The Permutation-Path Coloring Problem on Trees

Dominique Barth[‡], Sylvie Corteel[†], Alain Denise[†], Daniele Gardy[‡], Mario Valencia-Pabon[†]

[†] L.R.I., Bât 490, Université Paris-Sud, 91405 ORSAY, France.
[‡] PRISM, Université de Versailles, 45 Av. des Etats Unis, 78035 VERSAILLES, France

Abstract

The paper deals with the problem of routing a set of communication requests representing a permutation of the nodes of an all-optical tree shaped network employing the *wavelength division multiplexing* (or WDM) technology. In such networks, information between nodes is transmitted as light on fiber-optic lines without being converted to electronic form in between, and different messages may use the same link concurrently if and only if they are assigned distinct wavelengths. Thus, the goal of the routing problem on these networks is to assign a wavelength to each communication request in order to minimize the number of wavelengths needed to perform all communications in only one round. Such a routing problem can be modeled as a permutation-path coloring problem on trees. An instance of the permutation-path coloring problem on trees is given by a directed symmetric tree graph T on n nodes and a permutation σ of the vertex set of T. Moreover, we associate with each pair $(i, \sigma(i))$, $i \neq \sigma(i)$, $1 \leq i \leq n$, the unique directed path on T from vertex i to vertex $\sigma(i)$. Thus, the permutation-path coloring problem for this instance consists in assigning the minimum number of colors to such a permutation-set of paths in such a way that any two paths sharing a same arc of the tree are assigned different colors. In fact, the colors in the latest problem represents the wavelengths in the former one. In this paper we first show that the permutation-path coloring problem is NP-hard even in the case of *involutions* (resp. *circular permutations*), that are permutations which contain only cycles of length at most two (resp. contain exactly one cycle), on both binary trees and on trees having only two vertices with degree greater than two. Next, we calculate a lower bound on the average complexity of the permutation-path coloring problem on arbitrary networks. Then we give combinatorial and asymptotic results for the permutation-path coloring problem on linear networks in order to show that the average number of colors needed to color any permutation on a linear network on n vertices is $n/4 + o(n)$. We extend these results and obtain an upper bound on the average complexity of the permutation-path coloring problem on arbitrary trees, obtaining exact results in the case of generalized star trees. Finally we explain how to extend these results for the involutions-path coloring problem on arbitrary trees.

Key words: Average-Case Complexity, Routing Permutation, Path Coloring, Tree Networks, NP-completeness.

1 Introduction

Efficient communication is a prerequisite to exploit the performance of large parallel systems. The routing problem on communication networks consists in the efficient allocation of resources to connection requests. In this network, establishing a connection between two nodes requires *selecting* a path connecting the two nodes and *allocating* sufficient resources on all links along the paths associated to the collection of requests. In the case of *all-optical* networks, data is transmitted on lightwaves through optical fiber, and several signals can be transmitted through a fiber link simultaneously provided that different wavelengths are used in order to prevent interference (wavelength-division multiplexing) [5]. As the number of wavelengths is a limited resource, then it is desirable to establish a given set of connection requests with a minimum number of wavelengths.

In this context, it is natural to think in wavelengths as colors. Thus the routing problem for all-optical networks can be viewed as a path coloring problem: it consists in finding a desirable collection of paths on the network associated with the collection of connection requests in order to minimize the number of colors needed to color these paths in such a way that any two different paths sharing a same link of the network are assigned different colors. For simple networks, such as trees, the routing problem is simpler, as there is always a unique path for each communication request.

This paper is concerned with the routing permutations on trees by arc-disjoint paths, that is, the path coloring problem on trees when the collection of connection requests represents a permutation of the nodes of the tree network.

Previous and related work. Using a result of Leighton and Rao [19], Aumann and Rabani [1] have shown that $O(\frac{\log^2 n}{\beta^2})$ colors suffices for routing any permutation on any bounded degree network on n nodes, where β is the *arc expansion* of the network. The result of Aumman and Rabani almost matches the existential lower bound of $\Omega(\frac{1}{\beta^2})$ obtained by Raghavan and Upfal [24]. In the case of specific network topologies, Gu and Tamaki [16] proved that 2 colors are sufficient to route any permutation on any symmetric directed hypercube. Independently, Paterson et al. [23] and Wilfong and Winkler [28] have shown that the routing permutation problem on ring networks is NP-hard. Moreover, in [28] a 2-approximation algorithm is given for this problem on ring networks. Independently, Kumar et al. [18] and Erlebach and Jansen [7] have shown that computing a minimal coloring of any collection of paths on binary trees is NP-hard. Caragiannis et al. [4] consider the symmetric-path coloring problem on trees (i.e., for each path from vertex u to vertex v, there also exists its symmetric, a path from vertex v to vertex u) showing that this special instance is also

NP-hard for unbounded degree trees, and leaving as an open problem the complexity of such a symmetric instances on binary trees. To our knowledge, the routing permutation problem on tree networks by arc-disjoint paths has not been studied in the literature.

Our results. In Section 3 we show that the symmetric-path coloring problem on binary trees is NP-hard, answering to the open question asked in [4]. Moreover, we extend such a result in order to show that the permutation-path coloring problem remains NP-hard even in the case of *involutions* (resp. *circular-permutations*), that are permutations which contain only cycles of length at most two (resp. contain exactly one cycle), on both binary trees and on trees having only two vertices with degree greater than two. In Section 4 we compute a lower bound for the average number of colors needed to color any permutation-path set on arbitrary networks. In Section 5 we focus on linear networks. In this particular case, since the problem reduces to coloring an interval graph, the routing of any permutation is easily done in polynomial time [17]. We show that the average number of colors needed to color any permutation-path set on a linear network on n vertices is $n/4 + o(n)$. In Section 6, we extend the results obtained in Section 5, by giving an upper bound on the average number of colors needed to color any permutation-path set on arbitrary trees, obtaining exact results in the case of generalized star tree networks. As far as we know, this is the first result on the average-case complexity for routing permutations on all-optical networks. We finally show how to extend these results to the involution problem partly studied in [?]. We begin in Section 2 with the preliminaries.

2 Definitions and preliminary results

We model the tree network as a rooted labeled symmetric directed tree $T = (V, A)$ on n vertices, where processors and switches are vertices and links are modeled by two arcs in opposite directions. Let $\mathcal{P}$ be a collection of directed paths on T. We assume that the vertices of T are arbitrarily labeled by different integers in $\{1, 2, \ldots, n\}$. We denote by $i \rightsquigarrow j$ the unique directed path from vertex i to vertex j in T. The arc from vertex i to its father (resp. from the father of i to i), $1 \leq i \leq n-1$, is labeled by i^+ (resp. i^-). We call $T(i)$ the subtree of T rooted at vertex i, $1 \leq i \leq n$.The vertex labeled with the integer n is the root vertex of T. See Figure 1(a) for the linear network on $n = 6$ vertices rooted at vertex $i = 6$. Note that we will just draw an edge i rather than the arcs i^+ and i^- in the sequel.

For any i, $1 \leq i \leq n-1$, the *load* of an arc i^+ (resp. i^-) of T, denoted by $L_T(\mathcal{P}, i^+)$ (resp. $L_T(\mathcal{P}, i^-)$), is the number of paths in $\mathcal{P}$ using such an arc, and the *maximum load* among all arcs of T is denoted by $L_T(\mathcal{P})$. We call *the coloration number* and we denote by $R_T(\mathcal{P})$, the minimum number of colors needed to color the paths in $\mathcal{P}$ such that any two paths sharing a same arc in T are assigned different colors. Trivially, we have that $R_T(\mathcal{P}) \geq L_T(\mathcal{P})$.

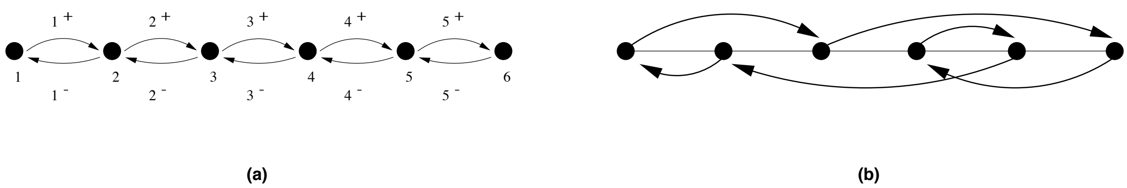

Figure 1: (a) Labeling of the vertices and the arcs for the linear network on $n = 6$ vertices rooted at vertex $i = 6$. (b) Representation of the permutation $\sigma = (3, 1, 6, 5, 2, 4)$ on the linear network given in (a).

Let $\mathcal{S}_n$ denotes the symmetric group of all permutations on $[n] = \{1, 2, \dots, n\}$. Let σ be a permutation in $\mathcal{S}_n$, then σ is called *involution* (resp. *circular*) if it contains only cycles of length at most two (resp. contains exactly one cycle of length n). Let I_{2n} be the set of involutions with no fixed point on $[2n]$. We say that $\mathcal{P}$ is a *permutation-path* set on T if $\mathcal{P}$ represents a permutation $\sigma \in \mathcal{S}_n$ of the vertex-set of T, where $\sigma(i) = j$ if and only if $i \rightsquigarrow j \in \mathcal{P}$. In the sequel we talk indifferently of a permutation-path set $\mathcal{P}$ or of the permutation $\sigma \in \mathcal{S}_n$ that $\mathcal{P}$ represents. Thus, given a permutation $\sigma \in \mathcal{S}_n$ and a tree T on n vertices, the load of the arc i^+ (resp. i^-), $1 \leq i \leq n-1$, can be expressed by $L_T(\sigma, i^+) = |\{j \in T(i) : \sigma(j) \notin T(i)\}|$ (resp. $L_T(\sigma, i^-) = |\{j \notin T(i) : \sigma(j) \in T(i)\}|$).

Lemma 1 *Let T be a tree on n vertices. For all $\sigma \in \mathcal{S}_n$ and for all $i \in \{1, 2, \dots, n-1\}$, $L_T(\sigma, i^+) = L_T(\sigma, i^-)$. Therefore, $L_T(\sigma) = \max\limits_i L_T(\sigma, i^+)$.*

Proof. We can prove this by induction. If vertex i is a leaf, we have two cases:

- $\sigma(i) = i$, then $L_T(\sigma, i^+) = L_T(\sigma, i^-) = 0$,
- $\sigma(i) \neq i$, then $L_T(\sigma, i^+) = L_T(\sigma, i^-) = 1$.

Otherwise, vertex i is an internal node. The result still holds for $i > 1$. Let $\{i_1, i_2, \dots, i_j\}$ be the sons of the vertex i, and $N(i_j)$ be the number of vertices $k \in T(i_j)$ such that $\sigma(k) \notin T(i_j)$ and $\sigma(k) \in T(i)$. Then it is easy to see that $L_T(\sigma, i^+)$ and $L_T(\sigma, i^-)$ satisfy the same recurrence relation for $i > 1$:

$$L_T(\sigma, i^{\pm}) = \begin{cases} \sum_{k=1}^{j} L_T(\sigma, i_k^{\pm}) - N(i_k) & \text{if } \sigma(i) = i \text{ or } \sigma(i) \in T(i) \text{ and } \sigma^{-1}(i) \notin T(i) \\ & \text{or } \sigma(i) \notin T(i) \text{ and } \sigma^{-1}(i) \in T(i) \\ 1 + \sum_{k=1}^{j} L_T(\sigma, i_k^{\pm}) - N(i_k) & \text{if } \sigma(i) \notin T(i) \text{ and } \sigma^{-1}(i) \notin T(i) \\ -1 + \sum_{k=1}^{j} L_T(\sigma, i_k^{\pm}) - N(i_k) & \text{if } \sigma(i) \in T(i) \text{ and } \sigma^{-1}(i) \in T(i) \end{cases} \tag{1}$$

□

This lemma tells us that in order to study the load of a permutation on a tree on n vertices, it suffices to consider only the load of the labeled arcs i^+. For example the permutation $\sigma = (3, 1, 6, 5, 2, 4)$ on the linear network in Figure 1(b) has load 2. The maximum is reached in the arcs $4^{\pm}$.

Definition 1 *Let T be a tree on n vertices. The average load of all permutations $\sigma \in \mathcal{S}_n$ on T, denoted by $\bar{L}_T$, is defined as* $\bar{L}_T = \frac{1}{n!} \sum_{\sigma \in \mathcal{S}_n} L_T(\sigma)$.

Proposition 1 [8] *There is a polynomial algorithm to color any collection $\mathcal{P}$ of paths on any tree T such that $L_T(\mathcal{P}) \leq R_T(\mathcal{P}) \leq \lceil \frac{5}{3} L_T(\mathcal{P}) \rceil$.*

Given a tree T on n vertices, we denote by $\bar{R}_T$ the average number of colors needed to color all permutations in $\mathcal{S}_n$ on T. Thus, by Proposition 1, we have the following lemma.

Lemma 2 *Let T be a tree on n vertices. Then $\bar{L}_T \leq \bar{R}_T \leq \frac{5}{3}\bar{L}_T + 1$.*

Proposition 2 [4] *There is a polynomial algorithm to color any collection $\mathcal{P}$ of symmetric paths on any tree T such that $L_T(\mathcal{P}) \leq R_T(\mathcal{P}) \leq \lfloor \frac{3}{2} L_T(\mathcal{P}) \rfloor$.*

Given a tree T on $2n$ vertices, we denote by $\tilde{R}_T$ the average number of colors needed to color all involutions in I_{2n} on T. Thus, by Proposition 2, we have the following lemma.

Lemma 3 *Let T be a tree on $2n$ vertices and let $\tilde{L}_T$ be the average load of all involutions in I_{2n} on T. Then $\tilde{L}_T \leq \tilde{R}_T \leq \frac{3}{2}\tilde{L}_T$.*

Definition 2 *Let T be a tree and let $\mathcal{P}$ be a collection of paths on T. We call conflict graph, and we denote by $G_T(\mathcal{P}) = (V, E)$, the undirected graph associated with T and $\mathcal{P}$, where each vertex $v_p \in V$ represents a path $p \in \mathcal{P}$, and two vertices v_p and v_q are joined by an edge in E if and only if its associated paths p and q respectively, share a same arc dans T.*

It is straightforward to see that the coloration number $R_T(\mathcal{P})$ is equal to the chromatic number of the conflict graph $G_T(\mathcal{P})$.

Definition 3 *Let T be a tree and let $\mathcal{P}$ be a collection of paths on T. The digraph associated with $\mathcal{P}$, denoted $\vec{G}_T(\mathcal{P})$, is the digraph with vertex set V', where $v \in V'$ if and only if v is a vertex of T and there is at least one path in $\mathcal{P}$ having v as ending-vertex, and with arc set $A' = \{(v, w) : v, w \in V'$ and $v \rightsquigarrow w \in \mathcal{P}\}$.*

A digraph $G = (V, A)$ is said *pseudo-symmetric*, if for any vertex $v \in V$, $d^+(v) = d^-(v)$, where $d^+(v)$ (resp. $d^-(v)$) denotes the in-degree (resp. out-degree) of vertex v.

Theorem 1 [11] *If G is a connected pseudo-symmetric digraph, then G is Eulerian and an Eulerian circuit of G can be found in linear time.*

Let P_n denotes the directed symmetric path graph on n vertices. Let $\mathrm{ST}(n)$ denotes the directed symmetric star graph on n vertices (i.e., the tree having only one internal vertex connected to $n-1$ leaves). Let $\mathrm{GST}(\lambda)$ denotes the *generalized star* graph on n vertices, where $\lambda = (\lambda_1, \dots, \lambda_k)$ is a partition of the integer $n-1$ into k parts ($k > 2$). In fact, the graph $\mathrm{GST}(\lambda)$ has k branches connected to each other by one vertex, where λ_i denotes the length of the i^{th} branch (i.e., a branch of length λ_i is a path graph on $\lambda_i + 1$ vertices).

3 Complexity of computing the coloration number

We begin this section by showing the NP-hardness of the symmetric-path coloring problem on binary trees, answering to an open question asked in [4]. Moreover, we extend such a result by showing that the permutation-path coloring problem remains NP-hard even in the case of *involutions* (resp. *circular-permutations*), that are permutations which contain only cycles of length at most two (resp. contain exactly one cycle), on both binary trees and on trees having only two vertices with degree greater than two. Finally, we discuss some polynomial cases of this problem.

3.1 NP-hardness results

This section shows that the path coloring problem on trees is difficult even for very restrictive cases. For this, we use a similar reduction with the one used in [7, 18] for proving the NP-hardness of the general path coloring problem on binary trees. We remark that the reduction used in [7, 18] can not be directly extended to obtain NP-hardness results on the restrictive instances of the problem that we consider in the following theorem.

Theorem 2 *Let T be a directed symmetric tree and let $\mathcal{P}$ be a collection of directed paths on T. Then, computing $R_T(\mathcal{P})$ is NP-hard in the following cases:*
(a) T is a binary tree and $\mathcal{P}$ is a collection of symmetric paths on T.
(b) T is a binary tree and $\mathcal{P}$ represents an involution of the vertices of T.
(c) T is a binary tree and $\mathcal{P}$ represents a circular-permutation of the vertices of T.
(d) T is a tree having only two vertices with degree greater than two and $\mathcal{P}$ represents an involution or a circular-permutation of the vertices of T.

Proof. We use a reduction from the ARC-COLORING problem [25]. The ARC-COLORING problem can be defined as follows : given a positive integer k, an undirected cycle C_n with vertex set numbered clockwise as $1, 2, \dots, n$, and any collection of paths F on C_n, where each path in F from vertex v to vertex w, denoted by $< v, w >$, is regarded as the path beginning at vertex v and ending at vertex w again in the clockwise direction, does F can be colored with k colors such that no two paths sharing an edge of C_n are assigned the same color ? It is well known that the ARC-COLORING problem is NP-complete [13]. W.l.o.g., we assume that each edge of C_n is traversed by exactly k paths in F. If some edge $[i, i+1]$ of C_n is traversed by $r < k$ paths, then we can add $k - r$ paths of the form $< i, i+1 >$ (or $< i, 1 >$ if $i = n$) to F without changing its k-colorability. We assume that no path in F covers entirely the cycle C_n. Let I be an instance of the ARC-COLORING problem defined as above. We construct from I an instance I' consisting of a symmetric directed tree T and a collection of paths $\mathcal{P}$ on T such I' verifies the constraints giving in (a) (resp. (b), (c), (d)), and such that F is k-colorable if and only if $\mathcal{P}$ is k'-colorable, for some integer $k' \geq k$. Let $< i, j >$ be any path in F, thus we say that $< i, j >$ is of *type 1* (resp. *type 2*) if $i < j$ (resp. $i > j$).

(a) T is constructed as follows : first, construct a graph on n vertices isomorphic to path graph P_n and denote its vertices by $v_1, v_2, \ldots, v_n$. Next, construct $2(n+k)$ different isomorphic graphs to star graph ST(4). Take $n+k$ of such $2(n+k)$ isomorphic graphs and denote theirs leaves by l_i, s_i and t_i, and denote the leaves of the $n+k$ other ones by r_i, x_i and z_i, $1 \leq i \leq n+k$. Finally, connect the vertex l_i (resp. r_i) to vertex l_{i+1} (resp. r_{i+1}), $1 \leq i \leq n+k-1$, and connect the vertex l_1 (resp. r_1) to vertex v_1 (resp. v_n) of P_n.
$\mathcal{P}$ is constructed as follows : for each path $< i, j > \in F$, if $< i, j >$ is of type 1 (i.e. $i < j$), then add to $\mathcal{P}$ the paths $A_{i,j} = v_i \rightsquigarrow v_j$ and $B_{j,i} = v_j \rightsquigarrow v_i$. Otherwise, if $< i, j >$ is of type 2 (i.e. $i > j$), then let p (resp. q) be an integer in $\{1, 2, \ldots, k\}$ such that no path in $\mathcal{P}$ uses the vertices x_p and z_p (resp. s_q and t_q) as ending vertices. So, add to $\mathcal{P}$ the path sets $\bar{A}_{i,j} = \{v_i \rightsquigarrow z_p \ , x_p \rightsquigarrow t_q \ , s_q \rightsquigarrow v_j\}$ and $\bar{B}_{j,i} = \{v_j \rightsquigarrow s_q \ , t_q \rightsquigarrow x_p \ , z_p \rightsquigarrow v_i\}$.
In order to make sure that the digraph associated to the collection of paths (see Def. 3) be connected (property that will be used to prove Part (c)), for each j, $k+1 \leq j \leq n+k$, we add to $\mathcal{P}$ the sets of paths $C_{j-k} = \{s_j \rightsquigarrow v_{j-k}, v_{j-k} \rightsquigarrow z_{j'}, x_{j'} \rightsquigarrow v_{j'-k}, v_{j'-k} \rightsquigarrow t_j\}$ and $D_{j-k} = \{v_{j-k} \rightsquigarrow s_j, z_{j'} \rightsquigarrow v_{j-k}, v_{j'-k} \rightsquigarrow x_{j'}, t_j \rightsquigarrow v_{j'-k}\}$, where $j' = j+1$ if $j < n+k$, otherwise $j' = k+1$.
In addition, for each i, $1 \leq i \leq n+k$, we add to $\mathcal{P}$ $2(n+k)-1$ identical paths $s_i \rightsquigarrow t_i$ (resp. $x_i \rightsquigarrow z_i$) and $2(n+k)-1$ identical paths $t_i \rightsquigarrow s_i$ (resp. $z_i \rightsquigarrow x_i$). Finally, set $k' = 2(n+k)$. In Figure 2 we present an example of this polynomial construction. By construction is easy to see that $\mathcal{P}$ is a collection of symmetric paths on T. Moreover, let $< i_1, j_1 >, < i_2, j_2 >, \ldots, < i_k, j_k >$ be the k paths of type 2 in F, and let $\bar{A}_{i_r,j_r} = \{v_{i_r} \rightsquigarrow z_{p_r} \ , x_{p_r} \rightsquigarrow t_{q_r} \ , s_{q_r} \rightsquigarrow v_{j_r}\}$ and $\bar{B}_{j_r,i_r} = \{v_{j_r} \rightsquigarrow s_{q_r} \ , t_{q_r} \rightsquigarrow x_{p_r} \ , z_{p_r} \rightsquigarrow v_{i_r}\}$ be the two sets of paths in $\mathcal{P}$ associated with the path $< i_r, j_r >$, $1 \leq r \leq k$. Then $\mathcal{P}$ verifies the following properties:
P1. All the paths in each one of the sets $\bar{A}_{i_r,j_r}$, $\bar{B}_{j_r,i_r}$, C_m, and D_m, $1 \leq r \leq k$, $1 \leq m \leq n$, are colored with the same color in any k'-coloring of $\mathcal{P}$. This is done by the $k'-1$ identical paths $s_i \rightsquigarrow t_i$ (resp. $t_i \rightsquigarrow s_i$) and the $k'-1$ identical paths $x_i \rightsquigarrow z_i$ (resp. $z_i \rightsquigarrow x_i$), $1 \leq i \leq n+k$, which make sure that all the paths in each one of these sets are colored with

Figure 2: Partial construction of I' from I, where $n = 4$ and $k = 2$.

the same color in any k'-coloring of $\mathcal{P}$.
P2. Each one of the sets $\bar{A}_{i_r,j_r}$, $\bar{B}_{j_r,i_r}$, C_m, and D_m, $1 \leq r \leq k$, $1 \leq m \leq n$, should be assigned a different color in any k'-coloring of $\mathcal{P}$. In fact, it is easy to see that by construction, each path $x_{p_r} \rightsquigarrow t_{q_r} \in \bar{A}_{i_r,j_r}$ (resp. $t_{q_r} \rightsquigarrow x_{p_r} \in \bar{B}_{j_r,i_r}$), $1 \leq r \leq k$, intersects with all the paths in $\cup_{m=1}^{k}\{x_{p_m} \rightsquigarrow t_{q_m} \in \bar{A}_{i_m,j_m} : m \neq r\}$ (resp. in $\cup_{m=1}^{k}\{t_{q_m} \rightsquigarrow x_{p_m} \in \bar{B}_{j_m,i_m} : m \neq r\}$), and with all the paths in $\cup_{m=1}^{k}\{z_{p_m} \rightsquigarrow v_{i_m}\ ,\ v_{j_m} \rightsquigarrow s_{q_m} \in \bar{B}_{j_m,i_m}\}$ (resp. in $\cup_{m=1}^{k}\{v_{i_m} \rightsquigarrow z_{p_m}\ ,\ s_{q_m} \rightsquigarrow v_{j_m} \in \bar{A}_{i_m,j_m}\}$). Moreover, each set of paths C_m (resp. D_m) intersects with all the paths in $\mathcal{P} \setminus C_m \setminus Q$ (resp. in $\mathcal{P} \setminus D_m \setminus Q$), where Q is the collection of all the $k'-1$ identical paths $s_i \rightsquigarrow t_i$ (resp. $t_i \rightsquigarrow s_i$) and the $k'-1$ identical paths $x_i \rightsquigarrow z_i$ (resp. $z_i \rightsquigarrow x_i$), $1 \leq i \leq k$. Therefore, by Property P1, this property follows.
P3. Each path $A_{a,b}$ (resp. $B_{b,a}$) in $\mathcal{P}$ associated with a path $< a, b >$ in F of type 1, intersects with all the paths in $\cup_{r=1}^{k}\{t_{q_r} \rightsquigarrow x_{p_r} \in \bar{B}_{j_r,i_r}\}$ (resp. in $\cup_{r=1}^{k}\{x_{p_r} \rightsquigarrow t_{q_r} \in \bar{A}_{i_r,j_r}\}$), and by Property P2, with at least one of the paths in each one of the sets C_m and D_m, $1 \leq m \leq n$.

Now we claim that there is a k-coloring of F if and only if there is a k'-coloring of $\mathcal{P}$. Assume that there is a k-coloring of F, and let $< i, j >$ be any path in F colored with the color γ, $1 \leq \gamma \leq k$. Thus a k'-coloring of $\mathcal{P}$ can be carried out as follows: if $< i, j >$ is of type 1, then we color the paths $A_{i,j} = v_i \rightsquigarrow v_j$ and $B_{j,i} = v_j \rightsquigarrow v_i$ in $\mathcal{P}$ with colors γ and $\gamma + k$ respectively. Otherwise, if $< i, j >$ is of type 2, then we color all its three associated paths in $\bar{A}_{i,j}$ (resp. in $\bar{B}_{j,i}$) with color γ (resp. $\gamma + k$). Next, for each i, $1 \leq i \leq n$, we assign to all the paths in the set C_i (resp. D_i) the color $2k + i$ (resp. $2k + n + i$). Finally, for each i, $1 \leq i \leq n + k$, we color the $k'-1$ identical paths $s_i \rightsquigarrow t_i$ (resp. $x_i \rightsquigarrow z_i$) and the $k'-1$ identical paths $t_i \rightsquigarrow s_i$ (resp. $z_i \rightsquigarrow x_i$) with the $k'-1$ available colors for each one of these $(k'-1)$-set of paths. Thus, by Properties P1, P2, and P3, it is easy to see that such a coloring is a proper k'-coloring of $\mathcal{P}$.
Conversely, assume that there is a k'-coloring of $\mathcal{P}$. By Properties P1, P2, and P3, it is easy to deduce two proper k-colorings for F as follows: if $< i, j >$ is a path in F of type 1, we assign to $< i, j >$ the color assigned to path $A_{i,j} = v_i \rightsquigarrow v_j$ (resp. $B_{j,i} = v_j \rightsquigarrow v_i$) in $\mathcal{P}$. Otherwise, if $< i, j >$ is a path in F of type 2, we assign to $< i, j >$ the same color assigned to the three paths in the set $\bar{A}_{i,j}$ (resp. $\bar{B}_{j,i}$). Thus, F is k-colorable if and only if $\mathcal{P}$ is k'-colorable which ends the proof of (a).

(b) It follows from (a). In fact, let T be the binary tree and $\mathcal{P}$ be the symmetric collection of paths constructed in Part (a). Let u and v be two adjacent vertices in T, and let $i(u, v)$ (resp. $o(u, v)$) be the set of paths traversing the arc (u, v) (resp. (v, u)) and having as final-vertex (resp. initial-vertex) the vertex v. As $\mathcal{P}$ is symmetric, it is clear that $|i(u, v)| = |o(u, v)|$. Then, replace the pair of arcs (u, v) and (v, u) by a path graph on $|i(u, v)| = \alpha$ vertices. Let P_α denote such a path graph, and $w_1, w_2, \ldots, w_\alpha$ denote its vertices. Replace each pair of symmetric paths $a \rightsquigarrow v \in i(u, v)$ and $v \rightsquigarrow a \in o(u, v)$ by the paths $a \rightsquigarrow w_j$ and $w_j \rightsquigarrow a$, where w_j is a vertex of P_α not yet used by any path as initial or final vertex. Using the previous transformation on each pair of adjacent vertices of T, we obtain an instance giving by an extended binary tree T' and a set of paths $\mathcal{P}$ which represents an involution of the vertices of T', which is equivalent (from the coloring point of view) to the one obtained in Part (a), ending the proof of (b).

(c) Let T be the binary tree and $\mathcal{P}$ be the symmetric collection of paths constructed in Part (a). Clearly, the digraph $\vec{G}_T(\mathcal{P})$ associated with $\mathcal{P}$ (see Def. 3) is a connected pseudo-symmetric digraph. By using a similar procedure as in Part (b), we can transform T and $\mathcal{P}$ in such a way that each vertex v_i in T (recall that vertex v_i belongs to the initial path graph P_n constructed in Part (a)), $1 \leq i \leq n$, be the initial and final vertex of exactly two paths in $\mathcal{P}$. However, we should be careful in order to maintain the connectness of the digraph associated with $\mathcal{P}$. For this, for each vertex v_i, $1 \leq i \leq n$, if u is an adjacent vertex to vertex v_i, and the pairs of arcs (u, v_i) and (v_i, u) should be replaced by a new path graph on α vertices denoted by $w_1, w_2, \ldots, w_\alpha$ (see Part (b)), where w_1 (resp. w_α) will be the new adjacent vertex to v_i (resp. u), then after this replacement we should add to $\mathcal{P}$ the paths $w_j \rightsquigarrow w_{j+1}$ and $w_{j+1} \rightsquigarrow w_j$, $1 \leq j < \alpha$, and the paths $w_1 \rightsquigarrow v_i$ and $v_i \rightsquigarrow w_1$. It is not difficult to see that this new instance is equivalent (from the coloring point of view) to the previous one. Let T' and $\mathcal{P}'$ denote the current tree and the current symmetric collection of paths respectively, after the previous transformation. Then, the digraph $\vec{G}_{T'}(\mathcal{P}')$ is a connected pseudo-symmetric digraph, and by Theorem 1, $\vec{G}_{T'}(\mathcal{P}')$ is Eulerian and an Eulerian circuit of it can be found in polynomial-time. Let $a_1, a_2, \ldots, a_\rho, a_1$ be an Eulerian circuit of $\vec{G}_{T'}(\mathcal{P}')$, where $\rho = |\mathcal{P}'|$. Note that each pair (a_i, a_{i+1}) in the Eulerian circuit represents a path $a_i \rightsquigarrow a_{i+1}$ of $\mathcal{P}'$. Moreover, let $w_1, w_2, \ldots, w_{n'}$ denote the vertices of the current path graph in T', where w_1 is adjacent to vertex l_1, and vertex $w_{n'}$ is adjacent to vertex r_1. By previous construction, each vertex w_i must be twice on the Eulerian circuit. Thus, following the Eulerian circuit in the order $a_1, a_2, \ldots, a_\rho, a_1$, for each vertex w_i, $1 \leq i \leq n'$, if w_i is found for the second time on the Eulerian circuit, we add a new vertex u_i to T' and connect it to vertex w_i, and we replace the paths $\beta \rightsquigarrow w_i$ and $w_i \rightsquigarrow \gamma$ in $\mathcal{P}'$, by the paths $\beta \rightsquigarrow u_i$ and $u_i \rightsquigarrow \gamma$, where β and γ are the immediately predecessor and successor of w_i respectively, on the Eulerian circuit. Indeed, by construction, each one of the vertices s_i, t_i, x_i, and z_i, $1 \leq i \leq n+k$, is found exactly k' times on the Eulerian circuit, and giving that each one of these vertices is a leaf of T', we can replace each one of these vertices by a new path on k' vertices and arrange the k' paths in $\mathcal{P}'$ ending and beginning in each one of these in agreement with the Eulerian circuit. Therefore, it is easy to prove that the obtained tree T' is binary and that the set of paths $\mathcal{P}'$ represents a circular-permutation of the vertices of T'. Thus, taking care of the initial paths $A_{i,j}$ and $B_{j,i}$ (resp. set of paths $\bar{A}_{i,j}$ and $\bar{B}_{j,i}$) associated with paths $< i, j >$ in F of type 1 (resp. 2), we obtain that the final circular-permutation set of paths $\mathcal{P}'$ on T' is k'-colorable if and only if F is k-colorable, which ends the proof of (c).

(d) The involution case follows from (a) and (b). In fact, let T be the binary tree and $\mathcal{P}$ be the symmetric collection of paths constructed in Part (a). Replace all the $n+k$ isomorphic star graphs whose leaves are labeled by l_i, s_i, and t_i (resp. r_i, x_i, and z_i) by only one star $\mathrm{ST}(2(n+k)+1)$ having as leaves the vertices s_i and t_i (resp. x_i and z_i), $1 \leq i \leq n+k$, and denote by l_1 (resp. r_1) its only vertex of degree $2(n+k)$. Next, connect the vertex l_1 (resp. r_1) to vertex v_1 (resp. v_n) of P_n, leaving $\mathcal{P}$ as in (a). Thus, it is easy to see that this new instance is equivalent to the one obtained in Part (a) (from the coloring point of view). Finally, using similar arguments as in Part (b), we proof the NP-hardness for the involution case. The circular-permutation case follows by similar arguments from Part (c). This ends the proof of (d) and the theorem is proved. □

By Proposition 1 (resp. Proposition 2), the best known approximation algorithm for coloring any collection of paths (resp. symmetrique paths) with load L on any tree network uses at most $\lceil \frac{5}{3}L \rceil$ (resp. $\lfloor \frac{3}{2}L \rfloor$) colors. Therefore it trivially also holds for any permutation-set (resp. involution-set) of paths with load L on any tree.

3.2 Some polynomial cases

Let $\mathcal{P}$ be any collection of paths on a tree network T. If T is a linear network then, the minimum number of colors $R_T(\mathcal{P})$ needed to color the paths in $\mathcal{P}$ is equal to the load $L_T(\mathcal{P})$ induced by $\mathcal{P}$. In fact, if T is a linear network then, the conflict graph of the paths in $\mathcal{P}$ is an *interval graph* (see [15]). Moreover, optimal vertex coloring for interval graphs can be computed efficiently [17]. When T is a star network, the equality between $R_T(\mathcal{P})$ and $L_T(\mathcal{P})$ also holds because the path coloring problem on these graphs is equivalent to find a minimum edge-coloring of an undirected bipartite graph. Moreover, the minimum number of colors needed to color the edges of a bipartite graph is equal to its maximum degree, and such an edge-coloring in these graphs can be found in polynomial time [3]. Combining these approaches for linear and star networks, Gargano et al. [14] show that if T is a generalized star network then, an optimal coloring of $\mathcal{P}$ on these networks can be computed efficiently in polynomial time, and that the equality between $R_T(\mathcal{P})$ and $L_T(\mathcal{P})$ also holds. Note that all the results in these three networks trivially hold when $\mathcal{P}$ is a permutation-set of paths. However, by Theorem 2, it suffice that the tree network T has two vertices with degree greater that two and the permutation-path coloring problem on these networks becomes NP-hard. Moreover, in binary tree networks having only two vertices with degree equal to 3, the equality between the load and the minimum number of colors for a permutation-path set does not always hold as we can see in Figure 3.

Figure 3(a) shows an example of a permutation $\sigma \in S_{10}$ to be routed on a tree T on 10 vertices, which load $L_T(\sigma)$ is equal to 2. Figure 3(b) shows the conflict graph $G = G_T(\sigma)$. Thus, clearly $R_T(\sigma)$ is equal to the chromatic number of G. Therefore, as the conflict graph

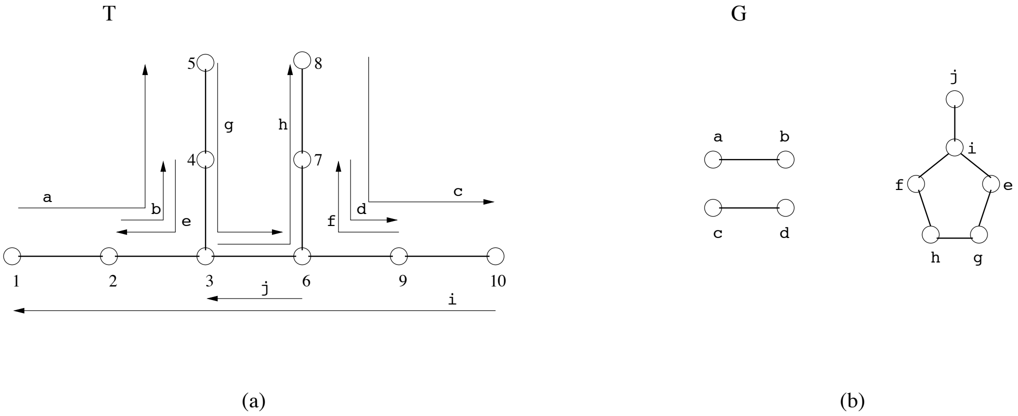

Figure 3: (a) A tree T on 10 vertices and a permutation $\sigma = (5, 4, 8, 2, 6, 3, 9, 10, 7, 1)$ to be routed on T. (b) The conflict graph G associated with permutation σ in (a).

G has the cycle C_5 as induced subgraph, then the chromatic number of G is equal to 3, and thus $R_T(\sigma) = 3$.

4 A lower bound for the average coloration number

We derive a lower bound for the average coloration number of permutations to be routed on arbitrary networks, by giving a lower bound for the average load of permutations to be routed on these ones. Let $G = (V, A)$ be a directed symmetric graph on n vertices (i.e. $|V| = n$) and r a routing function in G which assigns a set of paths on G to route any permutation $\sigma \in \mathcal{S}_n$. Let $\overline{L}_{G,r}$ be the average load of all permutations in $\mathcal{S}_n$ induced by the routing function r, and let $U \subseteq V$ be a subset of the vertex set of G. We denote by $c(U)$ the cut $(U, \bar{U})$, i.e., the set of arcs $\{(u, v) \in A : \ u \in U, \ v \in V \setminus U\}$.

Proposition 3 *For any graph $G = (V, A)$ on n vertices, and any routing function r in G,*

$$\overline{L}_{G,r} \geq \frac{1}{n} \cdot \max_{U \subseteq V} \left(\frac{|U| \cdot (n - |U|)}{|c(U)|} \right) .$$

Proof. Let $U \subseteq V$ be any subset of vertices in G and consider a permutation $\sigma \in \mathcal{S}_n$ to be routed on G by using the routing function r. The load of all the arcs in $c(U)$ induced by σ with the routing function r is defined by $L_r(U, \sigma) = |\{j \in U : \ \sigma(j) \notin U\}|$. The global load of $c(U)$ is then defined by $L_r(U) = \sum_{\sigma \in \mathcal{S}_n} L_r(U, \sigma)$. Thus, for any vertex $j \in U$ and for any vertex $k \in V \backslash U$, each permutation $\sigma \in \mathcal{S}_n$ such that $\sigma(j) = k$ contributes one unit to the global load of $c(U)$. So, the average load of $c(U)$ verifies $\overline{L}_r(U) \geq \frac{1}{n!} \sum_{j \in U} \sum_{k \in V \backslash U} |\{\sigma \in \mathcal{S}_n : \ \sigma(j) = k\}|$. Moreover, for all pair of vertices j and k in G, there exists $(n-1)!$ permutations $\sigma \in \mathcal{S}_n$

such that $\sigma(j) = k$. Therefore, $\overline{L}_r(U) \geq \frac{1}{n!} \sum_{j\in U} \sum_{k\in V\setminus U} (n-1)! = \frac{1}{n} \sum_{j\in U} \sum_{k\in V\setminus U} 1 = \frac{|U|(n-|U|)}{n}$. Thus, for any arc $\alpha \in c(U)$, the average load of α verifies $\overline{L}_r(\alpha) \geq \frac{\overline{L}_r(U)}{|c(U)|}$. So, $\overline{L}_{G,r} \geq \frac{1}{n} \cdot \max_{U\subseteq V} \left(\frac{|U|\cdot(n-|U|)}{|c(U)|}\right)$. □

Let us denote by $\mathcal{C}(G)$ the parameter $\max_{U\subseteq V} \left(\frac{|U|\cdot(|V|-|U|)}{|c(U)|}\right)$. It is not difficult to see that $\mathcal{C}(G)$ is equal to $\frac{1}{\mathcal{S}(G)}$, where $\mathcal{S}(G)$ denotes the *sparsest cut* of graph G. In fact, $\mathcal{S}(G)$ is defined by $\min_{U\subseteq V} \left(\frac{|c(U)|}{|U|\cdot(|V|-|U|)}\right)$. In [21] is shown that computing the sparsest cut of a graph is NP-hard, which implies that computing the parameter $\mathcal{C}(G)$ is also NP-hard and so, computing the lower bound given in Proposition 3 is NP-hard. Therefore, the method induced by the proof of this proposition can be used to give some lower bounds (but not necessarily the best one). However, for any constant c and given a graph $G = (V, A)$, if the graph bisection of G, denoted $\mathrm{Bis}(G)$, is at most equal to c, then to know whether $\mathcal{C}(G) \leq k$ can be computed efficiently : it is sufficient to know if the bound can be better than $\frac{(\lceil\frac{|V(G)|}{2}\rceil\cdot\lfloor\frac{|V(G)|}{2}\rfloor)}{c}$ by considering all the subsets of $A(G)$ disconnecting G with maximal cardinality c, i.e., a polynomial number of such subsets. For example, for any $2d$-mesh $M(2n, c)$ with $2n$ lines and a constant number c of columns, $\mathcal{C}(M(2n, c)) = c \cdot n^2$. For any ring C_n, $\mathcal{C}(C_n) = \frac{\lceil n/2\rceil\cdot\lfloor n/2\rfloor}{2}$. We end this section by giving a lower bound on the average number of colors needed to color any permutation $\sigma \in \mathcal{S}_n$ on any tree on n vertices as follows.

Let T be a tree on n vertices. By Proposition 3, we can deduce that the average load of any arc i^+ of T, $1 \leq i \leq n-1$, denoted by $\bar{L}_T(i)$, verifies $\bar{L}_T(i) = \dfrac{|T(i)|(n-|T(i)|)}{n}$. Moreover, for any vertex i of T, let $v_T(i) = |T(i)|/n$ and $\tilde{v}_T(i) = \min(v_T(i), 1 - v_T(i))$. Let $\tilde{v}_T = \max_i \tilde{v}_T(i)$. Then, it is clear that $\max_i\{\bar{L}_T(i)\} = n\tilde{v}_T(1-\tilde{v}_T)$. Indeed, it is straightforward that $\bar{L}_T \geq \max_i\{\bar{L}_T(i)\}$. Therefore, we obtain a lower bound for the average load $\bar{L}_T$.

Lemma 4 $\bar{L}_T \geq n\tilde{v}_T(1-\tilde{v}_T)$.

Moreover, as $\bar{R}_T \geq \bar{L}_T$, we obtain the following lower bound on the average number of colors needed to color any permutation $\sigma \in \mathcal{S}_n$ on any tree T on n vertices.

Lemma 5 $\bar{R}_T \geq n\tilde{v}_T(1-\tilde{v}_T)$.

5 Average coloration number on linear networks

The main result of this section is the following:

Theorem 3 *The average coloration number of the permutations in S_n to be routed on a linear network on n vertices is*

$$\frac{n}{4} + \frac{\lambda}{2} n^{1/3} + O(n^{1/6})$$

where $\lambda = 0.99615\ldots$.

To prove this result, we use enumerative and asymptotic combinatorial techniques. Our approach first uses the same methodology applied to permutations as Lagarias and al [?] for involutions with no fixed point. At first we recall in Subsection 5.1 a bijection between permutations in $\mathcal{S}_n$ and special walks in $\mathbb{N} \times \mathbb{N}$, called "Motzkin walks", which are labeled in a certain way [2]. The bijection is such that the height of the walk is equal to the load of the permutation. We get in Subsection 5.3 the generating function of permutations with coloration number k, for any given k. This gives rise to an algorithm to compute exactly the average coloration number of the permutations for any fixed n. Then we are able to combine these enumerative results with random walks techniques developped by Louchard [20] and Daniels and Skyrme [6] to prove Theorem 3. Note that this "random walk" approach was not developed in [?] and we therefore extend our results for permutations to involutions with no fixed point in Section 7.

5.1 A bijection between permutations and Motzkin walks

A **Motzkin walk** w of length n is a n-uple $(w_1, w_2, \ldots, w_n)$ of unitary steps (North-East, South-East or East). Let $\langle_i$ be the height of each step that is the difference between the number of North-East and South-East steps. Then the walk must satisfy the following conditions:

-
- $h_i \geq 0$, $1 \leq i \leq n$;
- $h_n = 0$;

The **height** of a Motzkin walk w is $H(w) = \max\limits_{i \in \{0,1,\ldots,n\}} h_i\}$.

Given two infinite sequences $\{\lambda_n\}_{n \geq 1}$ and $\{b_n\}_{n \geq 0}$, a **labeled Motzkin walk** of length n has the shape of a Motzkin walk and the South-East steps going from (i, y_i) to $(i+1, y_i - 1)$ can be labeled from 1 to λ_{y_i} and the East steps going from (i, y_i) to $(i+1, y_i)$ can be labeled from 1 to b_{y_i}. Moreover, given two sequences $\{\lambda_n\}_{n \geq 1}$ and $\{b_n\}_{n \geq 0}$, let $\mathcal{M}_n$ be the number of labeled Motzkin walks and $\mathcal{M}(z) = \sum_{n \geq 0} \mathcal{M}_n z^n$ the associated generating function.

Proposition 4 *[9, 10, 26] The generating function $\mathcal{M}(z)$ is a continued fraction. Its expression is*

$$\mathcal{M}(z) = \cfrac{1}{1 - b_0 z - \cfrac{\lambda_1 z^2}{1 - b_1 z - \cfrac{\lambda_2 z^2}{1 - b_2 z - \cfrac{\lambda_3 z^2}{1 - b_3 z - \cfrac{\lambda_4 z^2}{\vdots}}}}}.$$

Labeled Motzkin walks are in relation with several well-studied combinatorial objects [9, 26, 27] and in particular with permutations. The walks we will deal with are labeled as follows:

- each South-East step $(i, y_i) \to (i+1, y_i - 1)$ is labeled by an integer between 1 and ${y_i}^2$ (or, equivalently, by a pair of integers, each one between 1 and y_i);
- each East step $(i, y_i) \to (i+1, y_i)$ is labeled by an integer between 1 and $2y_i + 1$.

In Figure 4 we present an example of the labeled Motzkin walks we consider.

Let P_n be the set of such labeled Motzkin walks of length n. We recall that $\mathcal{S}_n$ is the set of permutations on $[n]$. The following result was first established by Françon and Viennot [12]:

Theorem (Françon-Viennot) *There is a one-to-one correspondence between the elements of P_n and the elements of $\mathcal{S}_n$.*

Several bijective proofs of this theorem are known. Biane's bijection [2] is particular, in the sense that it preserves the height: to any labeled Motzkin walk of length n and height k corresponds a permutation in $\mathcal{S}_n$ with load k (and so with coloration number k). We present in the following another version of the Biane's bijection in order to understand the relationship between the height of the labeled Motzkin walks and the load of the permutations.

The bijection We will explain the bijection ϕ from the labeled Motzkin walks of length n to the permutations in $\mathcal{S}_n$ on the linear network on n vertices. The reverse is easy and left to the reader. Consider a linear network P_n on n vertices such that these ones are labeled from left to right from 1 to n. Thus, Biane's correspondence between a labeled Motzkin walk $w = (w_1, w_2, \ldots, w_n)$ and a permutation $\phi(w) = \sigma = (\sigma(1), \sigma(2), \ldots, \sigma(n))$ on P_n is such that, for $1 \leq i \leq n$:

- w_i is an East step of height j and is labeled $2j+1$ if and only if $\sigma(i) = i$.

Figure 4: Example of a labeled Motzkin walk of length 11 and height 3.

Figure 5: From the shape of the path to the shape of the permutation

- w_i is an East step of height j and is labeled l with $1 \leq l \leq j$ if and only if $\sigma^{-1}(i) < i$ and $\sigma(i) > i$.
- w_i is an East step of height j and is labeled l with $j+1 \leq l \leq 2j$ if and only if $\sigma^{-1}(i) > i$ and $\sigma(i) < i$.
- w_i is a North-East step if and only if $\sigma(i) > i$ and $\sigma^{-1}(i) > i$.
- w_i is a South-East step if and only if $\sigma(i) < i$ and $\sigma^{-1}(i) < i$.

The recurrence (1) and the previous correspondence automatically give us that the height each step w_i is equal to $L_{P_n}(\sigma, i^+)$ for $1 \leq i \leq n-1$; as the load of the arc i^+ in P_n is equal to the number of integers $j \leq i$ such that $\sigma(j) > j$ and $\sigma^{-1}(j) > j$, minus the number of integers $j \leq i$ such that $\sigma(j) < j$ and $\sigma^{-1}(j) < j$.

Given a labeled Motzkin walk, it is easy to draw the shape of the permutation σ (beginning and end of the path $i \rightsquigarrow \sigma(i)$, $1 \leq i \leq n$), using the previous correspondence. The beginning of the path $i \rightsquigarrow \sigma(i)$ uses arc i^+ in P_n if and only if w_i is a North East step or an East step at height j with a label between $j+1$ and $2j$. The beginning of the path $i \rightsquigarrow \sigma(i)$ uses arc $(i-1)^-$ in P_n if and only if w_i is a South East step or an East step at height j with a label between 1 and j. The end of the path $\sigma^{-1}(i) \rightsquigarrow i$ uses arc i^- in P_n if and only if w_i is a North East step or an East step at height j with a label between 1 and j. The end of the path $\sigma^{-1}(i) \rightsquigarrow i$ uses arc $(i-1)^+$ in P_n if and only if w_i is a South East step

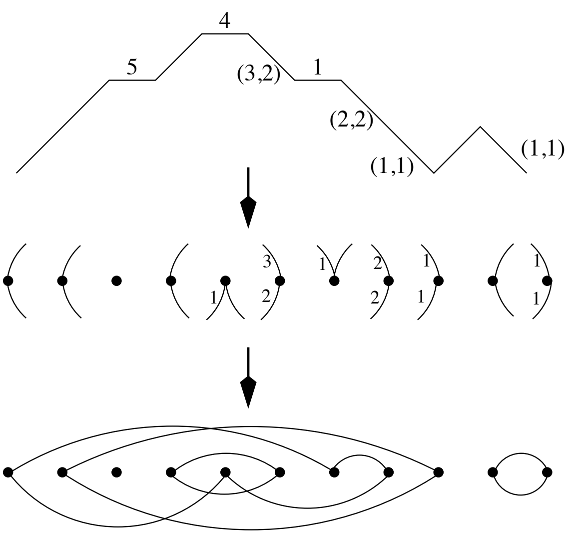

Figure 6: From a labeled Motzkin walk to a permutation

or an East step at height j with a label between $j+1$ and $2j$. An example is illustrated on Figure 5. Now we label the shape of the permutation to keep all the information of the labeled Motzkin path. For i from 2 to n, if w_i is a South-East step with label (x, y) then we label the end of the path $\sigma^{-1}(i) \rightsquigarrow i$ by x and the beginning of the path $i \rightsquigarrow \sigma(i)$ by y. For i from 2 to n, if w_i is an East step of height j labeled by l; if $j+1 \leq l \leq 2j$ then we label the beginning of the path $i \rightsquigarrow \sigma(i)$ by $l-j$; if $1 \leq l \leq j$ then we label the end of the path $\sigma^{-1}(i) \rightsquigarrow i$ by l. See Figure 6 for an example of labeling. Finally we associate free beginnings and ends of paths going from left to right. For any free end of a path (resp. beginning of the path) labeled x, we associate to it the x^{th} free unlabeled beginning of a path (resp. unlabeled end of a path) starting from the left. See Figure 6 to see an example of the construction of the permutation. □

5.2 Proof of Theorem 3

In [20], Louchard analyzes some list structures; in particular his "dictionary structure" corresponds to our labeled Motzkin walks. We will use his notation in order to refer directly to his article. From Louchard's theorem 6.2, we deduce the following lemma:

Lemma 6 *The height $Y^*([nv])$ of a random labeled Motzkin walk of length n after the step $[nv]$ ($v \in [0,1]$)) has the following behavior*

$$\frac{Y^*([nv]) - nv(1-v)}{\sqrt{n}} \Rightarrow X(v),$$

where "$\Rightarrow$" denotes the weak convergence and X is a Markovian process with mean 0 and covariance $C(s,t) = 2s^2(1-t)^2$, $s \le t$. $X(\cdot)$ can also be written as $X(v) = \sqrt{2}(1-v)^2 B\left(\frac{v^2}{(1-v)^2}\right)$, for some Brownian motion $B(\cdot)$.

Then the work of Daniels and Skyrme [6] gives us a way to compute the maximum of $Y^*([nv])$, that is the height of a random labeled Motzkin walk. Let $X(v)$ be a gaussian process with covariance and superposed on a curve $\tilde{y}(v)$. Assume that $\tilde{y}(v)$ is given by $\sqrt{n}y(v)$, $n \gg 1$ and it has a unique maximum at $\bar{v}$. Their result is the following :

Theorem 4 *[6] The random variable $m = max_v(X(v) + \tilde{y}(t))$ is asymptotically gaussian with mean and variance*

$$E(m) \sim \lambda n^{-1/6} A^{2/3} B^{-1/3}, \qquad \sigma^2(m) \sim c$$

where

$$c = C(\bar{v}, \bar{v}), \quad B = -y''(v), \quad \lambda = 0.99615\ldots, \quad A = \frac{\partial C}{\partial s}(\bar{v}, \bar{v}) - \frac{\partial C}{\partial t}(\bar{v}, \bar{v}).$$

From the Louchard's result we know that

$$Y^*([nv]) = \sqrt{n}\left(\sqrt{n}v(1-v) + X(v)\right) + O(1).$$

Therefore we have $y(v) = v(1-v)$ and the unique maximum is attained at $\bar{v} = 1/2$. The covariance of our gaussian process is $C(s,t) = 2s^2(1-t)^2$. We just then have to apply the theorem [?] and get :

$$c = 1/8, \quad B = 2, \quad A = 1/2.$$

We can write now our result :

Proposition 5 *The height of a random labeled Motzkin walk Y^* is*

$$\max_v Y^*([nv]) = \frac{n}{4} + m\sqrt{n} + O(n^{1/6}), \tag{2}$$

where m is asymptotically Gaussian with mean $E(m) \sim \lambda n^{-1/6}/2$ and variance $\sigma^2(m) \sim 1/8$ and $\lambda = 0.99615\ldots$.

In the formula (2) of the above Proposition 5, the only non-deterministic part is m which is Gaussian. So we just have to replace m by $E(m)$ the mean of the coloration number and hence to prove Theorem 3. We can also get directly the variance :

Theorem 5 *The variance of the coloration number of the permutations in S_n to be routed on a linear network on n vertices is $\frac{n}{8}$.*

Proof. The variance of $\max_v Y^*([nv])$ is just $n\sigma^2(m)$. □

5.3 An algorithm to compute exactly the average coloration number

From previous Biane's bijection and Proposition 4, we can get directly the generating function of the permutations in $\mathcal{S}_n$ of coloration number at most k to be routed on a linear network on n vertices,

$$H_{\leq k}(z) = \sum_{n\geq 0} \sum_{\sigma \in \mathcal{S}_{n,\leq k}} z^n$$

that is

$$\cfrac{1}{1 - z - \cfrac{z^2}{1 - 3z - \cfrac{4z^2}{1 - 5z - \cfrac{\vdots}{1 - (2k-1)z - \cfrac{k^2 z^2}{1 - (2k+1)z}}}}} .$$

Note that for any fixed k this generating function is rational. We can also use known results in enumerative combinatorics [9, 27] to get the generating function of the permutations of coloration number **exactly** k,

$$H_k(z) = \sum_{n\geq 0} \sum_{\sigma \in \mathcal{S}_{n,k}} z^n$$

that is

$$\frac{(k!)^2 z^{2k}}{P^*_{k+1}(z) P^*_k(z)}$$

with $P_0(z) = 1$, $P_1(z) = z - b_0$ and $P_{n+1}(z) = (t - b_n)P_n(z) - \lambda_n P_{n-1}(z)$ for $n \geq 1$, where P^* is the reciprocal polynomial of P, that is $P^*_n(z) = z^n P_n(1/z)$ for $n \geq 0$.

This generating function leads to a recursive algorithm to compute the number of permutations with coloration number k, denoted by $h_{n,k}$.

Proposition 6 *The number of permutations in $\mathcal{S}_n$ to be routed on a linear network on n vertices with coloration number k, follows the following recurrence*

$$h_{n,k} = \begin{cases} 0 & \text{if } n < 2k \\ (k!)^2 & \text{if } n = 2k \\ -\sum_{i=1}^{2h+1} p(i) h_{n-i,k} & \text{otherwise} \end{cases}$$

where $p(i)$ is the coefficient of z^i in $P^*_{k+1}(z)P^*_k(z)$.

From this result we are able to compute thanks to Maple the average coloration number of permutations in $\mathcal{S}_n$ to be routed on a linear network on n vertices as it is $\bar{h}(n) = \sum_{k\geq 0} k h_{n,k}/n!$. The first forty values are presented in Table 1.

n	$\bar{h}(n)$	n	$\bar{h}(n)$	n	$\bar{h}(n)$	n	$\bar{h}(n)$	n	$\bar{h}(n)$
1	0	9	2.60	17	4.82	25	7.00	33	9.13
2	0.5	10	2.88	18	5.10	26	7.27	34	9.4
3	0.83	11	3.16	19	5.37	27	7.53	35	9.66
4	1.12	12	3.44	20	5.65	28	7.80	36	9.93
5	1.42	13	3.72	21	5.92	29	8.07	37	10.19
6	1.73	14	4.00	22	6.19	30	8.33	38	10.46
7	2.02	15	4.27	23	6.46	31	8.60	39	10.72
8	2.31	16	4.55	24	6.73	32	8.86	40	10.98

Table 1: Average coloration number of permutations in $\mathcal{S}_n$.

6 Average coloration number on arbitrary tree networks

In this section, we extend the average complexity results on linear networks obtained in Section 4, to the case of arbitrary tree networks. Given a tree T on n vertices, by Theorem 2, we know that it is NP-hard to compute $R_T(\sigma)$ for a permutation σ even if T is a binary tree and σ is an involution or a circular permutation. By Proposition 1, we know that computing $R_T(\sigma)$ is $5/3-$approximable. The aim of this section is then to find the average coloration number required for this approximation algorithm.

By Lemma 2, we know that $\bar{L}_T \leq \bar{R}_T \leq \frac{5}{3}\bar{L}_T + 1$. Therefore, we will compute the average load $\bar{L}_T$ for any tree T and will obtain bounds on $\bar{R}_T$, the average number of colors needed to color any permutation-path set on T. In Section 5.1 we present an upper bound for the average coloration number on tree networks. In Section 5.2 we obtain exact results on the average number of colors needed to color any permutation-path set on generalized star tree networks.

6.1 Upper bound

Let's remark that for any tree T on n vertices and for each vertex i of T, there exists a relabeling of the vertex-set of T such that, for any permutation $\sigma \in \mathcal{S}_n$, $L_T(\sigma, i^+) = L_{P_n}(\sigma, |T(i)|)$. Such a relabeling is trivial. The vertices of $T(i)$ are relabeled with integers in $\{1, 2, \ldots, |T(i)|\}$, and the vertices in $T \backslash T(i)$ are relabeled with integers in $\{|T(i)|+1, \ldots, n\}$. Therefore, the Lemma **??** can be rewritten as follows.

Lemma 7 *Let T be a tree on n vertices and let σ^* be a random permutation in $\mathcal{S}_n$. The load of any arc i^+ of T induced by σ^*, denoted by $L_T^*(\sigma^*, i^+)$, has the following behavior*

$$L_T^*(\sigma^*, i^+) = nv_T(i)(1 - v_T(i)) + X(v_T(i))\sqrt{n} + O(1),$$

where X is a Markovian process with mean 0 and covariance $C(s,t) = s^2(1-t)^2$, $s \leq t$, and where $v_T(i) = |T(i)|/n$.

As defined before, for any vertex i of T, $\tilde{v}_T(i) = \min(v_T(i), 1 - v_T(i))$ and $\tilde{v}_T = \max_i \tilde{v}_T(i)$. Thus, by Theorem 4 and Theorem 3, we obtain the following theorem if T is a bounded degree tree.

Theorem 6 *The average load induced by all permutations $\sigma \in \mathcal{S}_n$ on T is*

$$\bar{L}_T = n\tilde{v}_T(1 - \tilde{v}_T) + o(n).$$

Proof. By Lemma 4, we have that $\bar{L}_T \geq n\tilde{v}_T(1 - \tilde{v}_T)$. By Lemma 7, we know that for all ϵ, there exists $n_0(\epsilon)$ such that, for all $n \geq n_0$ and any tree T on n vertices, $\bar{L}_T \leq n\tilde{v}_T(1 - \tilde{v}_T) + n^{1/2+\epsilon}$, which proves the theorem. □

From Lemma 2 and Theorem 6, we obtain the following upper bound on the average number of colors needed to color any permutation $\sigma \in \mathcal{S}_n$ on any bounded degree tree T on n vertices.

Theorem 7 *For all ϵ, there exists $n_0(\epsilon)$ such that, for all $n \geq n_0$ and any bounded degree tree T on n vertices, the average number of colors $\bar{R}_T$ needed to color any permutation $\sigma \in \mathcal{S}_n$ on T verifies, $\bar{R}_T \leq \left(\frac{5}{3} + \epsilon\right) n\tilde{v}_T(1 - \tilde{v}_T)$.*

6.2 The average number of colors in generalized star trees

Let k be a fixed integer, λ be a partition of $n-1$ in k parts and $\mathrm{GST}(\lambda)$ be the associated generalized star tree on n vertices. In this case, we have $\tilde{v}_{\mathrm{GST}(\lambda)} = \min(\lfloor n/2 \rfloor, \lambda_1)$. Moreover, in [14] has been shown that for any collection of paths $\mathcal{P}$ on a generalized star $\mathrm{GST}(\lambda)$, $R_{\mathrm{GST}(\lambda)}(\mathcal{P}) = L_{\mathrm{GST}(\lambda)}(\mathcal{P})$. Therefore, we obtain the following results.

Theorem 8 *The average number of colors needed to color any permutation $\sigma \in \mathcal{S}_n$ on a generalized star tree $GST(\lambda)$ having n vertices is:*

$$\bar{R}_{GST(\lambda)} = nm(1-m) + o(n),$$

where $m = \min(\lfloor n/2 \rfloor, \lambda_1)/n$.

In particular we can obtain the following result.

Theorem 9 *The average number of colors needed to color any permutation $\sigma \in \mathcal{S}_{nk+1}$ on a generalized star tree $GST(\lambda)$ having $nk+1$ vertices and k branches of length n is $n(n(k-1)+1)/nk + 1 + o(n)$.*

7 Average coloration number for involutions

Given a tree T on $2n$ vertices, by Theorem 2, we know that it is NP-hard to compute $R_T(\sigma)$ for an arbitrary involution σ in I_{2n} even if T is a binary tree. By Proposition 2, we know that computing $R_T(\sigma)$ is $3/2-$approximable. The aim of this section is then to find the average coloration number required for this approximation algorithm and therefore complete the work initiated in [?]. We will compute the average load $\tilde{L}_T$ for any tree T and will obtain bounds on $\tilde{R}_T$, the average number of colors needed to color any involution-path set on T.

We can compute easily the average load of any arc i^+ of T, $1 \leq i \leq 2n-1$: $\tilde{L}_T(i) = |T(i)|(2n - |T(i)|)/2n$. Therefore, we obtain a lower bound for the average load $\tilde{L}_T$ and the following lower bound on the average number of colors needed to color any involution $\sigma \in I_{2n}$ on any tree T on $2n$ vertices.

Lemma 8 $\tilde{R}_T \geq 2n\tilde{v}_T(1 - \tilde{v}_T)$.

By using a classical bijection between the involutions in I_{2n} and the set V_{2n} of special walks in $N \times N$ called *labeled Dyck walks* of length $2n$ [?, ?] which preserves the load as the Biane's bijection, and from Louchard's theorem 5.3 [20], we get the following results that can be obtained applying the same methods as in the previous sections for arbitrary permutations.

Lemma 9 [20] *Let P_{2n} be the path graph on $2n$ vertices and let σ^* be a random involution in I_{2n}. The load $L^*_{P_{2n}}(\sigma^*, \lfloor nv \rfloor)$ of arc $\lfloor nv \rfloor$ ($v \in [0, 2]$) of P_{2n} has the following behavior*

$$L^*_{P_{2n}}(\sigma^*, \lfloor nv \rfloor) = nv(2 - v)/2 + X(v)\sqrt{n} + O(1),$$

where X is a Markovian process with mean 0 and covariance $C(s, t) = s^2(2 - t)^2/4$, $s \leq t$.

Lemma 10 *Let T be a tree on $2n$ vertices and let σ^* be a random involution in I_n. The load of any arc i^+ of T induced by σ^*, denoted by $L^*_T(\sigma^*, i^+)$, has the following behavior*

$$L^*_T(\sigma^*, i^+) = nv_T(i)(2 - v_T(i))/2 + X(v_T(i))\sqrt{n} + O(1),$$

where X is a Markovian process with mean 0 and covariance $C(s, t) = s^2(2 - t)^2/4$, $s \leq t$, and where $v_T(i) = |T(i)|/2n$.

Theorem 10 *Let T be a tree on $2n$ vertices. The average load induced by all involutions $\sigma \in I_{2n}$ on T is $\tilde{L}_T = 2n\tilde{v}_T(1 - \tilde{v}_T) + o(n)$.*

Theorem 11 *For all ϵ, there exists $n_0(\epsilon)$ such that, for all $n \geq n_0$ and any tree T on $2n$ vertices, the average number of colors $\tilde{R}_T$ verifies, $\tilde{R}_T \leq \left(\frac{3}{2} + \epsilon\right) 2n\tilde{v}_T(1 - \tilde{v}_T)$.*

We just have to apply Theorem 10 to get the average coloration number for involutions on linear networks and generalized star tree networks obtaining exactly the same asymptotic behavior as for arbitrary permutations.

Corollary 1 *The average coloration number for involutions on* P_{2n} *induced by all involutions* $\sigma \in I_{2n}$ *on* T *is* $\tilde{R}_{P_{2n}} = n/2 + o(n)$.

Note that the average complexity for involutions is the same as for permutations. Let k be a fixed integer greater than 2 and $\lambda = (\lambda_1, \ldots, \lambda_k$ a partition of $n-1$ with k parts.

Corollary 2 *The average coloration number for involutions on* $GST(\lambda)$ *induced by all involutions* $\sigma \in I_{2n}$ *on* T *is* $\tilde{R}_\lambda = 2n\tilde{v}_\lambda(1-\tilde{v}_\lambda) + o(n)$, *with* $\tilde{v}_\lambda = \min(n/2, \lambda_1)$.

References

[1] Y. Aumann, Y. Rabani. *Improved bounds for all optical routing.* In Proc. of the 6th ACM-SIAM Symposium on Discrete Algorithms (SODA), pp 567-576, 1995.

[2] Ph. Biane. *Permutations suivant le type d'excédance et le nombre d'inversions et interprétation combinatoire d'une fraction continue de Heine.* Eur. J. Comb., 14(4):277-284, 1993.

[3] B. Bollobàs. *GRAPH THEORY. An Introductory Course.* Graduate Texts in Mathematics no. 63, Springer-Verlag, 1979.

[4] I. Caragiannis, Ch. Kaklamanis, P. Persiano. *Wavelength Routing of Symmetric Communication Requests in Directed Fiber Trees.* In Proc. of SIROCCO'98, 1998.

[5] N. K. Cheung, K. Nosu, and G. Winzer, editors. *Special Issue on Dense Wavelength Division Multiplexing Techniques for High Capacity and Multiple Access Communications Systems.* IEEE Journal on Selected Areas in Communications, 8(6), 1990.

[6] H. E. Daniels, T. H. R. Skyrme. *The maximum of a random walk whose mean path has a maximum.* Adv. Appl. Probab., 17:85-99, 1985.

[7] T. Erlebach, K. Jansen. *Call scheduling in trees, rings and meshes.* In Proc. of the 30th Hawaii International Conference on System Sciences HICSS-30, vol. 1, pp 221-222. IEEE CS Press, 1997.

[8] T. Erlebach, K. Jansen, Ch. Kaklamanis, M. Mihail, P. Persiano. *Optimal wavelength routing on directed fiber trees.* Theoret. Comput. Sci. 221(1-2):119-137, 1999.

[9] Ph. Flajolet. *Combinatorial aspects of continued fractions.* Discrete Math., 32:125-161, 1980.

[10] Ph. Flajolet. *Combinatorial aspects of continued fractions.* Ann. Discrete Math., 9:217-222, 1980.

[11] H. Fleischner. *Eulerian graphs and related topics.* Annals of Discrete Mathematics, no. 45, vol. 2, Elsevier Science Publisher B.V., 1991.

[12] J. Françon, X. Viennot. *Permutations selon leurs pics, creux, doubles montées et doubles descentes, nombres d'Euler et nombres de Genocchi.* Discrete Math., 28:21-35, 1979.

[13] M. R. Garey, D. S. Johnson, G. L. Miller, Ch. Papadimitriou. *The complexity of coloring circular arcs and chords.* SIAM J. Alg. Disc. Meth., 1(2):216-227,1980.

[14] L. Gargano, P. Hell, S. Perennes. *Coloring all directed paths in a symmetric tree with applications to WDM routing.* In Proc. of 24th International Colloquium on Automata, Languages, and Programming (ICALP 97), LNCS 1256, pp 505-515, 1997.

[15] M. C. Golumbic. *Algorithmic graph theory and perfect graphs.* Academic Press, New York, 1980.

[16] Q.-P. Gu, H. Tamaki. *Routing a permutation in the Hypercube by two sets of edge-disjoint paths.* In Proc. of 10th IPPS. IEEE Computer Society Press, 1996.

[17] U. I. Gupta, D. T. Lee, J. Y.-T. Leung. *Efficient algorithms for interval graphs and circular-arc graphs.* Networks, 12:459-467, 1982.

[18] S. R. Kumar, R. Panigrahy, A. Russel, R. Sundaram. *A note on optical routing on trees.* Inf. Process. Lett., 62:295-300, 1997.

[19] F. T. Leighton, S. Rao. *An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms.* In Proc. of 29th IEEE FOCS, pp 422-431, 1988.

[20] G. Louchard. *Random walks, Gaussian processes and list structures.* Theoret. Comput. Sci., 53:99-124, 1987.

[21] D. W. Matula, F. Shahrokhi. *The Maximum Concurrent Flow Problem and Sparsest Cuts.* Technical Report, Southern Methodist University, March 1986.

[22] J.C. Lagarias, J.C., A.M. Odlyzko and D.B. Zagier. *On the capacity of disjointly shared networks.* Comput. Networks ISDN Syst. 10:275-285,1985.

[23] M. Paterson, H. Schröder, O. Sýkora, I. Vrťo. *On Permutation Communications in All-Optical Rings.* In Proc. of SIROCCO'98, 1998.

[24] P. Raghavan, U. Upfal. *Efficient routing in all optical networks.* In Proc. of the 26th ACM Symposium on Theory of Computating, pp 133-143, 1994.

[25] A. Tucker. *Coloring a family of circular arcs.* SIAM J. Appl. Maths., 29(3):493-502, 1975.

[26] X. Viennot. *A combinatorial theory for general orthogonal polynomials with extensions and applications.* Lect. Notes Math., 1171:139-157, 1985. Polynômes orthogonaux et applications, Proc. Laguerre Symp., Bar-le-Duc/France 1984.

[27] X. Viennot. *Une théorie combinatoire des polynômes orthogonaux généraux.* Notes de conférences, Univ. Quebec, Montreal, 1985.

[28] G. Wilfong, P. Winkler. *Ring routing and wavelength translation.* In Proc. of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'98, pp 333-341, 1998.