

LISTING SPANNING TREES OF OUTERPLANAR GRAPHS BY PIVOT-EXCHANGES

NASTARAN BEHROOZANIA AND TORSTEN MÜTZE

ABSTRACT. We prove that the spanning trees of any outerplanar triangulation G can be listed so that any two consecutive spanning trees differ in an exchange of two edges that share an end vertex. For outerplanar graphs G with faces of arbitrary lengths (not necessarily 3) we establish a similar result, with the condition that the two exchanged edges share an end vertex or lie on a common face. These listings of spanning trees are obtained from a simple greedy algorithm that can be implemented efficiently, i.e., in time $\mathcal{O}(n \log n)$ per generated spanning tree, where n is the number of vertices of G . Furthermore, the listings correspond to Hamilton paths on the 0/1-polytope that is obtained as the convex hull of the characteristic vectors of all spanning trees of G .

1. INTRODUCTION

For a given graph G , let $\mathcal{T}(G)$ denote the set of all spanning trees of G . Two spanning trees of G differ in an *edge exchange* if the symmetric difference of their edge sets is a 2-element set, i.e., each of the two spanning trees is obtained from the other one by removing one edge and adding another. The *flip graph* $\mathcal{F}(G)$ has the set $\mathcal{T}(G)$ as vertex set, and an edge between any two spanning trees that differ in an edge exchange; see Figure 1. It is well-known that $\mathcal{F}(G)$ is the skeleton of the 0/1-polytope that is obtained as the convex hull of all characteristic vectors $\chi(T)$ of spanning trees $T \in \mathcal{T}(G)$ (see [Sch03, Thm. 40.6]). Specifically, if the edge set of G is $\{1, \dots, m\}$, then for all $e \in \{1, \dots, m\}$ the characteristic vector $\chi(T)$ has a 1-bit at position e if the edge e belongs to T , and a 0-bit at position e otherwise.

We are interested in computing Hamilton paths or cycles in the flip graph $\mathcal{F}(G)$, i.e., we aim to list all spanning trees of G such that any two consecutive spanning trees differ in an edge exchange. Such a listing of combinatorial objects subject to some small-change condition is generally referred to as a *Gray code* [Sav97, Müt23]. A Gray code is called *cyclic* if the first and last object also differ in the small-change condition, i.e., this corresponds to a Hamilton cycle in the flip graph.

Cummins [Cum66] first proved that $\mathcal{F}(G)$ admits a Hamilton cycle for any graph G . He showed more generally that for any prescribed edge of $\mathcal{F}(G)$, there is a Hamilton cycle containing this edge. Similar results were obtained by Shank [Sha68], Kamae [Kam67], and Kishi and Kajitani [KK68].

Harary and Holzmänn [HH72] proved more generally that the base exchange flip graph of any matroid has a Hamilton cycle. They showed this in the stronger sense that any edge of this flip graph can be prescribed to be contained in the cycle, and to be avoided by another cycle.

(Nastaran Behrooznia) DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF WARWICK, UNITED KINGDOM
(Torsten Mütze) INSTITUT FÜR MATHEMATIK, UNIVERSITÄT KASSEL, GERMANY & DEPARTMENT OF THEORETICAL COMPUTER SCIENCE AND MATHEMATICAL LOGIC, CHARLES UNIVERSITY, PRAGUE, CZECH REPUBLIC
E-mail addresses: `nastaran.behrooznia@warwick.ac.uk`, `tmuetze@mathematik.uni-kassel.de`.

This work was supported by Czech Science Foundation grant GA 22-15272S. Both authors participated in the workshop ‘Combinatorics, Algorithms and Geometry’ in March 2024, which was funded by German Science Foundation grant 522790373.




FIGURE 1. The spanning tree flip graph $\mathcal{F}(G)$ for the ‘diamond’ graph G on the left, with a Hamilton cycle highlighted. For each spanning tree, the set of edges is shown above, and the characteristic vector is shown below. Edges of $\mathcal{F}(G)$ are labelled by the two edges of G being exchanged.

Even more generally, Naddef and Pulleyblank [NP81, NP84] proved that the skeleton of any 0/1-polytope either admits a Hamilton path between any two prescribed end vertices, or it is a hypercube, in which case it admits a Hamilton path between any two end vertices of opposite parity.


Algorithmically, a Hamilton path in $\mathcal{F}(G)$ can be computed in time $\mathcal{O}(1)$ on average per generated tree using Smith’s algorithm [Smi97] (this is Algorithm S in [Knu11, Sec. 7.2.1.6]); see Figure 2 (a).

Given these strong Hamiltonicity properties of the graph $\mathcal{F}(G)$, there has recently been interest to strengthen them by restricting the allowed edge exchanges.


1.1. The pivot-exchange property. We introduce some more notation. For a graph G , we write $V(G)$ and $E(G)$ for set of vertices and edges of G , respectively. For a subgraph $H \subseteq G$ and edges $e \in E(H)$ and $f \in E(G) \setminus E(H)$, we write $H - e$ and $H + f$ for the graphs obtained from H by removing and adding the edges e and f , respectively. We will think of a subgraph H as a subset of edges from $E(G)$, so the operations $H - e$ and $H + f$ remove and add an element from the set, respectively. Formally, an *edge exchange* for a spanning tree $T \in \mathcal{T}(G)$ is a pair $\{e, f\}$ of edges such that $e \in E(T)$ and $f \in E(G) \setminus E(T)$ with $T - e + f = T \Delta \{e, f\} \in \mathcal{T}(G)$.

Cameron, Grubb, and Sawada [CGS24] introduced the stronger notion of a *pivot-exchange*, which is an exchange $\{e, f\}$ with the additional property that e and f have a common end vertex. For example, in the graph G shown in Figure 1, all exchanges except $\{1, 5\}$ and $\{2, 4\}$ are pivot-exchanges. They raised the following problem.

(a) Smith [Smi97]



(b) Cameron, Grubb, Sawada [CGS24]



(c) Merino, Mütze, Williams [MMW22]




FIGURE 2. Three different edge exchange Gray codes for listing the 21 spanning trees of the fan graph F_5 . In each spanning tree, the edge removed to reach the next tree is highlighted (prefixed by $-$ in (c)), and the non-edge being added is dashed (prefixed by $+$ in (c)). In (b) and (c), the common end vertex of each pivot-exchange operation is highlighted, and in (c), the common face of each face-exchange operation is highlighted. The right-hand side of (c) shows the characteristic vectors of each spanning tree.

Problem 1. *Does every graph G admit a pivot-exchange Gray code of its spanning trees?*

Additionally, they asked if such a listing can be computed using a greedy strategy, and possibly by an efficient algorithm.

Problem 1 is a special case of a more general question raised by Knuth in Vol. 4A of his seminal series ‘The Art of Computer Programming’ [Knu11], stated as problem 102 in Section 7.2.1.6 with a difficulty rating of 46/50: ¹

¹A flawed attempt at settling this problem was published in [Che67] (see [RR72]).

Problem 2. *Does every directed graph admit an edge exchange Gray code of its oriented spanning trees, also known as arborescences, i.e., spanning trees in which all arcs are oriented away from a fixed root vertex r ?*

Note that for any exchange $\{e, f\}$ in this directed setting, in order to preserve the arborescence property, the arcs e and f must point to the same vertex, i.e., it must be a pivot-exchange. Furthermore, a positive answer to Problem 2 would imply an affirmative answer to Problem 1: Indeed, given an undirected graph G , construct the directed graph G^{\leftrightarrow} by replacing each undirected edge of G by two oppositely oriented arcs, and pick an arbitrary vertex as root r . Listing the oriented spanning trees of G^{\leftrightarrow} by arc exchanges produces every spanning tree of G exactly once, namely in the orientation forced by the choice of r .

1.2. **Fan graphs.** As a first step towards Problem 1, Cameron, Grubb, and Sawada [CGS24] provided a pivot-exchange Gray code for listing the spanning trees of *fan graphs* F_n , which are obtained by joining an extra vertex to all vertices of a path on $n - 1$ vertices; see Figure 3.




FIGURE 3. The fan graphs F_n .

Theorem 3 ([CGS24, Thm. 4]). *For any $n \geq 3$, there is a pivot-exchange Gray code for the spanning trees of F_n .*

The output of their algorithm for the case $n = 5$ is shown in Figure 2 (b). The algorithm uses a greedy strategy that prioritizes exchanges based on vertex labels.

1.3. **A simple greedy algorithm.** Merino, Mütze and Williams [MMW22] discovered a simple greedy algorithm for listing the spanning trees of a graph G by (arbitrary) edge exchanges. The algorithm operates based on a total ordering of the edges of G , which is captured by labeling the edges by integers. Specifically, if m denotes the number of edges of G , then an *edge labeling* is a bijection $\ell : E(G) \rightarrow \{1, \dots, m\}$. For an edge $e \in E(G)$, we refer to $\ell(e)$ as the *label* of the edge e . In the following examples, we will often identify edges by their labels. In particular, edge exchanges $\{e, f\}$ will be denoted by the pairs of labels $\{\ell(e), \ell(f)\}$. We will also use the abbreviation $[m] := \{1, \dots, m\}$.

Algorithm G (*Greedy edge exchanges*). This algorithm greedily generates the spanning trees of a graph G with m edges via edge exchanges, using an edge labeling $\ell : E(G) \rightarrow [m]$ and an initial spanning tree $\tilde{T} \in \mathcal{T}(G)$.

G1. [Initialize] Visit the initial spanning tree \tilde{T} .

G2. [Exchange] Perform an edge exchange in the current spanning tree that minimizes the larger of the two edge labels in the exchange and that yields an unvisited spanning tree from $\mathcal{T}(G)$. If no such exchange exists, then terminate. Otherwise visit this new spanning tree and repeat G2.

The output of Algorithm G when applied to the fan graph F_5 is shown in Figure 2 (c), using the edge labeling displayed on the left. To illustrate the greedy rule in step G2, when reaching the sixth spanning tree $T_6 = \{1, 2, 5, 6\}$, there are seven possible edge exchanges to obtain another spanning tree in $\mathcal{T}(F_5)$, namely $\{1, 3\}$, $\{2, 3\}$, $\{1, 4\}$, $\{2, 4\}$, $\{4, 5\}$, $\{5, 7\}$ and $\{6, 7\}$; see Figure 4. Only $\{1, 4\}$, $\{2, 4\}$, $\{5, 7\}$ and $\{6, 7\}$ give an unvisited spanning tree, and among those, $\{1, 4\}$ and $\{2, 4\}$, minimize the larger label, which is 4. In this case, the exchange $\{2, 4\}$ is applied, but $\{1, 4\}$ would be a valid alternative for the algorithm.

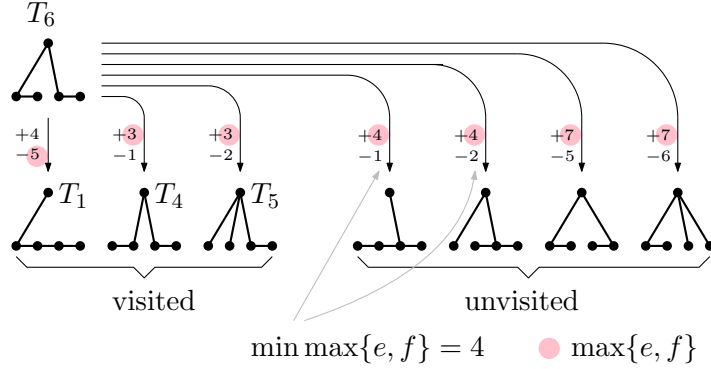


FIGURE 4. Illustration of the sixth iteration of Algorithm G in the run shown in Figure 2 (c).

In general, in step G2 there may be several edge exchanges

$$\{e_1, f\}, \{e_2, f\}, \dots, \{e_t, f\} \text{ with } \ell(e_1) < \ell(e_2) < \dots < \ell(e_t) < \ell(f), \quad (1)$$

applicable to the current spanning tree to give an unvisited spanning tree, which all have $\ell(f)$ as the larger label, differing only in the smaller label $\ell(e_i)$, $i \in [t]$. We refer to such a situation as a *tie*. A *tie-breaking rule* is a procedure that determines which exchange to apply in case of a tie.

By definition of the step G2, Algorithm G only selects edge exchanges that result in a previously unvisited spanning tree of G , i.e., the produced listing of spanning trees will not contain repetitions. However, it could be that the algorithm terminates before having visited the entire set $\mathcal{T}(G)$, a situation that is ruled out by the next theorem.

Theorem 4 ([MMW22, Thm. 10]). *For any graph G , for any edge labeling of G , for any initial spanning tree $\tilde{T} \in \mathcal{T}(G)$, and for any tie-breaking rule, Algorithm G yields a genlex listing of all spanning trees of G .*

A listing of bitstrings is called *genlex* if all strings with the same suffix appear consecutively. In particular, all strings ending with 0 appear before all strings ending with 1, or vice versa. A genlex listing of bitstrings is also sometimes referred to as suffix-partitioned in the literature. Note that genlex order generalizes colexicographic order. In the context of spanning trees, the *genlex* property refers to the corresponding characteristic vectors $\chi(T)$ of spanning trees $T \in \mathcal{T}(G)$; see the right-hand side of Figure 2 (c). In other words, all spanning trees not containing the highest-labeled edge appear before all spanning trees containing this edge, or vice versa, and this property is true recursively within the blocks.

As evidenced by this theorem, the greedy Algorithm G is very powerful and versatile. In fact, the algorithm can be generalized for listing the bases of any matroid by base exchanges, and even more generally, for traversing a Hamilton path on the skeleton of any 0/1-polytope [MM24].

1.4. The face-exchange property. The paper [MMW22] also introduced another closeness condition for edge exchanges, which is well-defined only for plane graphs. Specifically, a *face-exchange* between two spanning trees of a plane graph is an exchange of two edges that lie on a common face. If an edge exchange is both a pivot-exchange *and* face-exchange, then we refer to it as a *pivot \wedge face-exchange*. A weaker requirement is that it is a pivot-exchange *or* a face-exchange, and then we refer to it as a *pivot \vee face-exchange*. These notions give rise to the following question.

Problem 5. *Does every plane graph G admit a pivot \wedge face-exchange or a pivot \vee face-exchange Gray code of its spanning trees?*

Pivot- and face-exchanges are connected through the well-known concept of the dual graph; see Figure 5 (a). For a plane graph G , we write $F(G)$ for the set of faces of G . The *dual graph* G' is the plane graph obtained from G as follows: For every face $\alpha \in F(G)$, the dual graph G' has a vertex $\alpha' \in V(G')$, and for every edge e of G between faces α and β of G , the dual graph G' has the edge $e' = (\alpha', \beta')$. For every vertex $v \in V(G)$, we write $v' \in F(G')$ for the corresponding face of G' dual to it. For a spanning tree $T \in \mathcal{T}(G)$, the *dual spanning tree* $T' \in \mathcal{T}(G')$ has the dual edge $e' \in E(T')$ for every non-edge $e \in E(G) \setminus E(T)$ and a dual non-edge $e' \notin E(T')$ for every edge $e \in E(T)$; see Figure 5 (b). Observe that a pivot-exchange $\{e, f\}$ in T is a face-exchange $\{e', f'\}$ in the dual spanning tree T' ; see Figure 5 (c). Symmetrically, a face-exchange $\{e, f\}$ in T is a pivot-exchange $\{e', f'\}$ in T' .




FIGURE 5. Connection between pivot- and face-exchanges via the dual spanning tree.

Merino, Mütze and Williams [MMW22] strengthened Theorem 3, by showing that for a suitable edge labeling of the fan graph F_n and a suitable tie-breaking rule, Algorithm G yields a pivot \wedge face-exchange Gray code. Specifically, the edges of F_n are labeled from *left-to-right*, as shown in Figure 2 (c), and ties are broken according to the *closest* tie-breaking rule, which in case of a tie as in (1) selects the exchange $\{e_t, f\}$, i.e., the exchange that maximizes the smaller of the edge labels $\ell(e_t)$ (equivalently, the one for which the smaller label is closest to the larger label $\ell(f)$).

Theorem 6. *For any $n \geq 3$, for the left-to-right labeling of the edges of F_n , for any initial spanning tree $\tilde{T} \in \mathcal{T}(F_n)$, and for the closest tie-breaking rule, Algorithm G yields a genlex pivot \wedge face-exchange Gray code for the spanning trees of F_n .*

In fact, the Gray code mentioned in Theorem 3 also has the stronger pivot \wedge face-exchange property.




FIGURE 6. (a) Pivot-exchange Gray code for the spanning trees of the outerplane triangulation G . (b) Pivot/face-exchange Gray code for the spanning trees of the outerplane graph G' , which has the four marked face-exchanges $\{1, 7\}$, and the rest pivot-exchanges. Both listings were computed by Algorithm G, and the spanning trees are represented by their characteristic vectors, with 1-bits and 0-bits drawn as black and white squares, respectively. The initial spanning tree is highlighted in both graphs.

1.5. **Our results.** In this work, we solve Problems 1 and 5 for certain families of plane graphs that generalize fan graphs, thus generalizing Theorems 3 and 6.

An *outerplane* graph is a plane graph in which all vertices are incident to the outer face. An outerplane graph is a *triangulation* if all of its faces, except possibly the outer face, are triangles. Note that fan graphs are a very special case of outerplane triangulations.

Theorem 7. For any outerplane triangulation G , there is an edge labeling of G , so that for any initial spanning tree $\tilde{T} \in \mathcal{T}(G)$, there is a tie-breaking rule for which Algorithm G yields a genlex pivot-exchange Gray code for the spanning trees of G .

Theorem 8. For any outerplane graph G , there is an edge labeling of G , so that for any initial spanning tree $\tilde{T} \in \mathcal{T}(G)$, there is a tie-breaking rule for which Algorithm G yields a genlex pivot\face-exchange Gray code for the spanning trees of G .

These theorems directly yield efficient algorithms. Specifically, using the techniques described in [MM24, Sec. 7.2+Cor. 32], Algorithm G can be implemented to output each spanning tree in time $\mathcal{O}(n \log n)$, where n is the number of vertices of G . The required space for the algorithm is $\mathcal{O}(n)$. In particular, this implementation is *history-free* in the sense that no previously computed spanning trees apart from the current spanning tree need to be stored, plus some simple data structures for bookkeeping.

Two examples of Gray code listings of spanning trees obtained from these theorems are shown in Figure 6.

Corollary 9. Any outerplane triangulation admits a genlex pivot-exchange Gray code of its spanning trees. Any outerplane graph admits a genlex pivot\face-exchange Gray code of its spanning trees.

The *weak dual* graph of a plane graph G , denoted G^- is the graph obtained from the dual graph G' by removing the vertex corresponding to the outer face of G ; see Figure 7 (a).




FIGURE 7. Illustration of (a) weak dual graph; (b) 2-connected outerplane triangulation and its weak dual tree; (c) 2-connected outerplane triangulation whose weak dual tree is a path.

Note that G is a 2-connected outerplane graph if and only if the weak dual G^- is a tree; see Figure 7 (b). Also note that G^- is a path if and only if G is 2-connected and all but at most two sides of every inner face are incident with the outer face; see Figure 7 (c). In particular, for a triangulation G , we have that G^- is a path if and only if G is 2-connected and at least one side of every triangle touches the outer face. This is true in particular for fan graphs F_n .

We show that fan graphs, and more generally outerplane triangulations for which the weak dual is a path, are exactly the outerplane graphs that have the maximum number of spanning trees for a fixed number of edges. Moreover, the counts are the Fibonacci numbers, defined as $f_0 := 0$, $f_1 := 1$ and

$$f_{m+1} := f_m + f_{m-1} \tag{2}$$

for all $m \geq 1$. We write $t(G) := |\mathcal{T}(G)|$ for the number of spanning trees of G .

Theorem 10. *For any outerplane graph G with m edges we have $t(G) \leq f_{m+1}$, with equality if and only if G is a triangulation such that G^- is a path.*

The identity $t(G) = f_{m+1}$ when G is a triangulation for which G^- is a path was already noticed by Slater [Sla77, Prop. 1].

Remark 11. We remark that Problems 1, 2 and 5 are perfectly valid also for multigraphs, i.e., graphs that may have parallel edges and/or loops (though loops are irrelevant in the context of spanning trees). In fact, Theorems 4, 7 and 8, and hence Corollary 9, also hold in this more general setting. An ‘outerplane triangulation’ in this case has all inner faces of lengths at most 3, instead of exactly 3. The time and space bounds stated after Theorem 8 change to $\mathcal{O}(m \log n)$ and $\mathcal{O}(m)$, respectively, where m is the number of edges of the multigraph. However, for simplicity we do our proofs in the setting of simple graphs, where no parallel edges nor loops are allowed. Nonetheless, our proof of Theorem 10 will actually use multigraphs.

Remark 12. The listings of spanning trees produced by Algorithm G are in general not Hamilton cycles in $\mathcal{F}(G)$, but only Hamilton paths, i.e., the first and last spanning tree in general do not differ in an edge exchange. This remark applies to all the Gray codes mentioned in Theorems 4, 6, 7, and 8, and in Corollary 9. For some graphs, some edge labelings, and some initial spanning trees, however, the Gray codes are cyclic, such as the one mentioned in Theorem 6 for the L-shaped initial spanning tree shown in Figure 2 (c), which ends with the mirrored L.

1.6. Related work. There has been an extensive amount of work [RT75, GM78, Mat97, KR91, ST95, STU97] on efficiently generating the set $\mathcal{T}(G)$ of all spanning trees of G , *without* the requirement that any two consecutive trees differ in an edge exchange, i.e., the computed listings are *not* Hamilton paths or cycles in the flip graph $\mathcal{F}(G)$. The algorithms in the last three of the aforementioned papers achieve this in time $\mathcal{O}(1)$ on average per generated tree. See the survey [CCC⁺19] for a comparison of the different algorithms.

If instead of listing all spanning trees of G , we want to count them, then this can be achieved by Kirchhoff’s Matrix-Tree Theorem, which reduces the problem to computing a determinant, which can be done efficiently. This problem is closely connected to finding the so-called *most vital edge* of G , which is the edge contained in the most spanning trees from $\mathcal{T}(G)$. Random sampling [Bro89, Ald90] and ranking/unranking [CDN89, CMN96] of spanning trees have also been considered.

1.7. Outline of this paper. In Section 2 we prove Theorems 7 and 8. In Section 3 we prove Theorem 10. We conclude with some open questions in Section 4, and there we also report on some experimental evidence.

2. PROOF OF THEOREMS 7 AND 8

For the proofs of Theorems 7 and 8, we will assume w.l.o.g. that the outerplane graph G is 2-connected. If G is not 2-connected, then all the arguments presented in the following apply to each of its blocks, and the spanning trees of G are obtained by combining the spanning trees in each block in all possible ways, and pivot- or face-exchanges remain valid within the blocks.

We first describe the labeling of edges of a 2-connected outerplane graph G that we use in order to run Algorithm G on the graph G . We then establish two important properties of these labelings (Lemmas 13 and 14) that are crucial to show that whenever an arbitrary edge exchange becomes applicable in step G2 of Algorithm G , then ties can be broken to instead use a pivot- or face-exchange (Lemma 15).

2.1. The edge labeling. To define the edge labeling, we need another definition. Let v be the vertex of G' corresponding to the outer face of G , and let d be the degree of v . The *split dual* graph, denoted G^* , is obtained from G' by splitting v into d many degree-1 vertices, one adjacent to each neighbor of v ; see Figure 8 (a).

Note that the weak dual graph G^- is a tree if and only if the split dual graph G^* is a tree; see Figures 7 (b) and 8 (b). Furthermore, G^- is a path if and only if G^* is a caterpillar, i.e., a tree in which all vertices are in distance ≤ 1 from a central path; see Figures 7 (c) and 8 (c).

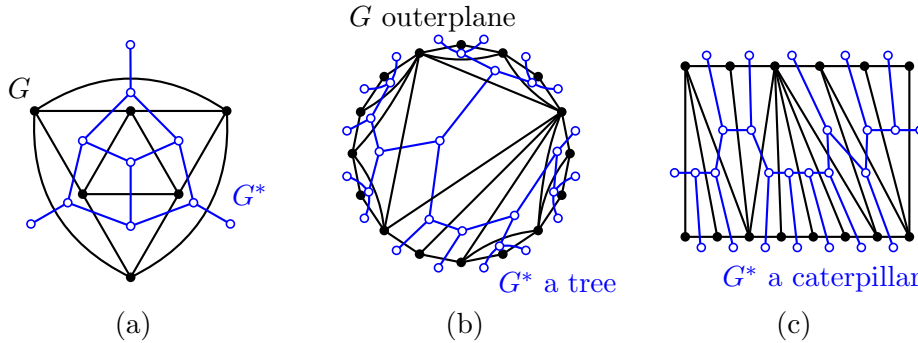


FIGURE 8. Illustration of (a) split dual graph; (b) 2-connected outerplane triangulation and its split dual tree; (c) 2-connected outerplane triangulation whose split dual tree is a caterpillar (cf. Figure 7).

Let G be a 2-connected outerplane graph with m edges. We consider the split dual tree G^* , we pick a leaf r of this tree as root, and we orient all of its edges away from r , giving an oriented tree \vec{G}^* , referred to as *oriented split dual*. We label the edges of \vec{G}^* in a depth-first-search manner with integers $1, \dots, m$, starting at the root r and processing subtrees in counterclockwise (ccw) order; see Figure 9. The labeling of the edges of \vec{G}^* induces a labeling of the corresponding dual edges of G with integers from $[m]$. We refer to this labeling of the edges of G as *dual-tree labeling*. Note that there is freedom in the choice of the root in the tree G^* , and different choices yield different oriented split dual trees \vec{G}^* and hence different edge labelings of the same graph G . Note that the left-to-right labeling of the fan graph $G = F_n$ shown in Figure 2 (c) is one particular dual-tree labeling.




FIGURE 9. Illustration of the dual-tree labeling procedure.

2.2. Proofs of theorems. Throughout this section, we let G be a 2-connected outerplane graph with m edges, \vec{G}^* an oriented split dual tree, and $\ell : E(G) \rightarrow [m]$ the corresponding dual-tree labeling of the edges of G . Before proving Theorems 7 and 8, we first derive two properties of the edge labelings defined in the previous section, stated in Lemmas 13 and 14 below.




FIGURE 10. Illustration of lobes.

The following definitions are illustrated in Figure 10. We denote a face α of G of length t by the sequence of edges (e_1, \dots, e_t) bounding this face in ccw order, starting with the edge e_1 whose dual edge in \vec{G}^* is oriented towards α' (all other edges of \vec{G}^* are oriented away from α'). For a face $\alpha = (e_1, \dots, e_t)$ and an index $i \in [t]$, the (α, e_i) -lobe L_{α, e_i} is the set of edges of G that are dual to the edges in the maximal subtree of G^* that contains the edge e'_i , but none of the other edges e'_j , $j \in [t] \setminus \{i\}$.

Lemma 13. *For any face $\alpha = (e_1, \dots, e_t)$ of G , we have $\ell(e_1) < \ell(e_2) < \dots < \ell(e_t)$. Furthermore, for any $i \in \{2, \dots, t\}$ and $f \in L_{\alpha, e_i} \setminus \{e_i\}$ we have $\ell(e_i) < \ell(f)$, and $\ell(f) < \ell(e_{i+1})$ if $i < t$.*

Proof. This is an immediate consequence of the labeling procedure which processes subtrees in the oriented split dual in ccw order and in a depth-first-search manner. Specifically, for all $i \in \{2, \dots, t\}$, all edges of $L_{\alpha, e_i} \setminus \{e_i\}$ are labeled directly after e_i , and directly before e_{i+1} if $i < t$. \square

For a vertex v of G , the *incidence list* E_v is the sequence of edges incident to v in clockwise order around v , starting and ending with the two edges that bound the outer face; see Figure 11. We say that an edge e incident with v is a *cw-edge* or *ccw-edge*, respectively, according to the clockwise or counterclockwise orientation of the dual edge e' in \vec{G}^* around v .

Lemma 14. *For any vertex v , let $E_v =: (e_1, \dots, e_t)$ be its incidence list. Then there is an index $i \in \{0, \dots, t\}$ such that e_1, \dots, e_i are ccw-edges, e_{i+1}, \dots, e_t are cw-edges, and we have $\ell(e_i) < \ell(e_{i-1}) < \dots < \ell(e_1) < \ell(e_{i+1}) < \ell(e_{i+2}) < \dots < \ell(e_t)$. In particular, for any cw-edge e_j and any $i \in [j-1]$ we have $\ell(e_i) < \ell(e_j)$.*

Proof. For any sequence of edges $P = (a_1, \dots, a_t)$ that form a path in the oriented tree \vec{G}^* , there is an index $i \in \{0, \dots, t\}$ such that a_1, \dots, a_i are oriented towards the start vertex of P and a_{i+1}, \dots, a_t are oriented towards the end vertex of P . The special cases $i = 0$ and $i = t$ mean that the path is oriented conformly, towards the end or start vertex, respectively.

The sequence of dual edges $P := (e'_1, \dots, e'_t)$ is a path in \vec{G}^* to which the aforementioned observation applies. For $j \in \{1, \dots, t-1\}$, let f_j be the face of G bounded by e_j and e_{j+1} . Furthermore, we denote the outer face of G by $f_0 = f_t$. Note that e'_j is oriented towards f'_{j-1} in \vec{G}^* for $j = 1, \dots, i$, i.e., these are all ccw-edges w.r.t. v . Furthermore, e'_j is oriented towards f'_j in \vec{G}^* for $j = i+1, \dots, t$, i.e., these are all cw-edges w.r.t. v .




FIGURE 11. Illustration of Lemma 14.

Applying Lemma 13 to the faces f_j , $j = 1, \dots, i-1$ and $j = i+1, \dots, t-1$, yields $\ell(e_i) < \ell(e_{i-1}) < \dots < \ell(e_1)$ and $\ell(e_{i+1}) < \ell(e_{i+2}) < \dots < \ell(e_t)$, respectively. Furthermore, applying Lemma 13 to L_{f_i, e_i} and $L_{f_i, e_{i+1}}$ proves that $\ell(e_1) < \ell(e_{i+1})$. Combining these inequalities proves the lemma. \square

Lemma 15. *Let $T \in \mathcal{T}(G)$ and let $\{e, f\}$ with $\ell(e) < \ell(f)$ be an edge exchange for T . Then there is a pivot \vee face-exchange $\{d, f\}$ with $\ell(d) < \ell(f)$. Furthermore, if G is a triangulation, then $\{d, f\}$ is a pivot-exchange.*

Proof. We let $\alpha = (e_1, \dots, e_t)$ be the face incident with f in G for which the dual edge f' in $\overrightarrow{G^*}$ is oriented away from α' . As $\ell(f) > \ell(e) \geq 1$ we have $\ell(f) > 1$, and consequently α is not the outer face, but an inner face. We have $f = e_i$ for some $i \in \{2, \dots, t\}$. The graph $T \cup \{e, f\}$ contains exactly one cycle C , which contains both edges e and f . As $\ell(e) < \ell(f)$, Lemma 13 implies that $e \notin L_{\alpha, f}$, and as $e \in C$, we obtain that C contains edges from every lobe L_{α, e_i} for all $i \in [t]$, and furthermore $e \in L_{\alpha, e_j}$ for some $j \in \{1, \dots, i-1\}$.

We now distinguish the cases whether $f \in T$ or $f \notin T$, i.e., whether the exchange $\{e, f\}$ adds the edge e and removes f , or removes e and adds f , respectively. These two cases are illustrated in Figure 12 (a) and (b), respectively.

Case (a): $f \in T$. Note that $\{d, f\}$ with $d := e_j$ is a valid face-exchange for T , and we have $\ell(d) < \ell(f)$ by Lemma 13. Furthermore, if G is a triangulation, then we have $t = 3$ and consequently $f = e_i$ and $d = e_j$ share an end vertex, so the exchange $\{d, f\}$ is a pivot-exchange.

Case (b): $f \notin T$. Let d be the edge of C incident with f in the lobe $L_{\alpha, e_{i-1}}$. Note that $\{d, f\}$ is a valid pivot-exchange for T . If $i > 2$, then we have $\ell(d) < \ell(f)$ by Lemma 13. If $i = 2$, then let v be the common end vertex of $e_1 = e_j$ and $e_2 = e_i$, and consider the incidence list E_v of v . The edge e_1 is a cw-edge and d comes before it (or is equal to e_1) on the list E_v , and so Lemma 14 yields $\ell(d) \leq \ell(e_1)$. By Lemma 13 we have $\ell(e_1) < \ell(f)$ and therefore $\ell(d) < \ell(f)$. \square

We are now in position to prove Theorems 7 and 8.

Proof of Theorems 7 and 8. By Lemma 15, whenever Algorithm G considers an edge exchange $\{e, f\}$ with $\ell(e) < \ell(f)$, there is a pivot \vee face-exchange $\{d, f\}$ with $\ell(d) < \ell(f)$. By Theorem 4, the listing of spanning trees produced by Algorithm G has the genlex property, which implies that if $T \triangle \{e, f\}$ is unvisited, then $T \triangle \{d, f\}$ is also unvisited. It follows that there is a tie-breaking rule for Algorithm G to only ever use pivot \vee face-exchanges.

Furthermore, if G is a triangulation, then by Lemma 15 the alternative exchange $\{d, f\}$ is a pivot-exchange, so there is a tie-breaking rule for Algorithm G to only ever use pivot-exchanges. \square




FIGURE 12. Illustration of the two cases in the proof of Lemma 15.

According to the proof, the tie-breaking rule for Algorithm G that works is to select an arbitrary pivot-exchange or pivot \vee face-exchange, respectively.

3. PROOF OF THEOREM 10

We prove Theorem 10 in the more general setting of multigraphs (recall Remark 11), i.e., from now on we allow multiple edges between pairs of vertices. Two edges between the same pair of vertices are sometimes referred to as *parallel* edges. Loops may be present, but are never contained in a spanning tree and hence irrelevant for us. However, the distinction of parallel edges is important. For example, the plane graph formed by two vertices connected by m parallel edges has m different spanning trees, each containing exactly one of the m edges.

All notions introduced in Section 1 are valid in the more general context of multigraphs. The only exception is the definition of an *outerplane triangulation*, which we change from ‘face length 3’ to ‘face length ≤ 3 ’. A length-2 face comes from two parallel edges, and is called a *digon*.

In addition to edge removal and edge addition, we now also introduce the operation of *edge contraction*. Given a graph G and an edge $e \in E(G)$, we write G/e for the graph obtained from G by contracting the edge e . Note that even if G is a simple graph, G/e may still contain parallel edges, and if G has parallel edges, then G/e may contain loops (which can be ignored when counting spanning trees). The number of spanning trees $t(G)$ of a graph G obeys the well-known recursive relation

$$t(G) = t(G - e) + t(G/e), \quad (3)$$

valid for any (non-loop) edge $e \in E(G)$ (see [Min65]).

We need the following auxiliary lemma about Fibonacci numbers.

Lemma 16. *For any two integers $i, j \geq 1$, we have $f_i \cdot f_j \leq f_{i+j-1}$, with equality if and only if $i = 1$ or $j = 1$.*

Proof. If $i = 1$ or $j = 1$, then one can check directly that the claimed equality holds.

Now assume that $i, j \geq 2$. We prove that $f_i \cdot f_j < f_{i+j-1}$ by induction on $i + j$. In the base case $i + j = 4$ we have $i = j = 2$ and therefore $f_i = f_j = f_2 = 1$ and $f_{i+j-1} = f_3 = 2$, so the claim is true. For the induction step we distinguish the cases $j \in \{2, 3\}$ and $j > 3$. If $j = 2$ we have $f_i \cdot f_j = f_i < f_{i+1} = f_{i+j-1}$ (recall that $i \geq 2$). If $j = 3$ we have $f_i \cdot f_j = 2f_i < f_i + f_{i+1} = f_{i+2} = f_{i+j-1}$. It remains to consider the case $j > 3$. We have

$$f_i \cdot f_j = f_i \cdot (f_{j-2} + f_{j-1}) = f_i \cdot f_{j-2} + f_i \cdot f_{j-1} < f_{i+j-3} + f_{i+j-2} = f_{i+j-1},$$

where the strict inequality in the third step holds by induction, using that $j - 2 > 1$. \square

Theorem 10 is an immediate consequence of the following more general statement for multi-graphs.

Theorem 17. *For any outerplane (multi)graph G with m edges we have $t(G) \leq f_{m+1}$, with equality if and only if G is a triangulation such that G^- is a path and all digons are incident with the outer face.*

Note that digons incident with the outer face correspond to degree 1 vertices in the dual graph G^- , so if G^- is a path, then there are at most two such digons corresponding to end vertices of the path.

Theorem 17 is illustrated in Figure 13.




FIGURE 13. Sequence of outerplane graphs that maximize the number of spanning trees, with the corresponding counts, the Fibonacci numbers. They are all triangulations for which G^- is a path, and they either have no digons (red) or one digon incident with the outer face (blue). For each graph, deleting or contracting the bold edge yields the two preceding graphs.

Proof. We argue by induction on m . The induction basis $m = 0$ is trivial. For the induction step we assume that $m \geq 1$.

If G is not 2-connected, then it has at least two blocks A and B with $a \geq 1$ and $b \geq 1$ edges, respectively, where $a + b = m$. We have $t(G) = t(A) \cdot t(B)$ and therefore $t(G) \leq f_{a+1} \cdot f_{b+1} < f_{a+b+1} = f_{m+1}$, where the strict inequality follows from Lemma 16.

It remains to consider the case that G is 2-connected. If G is a single edge, then $t(G) = 1 = f_2$. If G has only one inner face, then it is a cycle and we have $t(G) = m \leq f_{m+1}$, and this inequality is tight for $m \in \{2, 3\}$ and strict for $m \geq 4$. The same estimates hold if all inner faces of G are digons, i.e., G has 2 vertices and all m edges are parallel to each other. For the rest of the proof we assume that G has at least two faces, not all of which are digons. This implies in particular that $m \geq 4$. We consider a face α of G for which the dual vertex α' has degree at most 1 in G^- . We denote the sequence of edges bounding α by e_1, \dots, e_t in clockwise order, ending with the edge e_t that has a dual edge in G^- . We distinguish two cases, namely whether α is a digon ($t = 2$) or not ($t \geq 3$), and we bound the number of spanning trees using (3) with respect to the edge $e := e_1$, i.e., the first edge of the face α . These two cases are illustrated in Figure 14 (a) and (b), respectively.

Case (a): $t = 2$. $G - e$ has $m - 1$ edges, and therefore

$$t(G - e) \leq f_m \tag{4}$$

by induction. Let $s \geq 2$ be the number of edges that are parallel to e_1 (including e_1 and e_2), and denote them by e_1, e_2, \dots, e_s in ccw order around the corresponding common end vertex; see

Figure 14 (a). Furthermore, we let β be the face incident with e_s but not e_{s-1} . In the graph G/e , the edges e_2, \dots, e_s are loops, and we can therefore remove all of them, so we have

$$t(G/e) = t(G/e - \{e_2, \dots, e_s\}) \leq f_{m-1} \quad (5)$$

by induction. Combining (2), (3), (4), and (5) proves that $t(G) \leq f_{m+1}$, as claimed. Furthermore, this inequality is tight if and only if (4) and (5) are tight, which happens if and only if the conditions stated in the theorem hold for $G - e$ and G/e , which happens if and only if $s = 2$ and β is a triangle that is incident with the outer face, if and only if the conditions stated in the theorem hold for G .

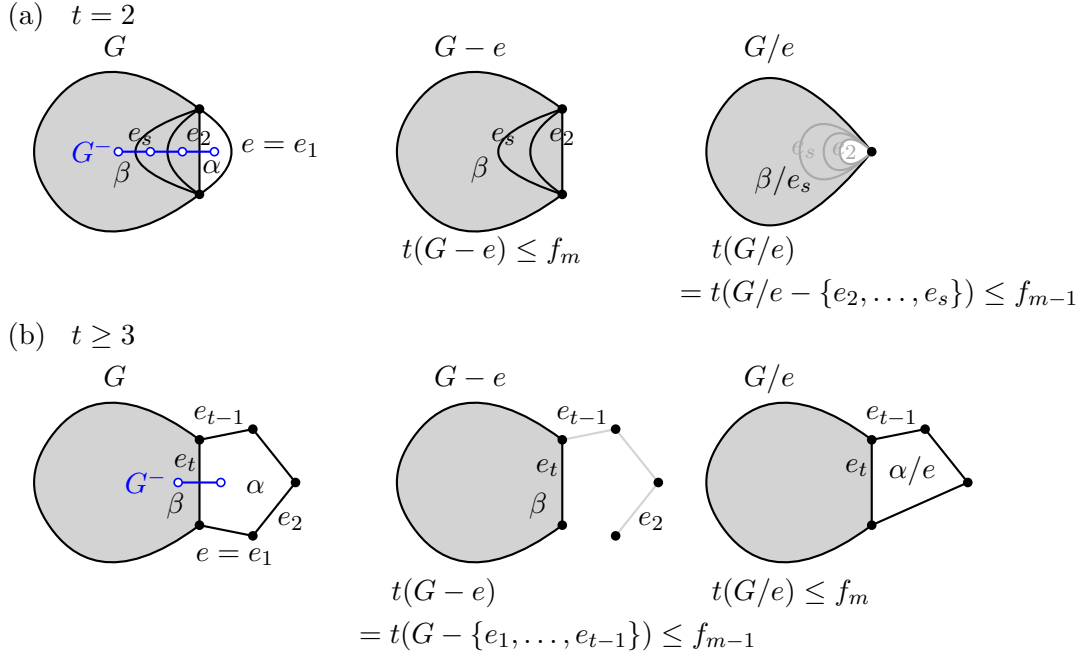


FIGURE 14. Illustration of the proof of Theorem 17.

Case (b): $t \geq 3$. Let β be the face incident with e_t but not e_1 . G/e has $m - 1$ edges, and therefore

$$t(G/e) \leq f_m \quad (6)$$

by induction. In the graph $G - e$, the edges e_2, \dots, e_{t-1} are present in every spanning tree, so we have

$$t(G - e) = t(G - \{e_1, \dots, e_{t-1}\}) \leq f_{m-1} \quad (7)$$

by induction. Combining (2), (3), (6) and (7) proves that $t(G) \leq f_{m+1}$, as claimed. Furthermore, this inequality is tight if and only if (6) and (7) are tight, which happens if and only if the conditions stated in the theorem hold for G/e and $G - e$, which happens if and only if $t = 3$ and β is a triangle that is incident with the outer face, if and only if the conditions stated in the theorem hold for G . \square

4. OPEN PROBLEMS

Problems 1 and 5 remain open for general graphs, i.e., for graphs not covered by Theorems 7 and 8 (recall Remark 11). Also, Knuth's more general Problem 2 for directed graphs remains very much open.

Concerning Problem 1, we checked experimentally that all 2-connected graphs on up to 6 vertices admit a pivot-exchange Gray code for their spanning trees, and these Gray codes are all cyclic. Regarding Problem 2, we checked that all directed graphs on up to 5 vertices (oppositely oriented arcs are allowed) whose underlying undirected graph is 2-connected admit an edge exchange Gray code for their arborescences, for any choice of fixed root. In some cases those flip graphs admit no Hamilton cycle, but only a Hamilton path (see [RR72]), so not all of these Gray codes are cyclic. Regarding Problem 5, we checked that all outerplane graphs on up to 6 vertices admit a cyclic pivot \wedge face-exchange Gray code for their spanning trees.

Furthermore, Knuth [Knu11] asked in Exercise 101 in Section 7.2.1.6 whether the complete graph K_n admits a ‘nice’ Gray code listing of its spanning trees. One interpretation of this question, in the spirit of Problem 1, could be: Is there a pivot-exchange Gray code for the spanning trees of K_n ? Alternatively, is there a Gray code listing that provides a more fine-grained explanation of Cayley’s formula n^{n-2} for the number of spanning trees? Similar questions can be asked for the complete bipartite graph $K_{n,n}$.

Does Theorem 10 hold without the outerplane requirement for m sufficiently large? The only complete graphs K_n that violate the bound $t(K_n) \leq f_{\binom{n}{2}+1}$ arise for $n = 4, 5, 6$.

ACKNOWLEDGEMENTS

We thank both reviewers of this paper for a number of very helpful suggestions that improved the writing.

REFERENCES

- [Ald90] D. J. Aldous. The random walk construction of uniform spanning trees and uniform labelled trees. *SIAM J. Discrete Math.*, 3(4):450–465, 1990.
- [Bro89] A. Z. Broder. Generating random spanning trees. In *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*, pages 442–447. IEEE Computer Society, 1989.
- [CCC⁺19] M. Chakraborty, S. Chowdhury, J. Chakraborty, R. Mehera, and R. K. Pal. Algorithms for generating all possible spanning trees of a simple undirected connected graph: an extensive review. *Complex & Intelligent Systems*, 5:265–281, 2019.
- [CDN89] C. J. Colbourn, R. P. J. Day, and L. D. Nel. Unranking and ranking spanning trees of a graph. *J. Algorithms*, 10(2):271–286, 1989.
- [CGS24] B. Cameron, A. Grubb, and J. Sawada. Pivot Gray codes for the spanning trees of a graph ft. the fan. *Graphs Combin.*, 40(4):Paper No. 78, 26 pp., 2024.
- [Che67] W.-K. Chen. Hamilton circuits in directed-tree graphs. *IEEE Trans. Circuit Theory*, CT-14:231–233, 1967.
- [CMN96] C. J. Colbourn, W. J. Myrvold, and E. Neufeld. Two algorithms for unranking arborescences. *J. Algorithms*, 20(2):268–281, 1996.
- [Cum66] R. L. Cummins. Hamilton circuits in tree graphs. *IEEE Trans. Circuit Theory*, CT-13:82–90, 1966.
- [GM78] H. N. Gabow and E. W. Myers. Finding all spanning trees of directed and undirected graphs. *SIAM J. Comput.*, 7(3):280–287, 1978.
- [HH72] C. A. Holzmann and F. Harary. On the tree graph of a matroid. *SIAM J. Appl. Math.*, 22:187–193, 1972.
- [Kam67] T. Kamae. The existence of a Hamilton circuit in a tree graph. *IEEE Trans. Circuit Theory*, CT-14:279–283, 1967.
- [KK68] G. Kishi and Y. Kajitani. On Hamilton circuits in tree graphs. *IEEE Trans. Circuit Theory*, CT-15(1):42–50, 1968.
- [Knu11] D. E. Knuth. *The Art of Computer Programming. Vol. 4A. Combinatorial Algorithms. Part 1*. Addison-Wesley, Upper Saddle River, NJ, 2011.
- [KR91] S. Kapoor and H. Ramesh. Algorithms for generating all spanning trees of undirected, directed and weighted graphs. In *Algorithms and data structures (Ottawa, ON, 1991)*, volume 519 of *Lecture Notes in Comput. Sci.*, pages 461–472. Springer, Berlin, 1991.

- [Mat97] T. Matsui. A flexible algorithm for generating all the spanning trees in undirected graphs. *Algorithmica*, 18(4):530–543, 1997.
- [Min65] G. Minty. A simple algorithm for listing all the trees of a graph. *IEEE Trans. Circuit Theory*, 12(1):120, 1965.
- [MM24] A. Merino and T. Mütze. Traversing combinatorial 0/1-polytopes via optimization. *SIAM J. Comput.*, 53(5):1257–1292, 2024.
- [MMW22] A. Merino, T. Mütze, and A. Williams. All your bases are belong to us: listing all bases of a matroid by greedy exchanges. In *11th International Conference on Fun with Algorithms*, volume 226 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Paper No. 22, 28. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2022.
- [Müt23] T. Mütze. Combinatorial Gray codes—an updated survey. *Electron. J. Combin.*, DS26(Dynamic Surveys):99 pp., 2023.
- [NP81] D. J. Naddef and W. R. Pulleyblank. Hamiltonicity and combinatorial polyhedra. *J. Combin. Theory Ser. B*, 31(3):297–312, 1981.
- [NP84] D. J. Naddef and W. R. Pulleyblank. Hamiltonicity in (0-1)-polyhedra. *J. Combin. Theory Ser. B*, 37(1):41–52, 1984.
- [RR72] V. Rao and N. Raju. On tree graphs of directed graphs. *IEEE Trans. Circuit Theory*, 19(3):282–283, 1972.
- [RT75] R. C. Read and R. E. Tarjan. Bounds on backtrack algorithms for listing cycles, paths, and spanning trees. *Networks*, 5(3):237–252, 1975.
- [Sav97] C. D. Savage. A survey of combinatorial Gray codes. *SIAM Rev.*, 39(4):605–629, 1997.
- [Sch03] A. Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. B*, volume 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003. Matroids, trees, stable sets, Chapters 39–69.
- [Sha68] H. Shank. A note on Hamilton circuits in tree graphs. *IEEE Trans. Circuit Theory*, 15(1):86, 1968.
- [Sla77] P. J. Slater. Fibonacci numbers in the count of spanning trees. *Fibonacci Quart.*, 15(1):11–14, 1977.
- [Smi97] M. J. Smith. Generating spanning trees. Master’s thesis, University of Victoria, 1997.
- [ST95] A. Shioura and A. Tamura. Efficiently scanning all spanning trees of an undirected graph. *J. Oper. Res. Soc. Japan*, 38(3):331–344, 1995.
- [STU97] A. Shioura, A. Tamura, and T. Uno. An optimal algorithm for scanning all spanning trees of undirected graphs. *SIAM J. Comput.*, 26(3):678–692, 1997.