267814`.

15 L. Lovász. Computing ears and branchings in parallel. In *Proceedings of the 26th Annual Symposium on Foundations of Computer Science (FOCS'85)*, pages 464–467, 1985. `doi:10.1109/SFCS.1985.16`.

16 W. Mader. A reduction method for edge-connectivity in graphs. In B. Bollobás, editor, *Advances in Graph Theory*, volume 3 of *Annals of Discrete Mathematics*, pages 145–164. North-Holland, 1978. `doi:10.1016/S0167-5060(08)70504-1`.

17 R. M. McConnell, K. Mehlhorn, S. Näher, and P. Schweitzer. Certifying algorithms. *Computer Science Review*, 5(2):119–161, 2011. `doi:DOI:10.1016/j.cosrev.2010.09.009`.

18 K. Mehlhorn, A. Neumann, and J. M. Schmidt. Certifying 3-edge-connectivity. *Algorithmica*, 77(2):309–335, 2017. `doi:10.1007/s00453-015-0075-x`.

19 L. F. Mondshein. *Combinatorial Ordering and the Geometric Embedding of Graphs*. PhD thesis, M.I.T. Lincoln Laboratory / Harvard University, 1971. Technical Report available at `http://www.dtic.mil/cgi-bin/GetTRDoc?AD=AD0732882`.

20   S. Nagai and S. Nakano. A linear-time algorithm to find independent spanning trees in maximal planar graphs. In *26th International Workshop on Graph-Theoretic Concepts in Computer Science (WG'00)*, pages 290–301, 2000. `doi:10.1007/3-540-40064-8_27`.

21   S. Nakano, Md. S. Rahman, and T. Nishizeki. A linear-time algorithm for four-partitioning four-connected planar graphs. *Inf. Process. Lett.*, 62(6):315–322, 1997. `doi:10.1016/S0020-0190(97)00083-5`.

22   H. E. Robbins. A theorem on graphs, with an application to a problem of traffic control. *The American Mathematical Monthly*, 46(5):281–283, 1939. URL: `http://www.jstor.org/stable/2303897`.

23   J. M. Schmidt. Construction sequences and certifying 3-connectedness. In *Proceedings of the 27th Symposium on Theoretical Aspects of Computer Science (STACS'10)*, pages 633–644, 2010. URL: `http://arxiv.org/abs/0912.2561`, `doi:10.4230/LIPIcs.STACS.2010.2491`.

24   J. M. Schmidt. A simple test on 2-vertex- and 2-edge-connectivity. *Information Processing Letters*, 113(7):241–244, 2013. `doi:10.1016/j.ipl.2013.01.016`.

25   J. M. Schmidt. The Mondshein sequence. In *Proceedings of the 41st International Colloquium on Automata, Languages and Programming (ICALP'14)*, pages 967–978, 2014. `doi:10.1007/978-3-662-43948-7_80`.

26   J. M. Schmidt. Mondshein sequences (a.k.a. (2,1)-orders). *SIAM Journal on Computing*, 45(6):1985–2003, 2016. `doi:10.1137/15M1030030`.

27   K. Wada and K. Kawaguchi. Efficient algorithms for tripartitioning triconnected graphs and 3-edge-connected graphs. In *19th International Workshop on Graph-Theoretic Concepts in Computer Science (WG'93)*, pages 132–143, 1993. `doi:10.1007/3-540-57899-4_47`.

28   K. Wada, A. Takaki, and K. Kawaguchi. Efficient algorithms for a mixed k-partition problem of graphs without specifying bases. *Theoretical Computer Science*, 201:233–248, 1998.

29   H. Whitney. Non-separable and planar graphs. *Transactions of the American Mathematical Society*, 34(1):339–362, 1932.