DIMACS Technical Report 96-13
July 1996

Independent Spanning Trees with Small Stretch Factors

by

Fred Annexstein
Department of ECECS
University of Cincinnati
Cincinnati, OH 45221
fred.annexstein@uc.edu

Ken Berman
Department of ECECS
University of Cincinnati
Cincinnati, OH 45221
ken.berman@uc.edu

Ram Swaminathan[1]
RUTCOR and DIMACS
Rutgers University
New Brunswick, NJ 08903
manian@dimacs.rutgers.edu

[1]Support rendered by RUTCOR/DIMACS Center is acknowledged.

DIMACS is a cooperative project of Rutgers University, Princeton University, AT&T Bell Laboratories and Bellcore.

DIMACS is an NSF Science and Technology Center, funded under contract STC–91–19999; and also receives support from the New Jersey Commission on Science and Technology.

ABSTRACT

A pair of spanning trees rooted at a vertex r are independent if for every vertex v the pair of unique tree paths from v to the root r are disjoint. This paper presents the first analysis of the path lengths involved in independent spanning trees in 2-edge-connected and 2-vertex-connected graphs. We present upper and lower bounds on the stretch factors of pairs of independent spanning trees, where the stretch factor of a spanning tree is defined to be the maximum ratio between the length of paths in the tree to the root to the length of the shortest path in the graph to the root. We prove that if G is a 2-edge-connected graph with the property that every edge lies on a cycle of size at most h than we can construct in linear time a pair of edge-independent spanning trees whose stretch factors are bounded by $O(h)$. In fact, we prove a more general result, namely that the stretch factor of both independent trees can be bounded by a minimax length of ears with respect to a certain class of ear decompositions of the graph. We demonstrate analogous constructions of vertex-independent spanning trees with bounded stretch factors for the class of 2-vertex-connected graphs. We show that our upper bounds are existentially optimal, that is, there are classes of graphs for which our bounds are tight.

1 Introduction

A pair of spanning trees rooted at a vertex r are *independent* if for every vertex v the pair of unique tree paths from v to the root r are (vertex or edge) disjoint. The quality of a particular choice of independent spanning trees can be measured by comparing the length of paths in the trees to that of the length of shortest paths in the graph. For a given graph G our interest is in the optimization problem that identifies multiple independent spanning trees rooted at a vertex for which the *stretch factors* of the individual trees are simultaneously minimized. The stretch factor of a spanning tree is defined to be the maximum ratio between the length of the path in the tree to the root to that of the length of the shortest path in the graph to the root vertex.

The problem addressed in this paper is motivated by the fact that it is convenient to restrict communication in a distributed system to a set of links defined by a spanning tree of the underlying network, particularly when the complexity of the routing scheme is a concern [1, 8, 4]. The use of multiple independent-spanning trees for data routing in networks provide a number of advantages, including increased fault-tolerance and bandwidth, when the connectivity of the underlying network permits such trees to be defined [5]. Hence in the design of routing schemes for data networks, constructing multiple independent spanning trees with small stretch factors is an important issue.

1.1 Previous Work

A number of papers have previously considered the existence and construction of independent spanning trees in graphs; see e.g., [10, 3, 6]. Itai and Rodeh [5] introduced the notion of a 2-tree protocol (for edges and vertices) for distributed networks. These are communication protocols that are resilient to a single (edge or vertex) failure. The idea is that in a 2-edge-connected graph, for example, one can identify a pair of edge-independent spanning trees. By routing on both of these trees simultaneously, no single edge failure can prevent a message from reaching its intended destination. However, the work in [5] does not analyze the lengths of the paths in these spanning trees. For certain classes of graphs, for example n-*wheel* graphs, there exist edge-independent spanning trees with stretch factors of $O(1)$; however, the construction given in [5] yields stretch factors that are $\Omega(n)$, the worst possible.

Stretch factors for routing schemes have been considered before in the work of [1]. But, to our knowledge, this paper is the first study of stretch factors for independent spanning trees.

1.2 Our Results

The main contributions of this paper are linear-time algorithms for the construction of edge-independent and vertex-independent spanning trees with provable bounds on the stretch factors. Our analysis involves a new unified treatment of 2-edge-connected and 2-vertex-connected graphs, based upon a common framework of ear decompositions. We show how such decompositions induce natural orderings on the edges and vertices. These orderings allow us to construct edge-independent spanning trees with stretch factors bounded by $O(h)$, and vertex-independent spanning trees with stretch factors bounded by $O(\delta h)$, where h is the minimax value of certain ear decomposition of the graph, and δ is the maximum degree of a vertex in the graphs. More precisely, h is the minimum over all breadth-first ear decompositions of the maximum size of the largest ear in the decomposition; below we define these terms precisely. We show that in certain cases our bounds can be translated in terms of more commonly used structural graph invariants. For example, our results imply that if G is a 2-edge-connected graph with the property that every edge lies on a cycle of size at most h, then we can construct in linear time a pair of edge-independent spanning trees whose stretch factors are bounded by $O(h)$. An analogous result holds for constructing vertex-independent spanning trees in 2-vertex-connected graphs.

2 Preliminaries

A connected graph is 2-*edge-connected* (respectively, 2-*vertex-connected*) if it has no cut edge (respectively, cut vertex). For undefined graph theoretic terminology we refer the reader to the standard reference by Bondy and Murty [2].

Let T be a spanning tree of G rooted at a vertex $r \in V(G)$. For every vertex $v \in V(G)$, let $dist_T(v,r)$ denote the distance from v to r in T. Similarly, define $dist_G(v,r)$. Then, the *stretch factor* $sf(T)$ of T is:

$$sf(T) = \max \{dist_T(v,r) \ / \ dist_G(v,r) \mid v \in V\}.$$

Our goal in this paper is to simultaneously bound the stretch factors of a pair of independent spanning trees. A pair of spanning trees T_1 and T_2 of G, both rooted at r, are *edge-independent* (respectively, *vertex-independent*) if for every $v \in V(G)$, the paths from v to r in T_1 and T_2 are edge-disjoint (respectively, internally vertex-disjoint).

2.1 Constructing Independent Spanning Trees

Our constructions of independent spanning trees are based on the induced paths given by certain orderings of edges and vertices in the graphs. We now define these orderings precisely.

Definition 1. *Let G be a 2-vertex-connected graph, and let s,t be distinct elements of $V(G)$. An st-*vertex-ordering *of G is a total ordering $<$ of $V(G)$ such that*

1) *s is the minimum and t is the maximum in the ordering, and*
2) *for every vertex $v \in V - \{s,t\}$ there exists a pair of vertices u,w adjacent to v such that $u < v < w$.*

Next we define the edge-ordering which is slightly different from the vertex-ordering.

Definition 2. *Let G be a 2-edge-connected graph, and let $st \in E(G)$. An st-*edge-ordering *of G is a total ordering $<$ of $(E(G) - \{st\})$ such that*

1) *for every edge $e \in E(G) - \{st\}$, either e is incident to s, or there exists another edge d adjacent to e such that $d < e$, and*

2) *for every edge $e \in E(G) - \{st\}$, either e is incident to t, or there exists another edge f adjacent to e such that $e < f$.* (Notice that d and f need not be incident to distinct ends of e.)

Following propositions, which are easily proved using induction on the number of ears, relate st-orderings with the connectivity of the graph.

Proposition 1. ([7]) *Every 2-vertex-connected graph G has an st-vertex-ordering, for any pair of vertices $s, t \in V(G)$.* □

Though the proposition below can also be proved using induction on the number of ears, we give a different proof here relating st-vertex-ordering and st-edge-ordering.

Proposition 2. *Every 2-edge-connected graph G has an st-edge-ordering, for any edge $st \in E(G)$.*

Proof. Let G be a 2-edge-connected graph with st in $E(G)$. Define the line graph $L(G)$ of G as follows. The vertex-set of $L(G)$ is $E(G)$ and two vertices are adjacent in $L(G)$ if and only if the respective edges are adjacent in G. Since G is 2-vertex-connected, clearly $L(G)$ is 2-vertex-connected. Pick any edge (other than st) that is incident to s in G and label the corresponding vertex in $L(G)$ with s. Similarly, pick any edge (other than st) that is incident to t in G and label the corresponding vertex in $L(G)$ with t. It follows from Proposition 1 that $L(G)$ has an st-vertex-ordering. This vertex-ordering of $L(G)$ when applied to the edges of G satisfies the definition of an st-edge-ordering of G. □

2.1.1 Vertex-independent trees

Let s be a distinguished vertex, and let t be a vertex adjacent to s in a graph G. Given an st-vertex-ordering in G, we can construct vertex-independent spanning trees, called the *trees induced by the st-vertex-ordering*, as follows. An edge of a graph is *st-oriented* by viewing the edge as a pair of opposing arcs and labeling one of the arcs s and the other t. Given an st-vertex-ordering in a graph G, the *induced* st-orientation of each edge will be to label each descending directed arc with s and each ascending directed arc with t. The set of all s-labeled arcs forms a directed graph with a single sink s, and the set of all t-labeled arcs forms a directed graph with a single sink t. A pair of spanning arborescences (i.e., spanning in-trees, one rooted at s and one at t) can be obtained by restricting each non-sink vertex to have out-degree one. Since all paths to s in the s-labeled spanning tree are descending while all paths to t in the t-labeled spanning tree ascending (with respect to the st-vertex-ordering), it follows that the pair of underlying (undirected) spanning trees (rooted at s) are vertex-independent. Note that we may have to convert the t-labeled spanning tree to a spanning tree rooted at s by replacing the t-labeled edge incident to s with the edge st.

2.1.2 Edge-independent trees

Let s be a distinguished vertex, and let t be a vertex adjacent to s in a graph G. Given an st-edge-ordering in G, we can construct edge-independent spanning trees, called the *trees induced by the st-edge-ordering*. Consider a vertex v, and suppose that the minimum edge incident to v is vu. We then s-label the arc vu, except when the minimum edge incident to u is also uv. In this case it follows, from the definition of st-edge ordering, that one of u or v is the vertex s, in which case we s-label the arc towards s. Likewise, suppose that the maximum edge incident to v is vw. We then t-label the arc vw, except when the maximum edge incident to w is also vw. In this case it follows, that one of u or v is the vertex t, in which case we t-label the arc towards t. Again, it follows that we can identify a pair of spanning arborescences rooted at the vertex s (one s-labeled and one t-labeled) where the underlying (undirected) spanning trees are edge independent, since all paths in the s-oriented spanning tree are descending while all paths in the t-oriented spanning tree are ascending with respect to the st-edge ordering. As in the vertex-case, here too we may have to convert the t-labeled spanning tree to a spanning tree rooted at s by replacing the t-labeled edge incident to s with the edge st.

2.2 Ear decompositions

We will use a particular method to construct st-orderings of graphs that will induce trees with small stretch factors. The method is based on a certain ear decomposition of the graph. An *ear decomposition* $\mathcal{D} = [P_0, P_1, \ldots, P_t]$ of an undirected graph G is a partition of edge set $E(G)$ into an ordered collection of edge-disjoint simple paths $P_0, P_1, \ldots, P_{t-1}$ such that P_0 is a single edge, $P_0 \cup P_1$ is a simple cycle, and each end vertex of P_i, for $i > 1$, is contained in some P_j, $j < i$ and none of the other vertices of P_i are contained in any P_j, $j < i$. The paths P_i are called *ears*. An ear is *trivial* if it is a single edge. The *size* or *length* of an ear is the number of vertices it contains, excluding the ends. The end vertices of the ear are called the *attachments*, and all other vertices are called *internal*. An ear is *open* if it has distinct attachment vertices; and when all ears are open we call $\mathcal{D}$ an *open-ear decomposition*.

An open ear decomposition $[P_0, P_1, \ldots, P_t]$ with $P_0 = st$ naturally induces an st-vertex-ordering by inductively ordering the internal vertices of each ear P_i immediately succeeding the minimum attachment vertex of P_i. Similarly, an open ear decomposition $[P_0, P_1, \ldots, P_t]$ with $P_0 = st$ naturally induces an st-edge-ordering. Hence (via transitivity) we say that the edges of ear P_i have an induced st-orientation, that is the orientation induced by the st-ordering which in turn is induced by the ear decomposition.

Suppose we are given a graph with an (open) ear decomposition $[P_0, P_1, \ldots, P_t]$ with $P_0 = st$. The following two propositions point out that the edges of ear P_i have a naturally induced st-orientation. This orientation of edges is induced by the natural st-ordering induced by the ear decomposition.

Proposition 3. *An open ear decomposition $[P_0, P_1, \ldots, P_t]$ with $P_0 = st$ naturally induces an st-vertex-ordering by inductively ordering the internal vertices of each ear P_i immediately* *succeeding the minimum attachment vertex of* P_i. □

Proposition 4. *An ear decomposition* $[P_0, P_1, \ldots, P_t]$ *with* $P_0 = st$ *naturally induces an* st*-edge-ordering by inductively ordering the edges of each ear* P_i *immediately succeeding the minimum edge incident to one of the attachment vertices of* P_i. □

To construct independent spanning trees with small stretch factors we look at a specific kind of ear decomposition of G. Given a distinguished vertex r, a *breadth-first ear decomposition rooted at* r of G, is an ear decomposition $\mathcal{D} = [P_0, P_1, \ldots, P_t]$ such that P_0 is an edge incident to r, and each non-trivial ear P_i, for $i > 0$, contains an internal vertex v of minimum distance in G to r over all vertices of G not in previously generated ears $P_0, P_1, \ldots, P_{i-1}$, and an edge on some shortest path from r to v. This implies that the other vertex incident to this shortest-path edge is an attachment vertex for the ear P_i. If this minimum distance of v to r is say k, then we call ear P_i a *level-*k *ear*. Let $\mathcal{D}[k]$ denote the set of all level-k ears in the decomposition $\mathcal{D}$. It follows from the definition that each ear in $\mathcal{D}[k]$ has at least one attachment vertex that is a member of an ear of $\mathcal{D}[l]$, for some $l < k$, and another attachment vertex that is a member of an ear of $\mathcal{D}[l']$, for some $l' \leq k$.

3 Bounding the Stretch Factors

The following proposition shows that it is not always possible to simultaneously bound (non-trivially) the stretch factors of both trees in a pair of vertex-independent spanning trees. An n*-wheel* is defined as a graph consisting of an n-vertex cycle and a distinct *hub* vertex that is connected to every vertex on the cycle.

Proposition 5. *Let* r *be any non-hub vertex of an* n*-wheel. Then for any pair of vertex-independent spanning trees rooted at* r*, one of the pair must have a stretch factor of* $\Omega(n)$.

Proof. Consider the vertex v on the n-cycle diametrically opposite to the root. Any path in T from v to root that does not use the hub must have length at least $\lfloor n/2 \rfloor$. Since two vertex-disjoint tree paths from v to r cannot both use the hub, the proposition follows. □

In [5] a method is given for constructing a pair of edge-independent spanning trees in a 2-edge-connected graph G based on concatenating the vertex-independent spanning trees of the 2-vertex-connected components of G. The previous proposition shows that in n-wheel graphs, trees defined in this way have stretch factors that are $\Omega(n)$. We show next that in fact we can do much better, namely, in linear time we can construct a pair of edge-independent spanning trees in a 2-edge-connected graph G such that the stretch factors of both trees are proportional to the maximum length of an ear in any given breadth-first ear decomposition for G. Interestingly, for n-wheel graphs this bound becomes $O(1)$.

We now state and prove the first main result of the paper.

Theorem 1. *Let G be a 2-edge-connected graph with a distinguished root vertex r, and let $\mathcal{D}_1$ be any given breadth-first ear decomposition of G rooted at r. Then there exists a pair of edge-independent spanning trees T_1 and T_2 rooted at r so that $sf(T_1) \leq sf(T_2) \leq 3h$, where h is size of the largest ear in $\mathcal{D}_1$. Moreover, by modifying $\mathcal{D}_1$ we can produce in linear time a new ear decomposition $\mathcal{D}_2$ that induces the pair T_1 and T_2.*

Proof. The idea for the proof is as follows. We take $\mathcal{D}_1$ and modify it to produce a new (not necessarily breadth-first) ear decomposition $\mathcal{D}_2$ that induces the pair T_1 and T_2 with stretch factors satisfying the theorem. In particular, each ear of $\mathcal{D}_2$ will be constructed from the edge-sets of at most two ears of $\mathcal{D}_1$, and the edge-set of each ear of $\mathcal{D}_1$ will be covered by at most two ears of $\mathcal{D}_2$.

Let us first give some necessary definitions. Recall that for each k, $\mathcal{D}_1[k]$ denotes the set of all level-k ears of $\mathcal{D}_1$. We will construct a set of ears denoted by $\mathcal{D}_2[k]$, called the level-k ears of $\mathcal{D}_2$, whose union is identical as a subgraph to the union of the ears in $\mathcal{D}_1[k]$. For each k, let G_k denote the subgraph of G given by the union of all ears in $\mathcal{D}_1[i]$, for all $1 \leq i \leq k$. Each ear defined in $\mathcal{D}_2[k]$ will be one of three possible types: either an *anchor ear*, a *tether ear*, or a *bridge ear*.

An *anchor ear* Q_a in $\mathcal{D}_2[k]$ has the following properties:

- $a1)$ both attachment vertices of Q_a (not necessarily distinct) are elements of G_{k-1}; and
- $a2)$ Q_a is either a single ear of $\mathcal{D}_1[k]$, or the concatenation of subpaths of at most two ears of $\mathcal{D}_1[k]$.

A *tether ear* Q_t in $\mathcal{D}_2[k]$ has the following properties:

- $t1)$ one attachment vertex of Q_t is an element G_{k-1}, and the other attachment of Q_t is an internal vertex of an anchor ear of $\mathcal{D}_2[k]$;
- $t2)$ each internal vertex of Q_t has degree 2 in the subgraph G_k; and
- $t3)$ Q_t is a single ear of $\mathcal{D}_1[k]$.

A *bridge ear* Q_b in $\mathcal{D}_2[k]$ has the following properties:

- $b1)$ both attachment vertices of Q_b are internal vertices of anchor ears of $\mathcal{D}_2[k]$;
- $b2)$ each internal vertex of Q_b has degree 2 in the subgraph G_k; and
- $b3)$ Q_b is a subpath of a single ear of $\mathcal{D}_1[k]$.

For the inductive hypothesis, assume for each level $1 \leq i \leq k-1$, the subgraph given by the set of ears in $\mathcal{D}_1[i]$ can be decomposed into a set of ears $\mathcal{D}_2[i]$, such that each ear is either an anchor, tether, or bridge ear. This implies that in the induced edge-independent trees (defined for the subgraph G_{i-1}) the stretch factor is at most $3h$, where h is the size of the largest ear in any $\mathcal{D}_1[i]$. This follows since anchor ears defined in $\mathcal{D}_2[i]$ induce tree paths to vertices in G_{i-1} of length at most $2h$, and tether and bridge ears in $\mathcal{D}_2[i]$ induce tree paths to vertices in G_{i-1} of length at most $3h$.

To extend the induction to the level-k ears of $\mathcal{D}_1$, we apply the algorithm **Th1-Process-Level** (k) which inductively processes all the ears of $\mathcal{D}_1[k]$. When this procedure terminates all ears of $\mathcal{D}_1[k]$ are marked and decomposed into a set of ears $\mathcal{D}_2[k]$. Moreover, each ear of $\mathcal{D}_2[k]$ is either an anchor ear, a tether ear, or a bridge ear. Below we present pseudocode for this procedure, and a subroutine, and proofs for their correctness. The correctness of the procedure extends the induction to level-k, and the proof of the theorem follows. □

Procedure **Th1-Process-Level** (k)

Input: The level-k ears $\mathcal{D}_1[k]$.

Output: The level-k ears $D_2[k]$ such that the pair of induced tree paths from any vertex of $V(G_k) - V(G_{k-1})$ to a vertex of G_{k-1} are each of length at most $3h$.

Step 1: Initially set $D_2[k]$ to be empty and all ears in $\mathcal{D}_1[k]$ unmarked.

Step 2: Repeat the following until all ears of $\mathcal{D}_1[k]$ are marked:
Pick an unmarked ear P from $\mathcal{D}_1[k]$. If P has an attachment on an ear of $D_2[k]$ then call **Th1-Process-Ear** (P);
else P has both attachments in G_{k-1}. Mark P and insert it in $D_2[k]$ as an *anchor* ear.

Procedure **Th1-Process-Ear** (P)

Input: An unmarked ear P in $\mathcal{D}_1[k]$ with attachment vertices p and q such that q is an internal vertex of an anchor ear Q of $D_2[k]$ and p is a vertex in G_{k-1}.

Output: The ear P, and possibly one other ear of $\mathcal{D}_1[k]$, is marked, and their union is decomposed into one or two ears that are inserted into $D_2[k]$.

Step 1: Determine the set S_P of all ears of $\mathcal{D}_1[k]$ with the property that the ear has an attachment vertex on P (see Figure 1).

Step 2: Order the set $S_P = \{S_1, S_2, \ldots, S_t\}$ as follows:
for all $1 \le i \le t-1$, if S_i has an attachment s_i on P, and S_{i+1} has an attachment s_{i+1} on P, then s_i is no farther from q on P than s_{i+1} is from q (see Figure 2).

Step 3: If the set S_P is empty, then mark P and insert P in $D_2[k]$ as a *tether* ear; else
a) Mark both P and S_1.
b) Insert as an *anchor* ear in $\mathcal{D}_2[k]$ the path defined by the concatenation of S_1 with the subpath in P to the attachment vertex p in G_{k-1}.
c) Insert as a *bridge* ear in $\mathcal{D}_2[k]$ the subpath of P from s_1 to q (see Figure 3).

Correctness: First consider the procedure **Th1-Process-Level** (k). We need only justify the else clause in Step 2. Suppose there are unmarked ears, but none have an attachment vertex on an ear of $\mathcal{D}_2$. By the definition of Breadth-first ear decomposition, all such ears must have both attachments on G_{k-1}. Hence the correctness follows from the correctness of the procedure **Th1-Process-Ear** (P).

Now consider the procedure **Th1-Process-Level** (P). For the inductive hypothesis assume that if P is any ear of $\mathcal{D}_1[k]$ that has an attachment internal to an ear Q of $\mathcal{D}_2[k]$, then Q must be an anchor ear of $\mathcal{D}_2[k]$. The base of the induction holds vacuously. Hence we can assume that the input ear P attaches on an anchor ear Q.

Now consider Step 3 of the procedure. Consider the case where S_P is empty. It is clear that P satisfies properties ($t1$) and ($t2$), and therefore P can be inserted as a tether ear in $\mathcal{D}_2[k]$. Property ($t2$) guarantees that the inductive hypothesis is extended.

If S_P is not empty, then the concatenated path defined in Step 3 clearly satisfies property ($a1$), and therefore P can be inserted as an anchor ear in $\mathcal{D}_2[k]$. Likewise, the remaining subpath satisfies properties ($b1$) and ($b2$), and thus can be inserted as a bridge ear in $\mathcal{D}_2[k]$. Property ($b2$) guarantees that the inductive hypothesis is extended. □

¿From Proposition 5 it follows that for some graphs it is not possible to simultaneously bound the stretch factors of vertex-independent spanning trees in terms solely involving the size of the largest ear in a breadth-first ear decomposition. The next theorem shows that we can obtain bounds that are proportional to the product of this minimax ear length with the degree of the graph. The construction given in the proof of Theorem 2 is similar in spirit to that of the Theorem 1. However, we note that we cannot directly apply the previous construction to obtain vertex-independent trees, since the anchor ears constructed in Theorem 1 are not guaranteed to be open ears, even when $\mathcal{D}_1$ is an open ear decomposition.

We now state and prove the second main result of the paper.

Theorem 2. *Let G be a 2-vertex-connected graph with a distinguished root vertex r, and let $\mathcal{D}_1$ be any given breadth-first open ear decomposition of G rooted at r. Then there exists a pair of vertex-independent spanning trees T_1 and T_2 rooted at r so that $sf(T_1) \leq sf(T_2) \leq h(\delta+2)$, where h is the size of the largest ear in $\mathcal{D}_1$ and δ is maximum degree of G. Moreover, by modifying $\mathcal{D}_1$ we can produce a new open ear decomposition $\mathcal{D}_2$ that induces the pair T_1 and T_2.*

Proof. We modify $\mathcal{D}_1$ to produce a new (not necessarily breadth-first) open ear decomposition $\mathcal{D}_2$ that induces the trees T_1 and T_2 with stretch factors satisfying the theorem. As in the case of Theorem 1, for each k, we will construct a set of ears $\mathcal{D}_2[k]$ whose union is identical as a subgraph to the union of ears in $\mathcal{D}_1[k]$. Each ear defined in $\mathcal{D}_2[k]$ will be one of three possible types: either an *open-anchor ear*, a *chain ear*, or a *bridge ear*.

An *open-anchor ear* Q_o in $\mathcal{D}_2[k]$ has the following properties:

- $o1$) both (distinct) attachment vertices of Q_o are elements of G_{k-1}; and
- $o2$) Q_o is either a single ear of $\mathcal{D}_1[k]$, or the concatenation of subpaths of at most two ears of $\mathcal{D}_1[k]$.

A *chain ear Q_c in $\mathcal{D}_2[k]$ on vertex w* for w in G_{k-1} has the following properties:

- $c1$) one attachment of Q_c is the vertex w, and the other attachment is either an internal vertex of an open-anchor ear or an internal vertex of another chain ear on w;
- $c2$) for all ears of $\mathcal{D}_1[k]$ that attach on Q_c, the other attachment is vertex w; and

$c3)$ Q_c is either a single ear of $\mathcal{D}_1[k]$, or the concatenation of subpaths of at most two ears of $\mathcal{D}_1[k]$.

A *bridge ear Q_b in $\mathcal{D}_2[k]$ on vertex w* for w in G_{k-1} has the following properties:

$b1)$ one attachment vertex of Q_b is an internal vertex of an open-anchor ear of $\mathcal{D}_2[k]$, and the other attachment vertex of Q_b is either an internal vertex of an open-anchor ear or a chain ear on vertex w of $\mathcal{D}_2[k]$;

$b2)$ for all ears of $\mathcal{D}_1[k]$ that attach on Q_c, the other attachment is vertex w; and

$b3)$ Q_b is a subpath of a single ear of $\mathcal{D}_1[k]$.

For the inductive hypothesis, assume for each level $1 \leq i \leq k-1$, the subgraph given by the set of ears in $\mathcal{D}_1[i]$ can be decomposed into a set of ears $\mathcal{D}_2[i]$, such that each ear is either an open-anchor ear or a chain ear. This implies that in the induced vertex-independent trees (defined for the subgraph G_{k-1}) the stretch factor is at most $h(\delta+2)$. This follows from the fact that for each $1 \leq i \leq k-1$, open-anchor ears in $\mathcal{D}_2[i]$ induce tree paths to vertices in G_{i-1} of length at most $2h$. Chain ears on vertex w induce tree paths to vertices in G_{i-1} of length at most $h(\delta+2)$, since from any internal vertex of a chain ear there is one induced tree path of length at most h to the vertex w, while the other induced tree path can overlap at most δ_w chain ears (each on vertex w, where δ_w is the degree of w) plus two ears associated with the open-anchor ear, before reaching a vertex in G_{k-1}. By a similar arguement, bridge ears on vertex w induce tree paths to vertices in G_{i-1} of length at most $h(\delta+2)$.

To extend the induction to the level-k ears of $\mathcal{D}_1$, we apply the algorithm **Th2-Process-Level** (k) which inductively processes all the ears of $\mathcal{D}_1[k]$. When this algorithm terminates all ears of $\mathcal{D}_1[k]$ are marked and decomposed into a set of open ears $\mathcal{D}_2[k]$. Moreover, each ear of $\mathcal{D}_2[k]$ is defined as either an open-anchor ear, or a chain ear or bridge ear on some vertex. Below we present pseudocode for this procedure, and a subroutine, and proofs for their correctness. The correctness of the procedure extends the induction to level-k, and the proof of the theorem follows. □

Procedure **Th2-Process-Level** (k)

Input: The level-k ears $D_1[k]$.

Output: The level-k ears $D_2[k]$ such that the induced tree paths from any vertex of $V(G_k) - V(G_{k-1})$ to a vertex of G_{k-1} are of length at most $h(\delta+2)$.

Step 1: Initially set $D_2[k]$ to be empty and all ears in $D_1[k]$ unmarked.

Step 2: Repeat the following until all ears of $\mathcal{D}_1[k]$ are marked:

Pick an unmarked ear P from $\mathcal{D}_1[k]$. If P has an attachment on an ear of $D_2[k]$ then call **Th2-Process-Ear** (P);
else P has both attachments in G_{k-1}. Mark P and insert it in $D_2[k]$ as an *open-anchor* ear.

Procedure **Th2-Process-Ear** (P)

Input: An unmarked ear P in $D_1[k]$ with distinct attachment vertices p and q such that q is an internal vertex of an ear Q of $D_2[k]$, and p is a vertex in G_{k-1}. Also, Q is either an open-anchor ear or a chain ear on vertex w for w in G_{k-1}.

Output: The ear P, and possibly one other ear of $D_1[k]$, are marked, and their union is decomposed into one or two ears that are inserted into $D_2[k]$.

Step 1: Determine the set S_P of ears of $\mathcal{D}_1[k]$ with attachment vertex on P.

Step 2: Order the set $S_P = \{S_1, S_2, \ldots, S_t\}$ as follows:
for all $1 \leq i \leq t-1$, if S_i has an attachment s_i on P, and S_{i+1} has an attachment s_{i+1} on P, then s_i is no farther from q on P than s_{i+1} is from q.

Step 3: If the set S_P is empty, then mark P and insert P in $D_2[k]$ as a *chain* ear on p; else let S_f be the first ear in S_P having an attachment vertex different from p (see Figure 4).

Step 4: If such an ear S_f exists (see Figure 5), then
a) Mark both P and S_f.
b) Insert as an *open-anchor* ear in $\mathcal{D}_2[k]$ the path defined by the concatenation of S_f with the subpath in P to the attachment vertex p in G_{k-1}.
c) Insert as a *bridge* ear on vertex p in $\mathcal{D}_2[k]$ the remaining subpath of P;
else mark P and insert P as a *chain* ear on p in $\mathcal{D}_2[k]$ (see Figure 6).

Correctness: The correctness proof of the procedure **Th2-Process-Level (k)** is similar to that of **Th1-Process-Level (k)** and so is omitted.

Now consider the procedure **Th2-Process-Ear (P)**. For the inductive hypothesis assume that if P is any ear of $\mathcal{D}_1[k]$ that has an attachment internal to an ear Q of $\mathcal{D}_2[k]$, then Q is an open-anchor ear or a chain ear on vertex w. Also if Q is a chain ear on vertex w, then it follows that the other attachment vertex of P is w. The base of the induction holds vacuously. Hence we can assume that the input ear P attaches on an ear Q that is either an open-anchor ear or a chain ear.

Now consider Step 3 of the procedure—the case where S_P is empty. It is clear that P satisfies properties ($c1$) and ($c2$), and therefore P can be inserted as a chain ear on p in $\mathcal{D}_2[k]$. Property ($c2$) (which holds vacuously) guarantees that the inductive hypothesis is extended.

If S_P is not empty, and S_f exists, then the concatenated path defined in Step 4(b) satisfies property ($o1$), and therefore this path can be inserted as an open-anchor ear in $\mathcal{D}_2[k]$. Likewise, the remaining subpath of P satisfies properties ($b1$) and ($b2$), and thus can be inserted as a bridge ear in $\mathcal{D}_2[k]$. Property ($b2$) guarantees that the inductive hypothesis is extended.

If, as in Step 4, S_P is not empty, but S_f does not exist, then P satisfies properties ($c1$) and ($c2$), and therefore can be inserted as a chain ear in $\mathcal{D}_2[k]$. Again, property ($c2$) guarantees that the inductive hypothesis is extended. $\square$

3.1 Stretch factors and cycle covers

We show in this section that our bounds can be translated in terms of a simple structural properties of the graph. Define a *cycle cover* for a graph to be a set of cycles such that each edge of the graph is contained in at least one cycle. If the maximum size cycle in a cycle cover is m then it follows that our bounds for the stretch factors of edge-independent spanning trees can be stated in terms of m.

Corollary 1. *If G is a 2-edge-connected graph, then there are a pair of independent spanning trees with stretch factors bounded by $O(m)$, where m is the size of the largest cycle in any given cycle cover of G.*

Proof. Given a cycle cover of G, it is straightforward to construct a breadth-first ear decomposition of G with the property that each ear is a subset of edges of a cycle of the cycle cover. The corollary follows from Theorem 1. □

Define an *open cycle cover* for a graph (with respect to some distinguished root vertex r) to be a set of cycles such that each edge uv of the graph is contained in at least one cycle, and that cycle contains at least one other vertex (different from u and v) at a distance no farther from the root r than both u and v.

Corollary 2. *If G is a 2-vertex-connected graph, then there are a pair of vertex-independent spanning trees rooted at a vertex r with stretch factors bounded by $O(\delta m)$, where m is the size of the largest cycle in any given open cycle cover of G with respect to r.* □

References

[1] B. Awerbuch, A. Bar-Noy, N. Linial, and D. Peleg, 1990. Improved routing strategies with succinct tables, *Journal of Algorithms* **11**, pp. 307–341.

[2] J.A. Bondy and U.S.R. Murty. 1976. *Graph Theory with Applications*, North-Holland, New York.

[3] J. Cheriyan and S.N. Maheshwari, 1988. Finding non-separating induced cycles and independent spanning trees in 3-connected graphs, *Journal of Algorithms* **9**, pp. 507–537.

[4] S. Dolev, E. Kranakis, D. Krizanc, and D. Peleg, 1995. Bubbles: Adaptive routing scheme for high-speed dynamic networks, 27*th ACM Symposium on the Theory of Computing*, pp. ???

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

[6] S. Khuller and B. Schieber, 1992. On independent spanning trees, *Information Processing Letters* **42**, pp. 321–323.

[7] A. Lempel, S. Even, and I. Cederbaum, 1966. An algorithm for planarity testing graphs, in *Theory of Graphs*, P. Rosentiehl, (ed)., pp. 215–232.

[8] D. Peleg and E. Upfal, 1989. A tradeoff between size and efficiency for routing tables, *JACM* **36**, pp. 510–530.

[9] H. Whitney. 1932. Non-separable and planar graphs, *Transactions of American Mathematical Society* **34**, pp. 339–362.

[10] A. Zehavi and A. Itai, 1989. Three tree-paths, *Journal of Graph Theory* **13**, pp. 175–188.

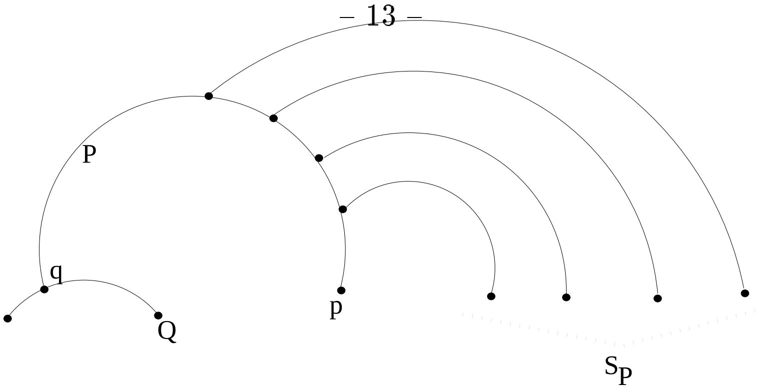

Figure 1: The ears P in $\mathcal{D}_1$ and Q in $\mathcal{D}_2$, and the set of ears S_P that attach on P.

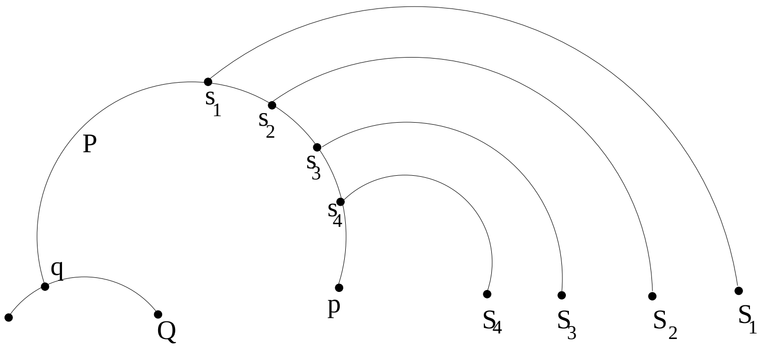

Figure 2: The ears in S_P correctly ordered as in Step 2.

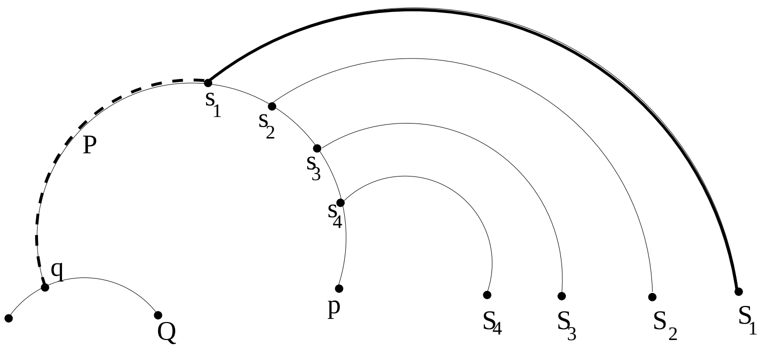

Figure 3: The decomposition of ears P and S_1 into an anchor ear and a bridge ear in $\mathcal{D}_2$.

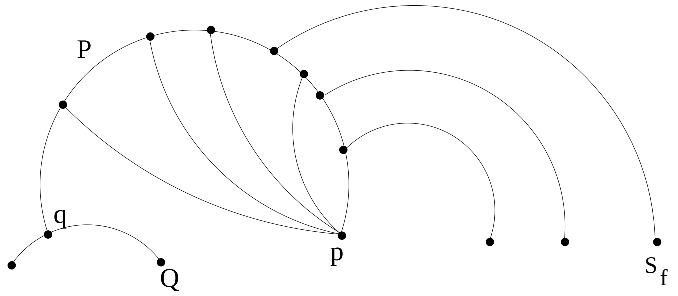

Figure 4: The ears S_f identified in the set S_P.

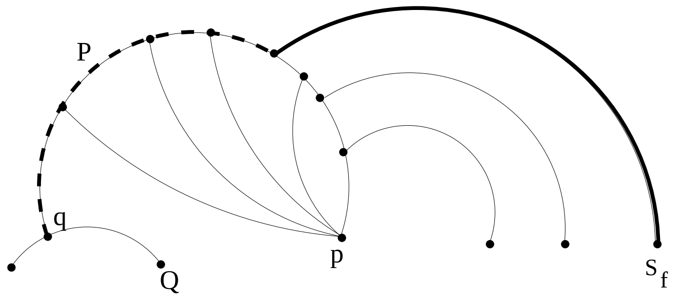

Figure 5: The decomposition of ears P and S_f into an open-anchor ear and a bridge ear in $\mathcal{D}_2$.

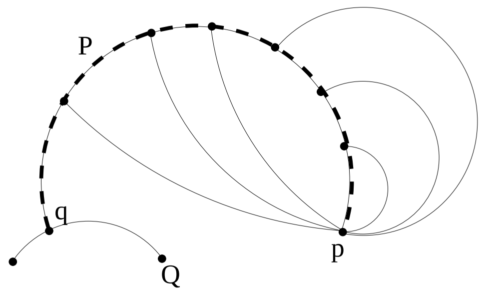

Figure 6: If S_f does not exist, P is inserted as a chain ear on p.