



Approximating Shortest Connected Graph Transformation for Trees

Nicolas Bousquet, Alice Joffard

► To cite this version:

Nicolas Bousquet, Alice Joffard. Approximating Shortest Connected Graph Transformation for Trees. 2019. hal-02358489

HAL Id: hal-02358489

<https://hal.science/hal-02358489>

Preprint submitted on 18 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Approximating SHORTEST CONNECTED GRAPH TRANSFORMATION for Trees^{*}

Nicolas Bousquet¹ and Alice Joffard²

¹ Univ. Grenoble Alpes, CNRS, Laboratoire G-SCOP, Grenoble-INP, Grenoble, France.

`nicolas.bousquet@grenoble-inp.fr`

² LIRIS, Université Claude Bernard, Lyon, France

`alice.joffard@liris.cnrs.fr`

Abstract. Let G, H be two connected graphs with the same degree sequence. The aim of this paper is to find a transformation from G to H via a sequence of flips maintaining connectivity. A flip of G is an operation consisting in replacing two existing edges uv, xy of G by ux and vy .

Taylor showed that there always exists a sequence of flips that transforms G into H maintaining connectivity. Bousquet and Mary proved that there exists a 4-approximation algorithm of a shortest transformation. In this paper, we show that there exists a 2.5-approximation algorithm running in polynomial time. We also discuss the tightness of the lower bound and show that, in order to drastically improve the approximation ratio, we need to improve the best known lower bounds.

1 Introduction

Sorting by reversals problem. The problem of sorting by reversals has been widely studied in the last twenty years in genomics. The reversal of a sequence of DNA is a common mutation of a genome, that can lead to major evolutionary events. It consists, given a DNA sequence that can be represented as a labelled path x_1, \dots, x_n on n vertices, in turning around a part of it. More formally, a reversal is a transformation that, given two integers $1 \leq i < j \leq n$, transforms the path x_1, \dots, x_n into $x_1, \dots, x_{i-1}, x_j, x_{j-1}, \dots, x_i, x_{j+1}, \dots, x_n$. It is easy to prove that, given two paths on the same vertex set (and with the same leaves), there exists a sequence of reversals that transforms the first into the second. Biologists want to find the minimum number of reversals needed to transform a genome (i.e. a path) into another in order to compute the evolutionary distance between different species.

An input of the SORTING BY REVERSALS problem consists of two paths P, P' with the same vertex set (and the same leaves) and an integer k . The output is positive if and only if there exists a sequence of at most k reversals that transforms P into P' . Caprara proved that the SORTING BY REVERSALS problem is NP-complete [4]. Kececiloglu and Sankoff first proposed an algorithm that computes a sequence of reversals of size at most twice the length of an optimal solution in polynomial time [10]. Then, Christie improved it into a 3/2-approximation algorithm [5]. The best polynomial time algorithm known so far is a 1.375-approximation due to Berman et al. [2].

A reversal can be equivalently defined as follows: given a path P and two edges ab and cd , a reversal consists in the deletion of the edges ab and cd and the addition of ac and bd that keeps the connectivity of the graph. Indeed, when we transform x_1, \dots, x_n into $x_1, \dots, x_{i-1}, x_j, x_{j-1}, \dots, x_i, x_{j+1}, \dots, x_n$, we have deleted the edges $x_{i-1}x_i$ and x_jx_{j+1} and have created the edges $x_{i-1}x_j$ and x_ix_{j+1} . In this paper, we study the generalization of the SORTING BY REVERSALS problem for trees and general graphs that has also been extensively studied in the last decades.

SHORTEST CONNECTED GRAPH TRANSFORMATION problem. Let $G = (V, E)$ be a graph where V denotes the set of vertices and E the set of edges. For basic definitions on graphs, the reader is referred to [6]. All along the paper, the graphs are loop-free but may admit multiple edges. A *tree* is a connected graph which does not contain any cycle (a multi-edge being considered as a cycle).

^{*} This work was supported by ANR project GrR (ANR-18-CE40-0032).

The *degree sequence* of a graph G is the sequence of the degrees of its vertices in non-increasing order. Given a non-increasing sequence of integers $S = \{d_1, \dots, d_n\}$, a graph $G = (V, E)$ whose vertices are labeled as $V = \{v_1, \dots, v_n\}$ realizes S if $d(v_i) = d_i$ for all $i \leq n$. Senior [12] gave necessary and sufficient conditions to guarantee that, given a sequence of integers $S = \{d_1, \dots, d_n\}$, there exists a connected multigraph realizing S . Hakimi [7] then proposed a polynomial time algorithm that outputs a connected (multi)graph realizing S if such a graph exists or returns no otherwise.

A *flip* σ (also called *swap* or *switch* in the literature) on two edges ab and cd consists in deleting the edges ab and cd and creating the edges ac and bd (or ad and bc)³. The flip operation that transforms the edges ab and cd into the edges ac and bd is denoted $(ab, cd) \rightarrow (ac, bd)$. When the target edges are not important we will simply say that we flip the edges ab and cd .

Let $S = \{d_1, \dots, d_n\}$ be a non-increasing sequence and let G and H be two graphs on n vertices v_1, \dots, v_n realizing S . The graph G can be *transformed* into H if there is a sequence $(\sigma_1, \dots, \sigma_k)$ of flips that transforms G into H . Note that since flips do not modify the degree sequence, all the intermediate graphs also realize S . Let $\mathcal{G}(S)$ be the graph whose vertices are the loop-free multigraphs realizing S and where two vertices G and H of $\mathcal{G}(S)$ are adjacent if G can be transformed into H via a single flip. Since the flip operation is reversible, the graph $\mathcal{G}(S)$ is an undirected graph called the *reconfiguration graph* of S . Note that there exists a sequence of flips between any pair of graphs realizing S if and only if the graph $\mathcal{G}(S)$ is connected. Hakimi [8] proved that, for any non-increasing sequence S , if the graph $\mathcal{G}(S)$ is not empty then it is connected.

One can wonder if the reconfiguration graph is still connected when we restrict to graphs with stronger properties. For a graph property Π , let us denote by $\mathcal{G}(S, \Pi)$ the subgraph of $\mathcal{G}(S)$ induced by the graphs realizing S that have the property Π . If we respectively denote by \mathcal{C} and \mathcal{S} the property of being connected and simple, Taylor proved in [13] that $\mathcal{G}(S, \mathcal{C})$, $\mathcal{G}(S, \mathcal{S})$ and $\mathcal{G}(S, \mathcal{C} \wedge \mathcal{S})$ are connected (where \wedge stands for “and”). Let G, H be two graphs of $\mathcal{G}(S, \Pi)$. A sequence of flips *transforms* G into H in $\mathcal{G}(S, \Pi)$ if the sequence of flips transforms G into H and all the intermediate graphs also have the property Π . In other words, a sequence of flips that transforms G into H in $\mathcal{G}(S, \Pi)$ is a path between G and H in $\mathcal{G}(S, \Pi)$. Since [13] ensures that $\mathcal{G}(S, \Pi)$ is connected, one can ask what is the minimum length of such a transformation between G and H . This problem is known to be NP-hard, see e.g. [4]. In this paper we will study the following problem:

SHORTEST CONNECTED GRAPH TRANSFORMATION

Input: Two connected multigraphs G, H with the same degree sequence.

Output: The minimum number of flips needed to transform G into H in $\mathcal{G}(S, \mathcal{C})$.

Note that SHORTEST CONNECTED GRAPH TRANSFORMATION is a generalization of SORTING BY REVERSALS since, when the degree sequence consists of $n - 2$ vertices of degree 2 and two vertices of degree 1, we simply want to find a sequence of reversals of minimum length between two paths. Bousquet and Mary [3] proposed a 4-approximation algorithm for SHORTEST CONNECTED GRAPH TRANSFORMATION. Our main result is the following:

Theorem 1. SHORTEST CONNECTED GRAPH TRANSFORMATION admits a 2.5-approximation algorithm.

Section 3 is devoted to the proof of Theorem 1. In order to prove it, we will mainly focus on the SHORTEST TREE TRANSFORMATION problem which is the same as SHORTEST CONNECTED GRAPH TRANSFORMATION except that the input consists of trees with the same degree sequence. Informally speaking, it is due to the fact that if an edge of the symmetric difference appears in some cycle, then we can reduce the size of the symmetric difference in one flip, as observed in [3].

When we desire to give some explicit bound on the quality of a solution, we need to compare it with the length of an optimal transformation. When we do not want to keep connectivity, Will [14] gives an explicit formula of the number of steps in a minimum transformation. When we want to keep connectivity, no such formula is known. Our 5/2-approximation algorithm is obtained by comparing it to the formula

³ In the case of multigraphs, we simply decrease by one the multiplicities of ab and cd and increase by one the ones of ac and bd .

of Will (which is a lower bound when we want to keep connectivity). In Section 4, we discuss the tightness of this lower bound. We exhibit two graphs G and H such that the length of a shortest transformation between G and H is at least 1.5 times larger than the bound given by [14], and even twice longer under some assumptions on the set of possible flips. In order to prove this result, we generalize some notions introduced for sorting by reversals in [5] to general graphs.

This example ensures that if we want to find an approximation algorithm with a ratio better than 1.5, we might have to improve the algorithm, but overall, we need to improve the lower bound. The formal point and the two graphs G and H can be found in Section 4.

Related works

Mass spectrometry. Mass spectrometry is a technique used by chemists in order to obtain the formula of a molecule. It provides the mass-to-charge (m/z) ratio spectrum of the molecule from which we can deduce how many atoms of each element the molecule has. With this formula, we would like to find out the nature of the molecule, i.e. the bonds between the different atoms. But the existence of structural isomers points out that there could exist several solutions for this problem. Thus, we would like to find all of them. Since the valence of each atom is known, this problem actually consists in finding all the connected loop-free multi-graphs whose degree sequence is the sequence of the valences of those atoms. The reconfiguration problem we are studying here can be a tool for an enumeration algorithm consisting in visiting the reconfiguration graph.

Flips and reconfiguration. The SHORTEST CONNECTED GRAPH TRANSFORMATION problem belongs to the class of reconfiguration problems that received a considerable attention in the last few years. Reconfiguration problems consist, given two solutions of the same problem, in transforming the first solution into the second via a sequence of "elementary" transformations (such as flips) maintaining some properties all along. For more information on reconfiguration problems, the reader is referred for instance to [11].

2 Preliminaries

2.1 Symmetric difference

Unless specified otherwise, we consider unoriented loop-free multigraphs. Let $G = (V(G), E(G))$ be a graph where $V(G)$ is the set of vertices of G and $E(G)$ is its set of edges. The *intersection* of two graphs G and H on the same set of vertices V is the graph $G \cap H$ with vertex set V , and such that $e \in E(G \cap H)$, with multiplicity m , if the minimum multiplicity of e in both graphs is m . Their *union*, $G \cup H$, has vertex set V , and $e \in E(G \cup H)$, with multiplicity m , if and only if the maximum multiplicity of e in G and H is m . Finally, the *difference* $G - H$ has vertex set V and $e \in E(G - H)$ with multiplicity m if and only if the difference between its multiplicities in G and H is $m > 0$. The *symmetric difference* of G and H is $\Delta(G, H) = (G - H) \cup (H - G)$. We denote by $\delta(G, H)$ the number of edges of $\Delta(G, H)$.

Let G, H be two graphs with the same degree sequence. An edge e of G is *good* if it is in $G \cap H$ and is *bad* otherwise. Note that since G and H have the same degree sequence, the graph $\Delta(G, H)$ has even degree on each vertex and the number of edges of G incident to v is equal to the number of edges of H incident to v .

Each flip removes at most 4 edges of the symmetric difference. Therefore, the length of a transformation from G to H is at least $\delta(G, H)/4$. In fact, it is possible to obtain a slightly better bound on the length of the transformation. A cycle C in $\Delta(G, H)$ is *alternating* if edges of G and H alternate in C . Since the number of edges of G incident to v is equal to the number of edge of H incident to v in $\Delta(G, H)$, the graph $\Delta(G, H)$ can be partitioned into a collection of alternating cycles. We denote by $mnc(G, H)$ the maximal number of cycles in a partition \mathcal{C} of $\Delta(G, H)$ into alternating cycles. Will [14] proved the following:

Theorem 2 (Will [14]). *Let G, H be two graphs with the same degree sequence. A shortest sequence of flips that transforms G into H (that does not necessarily maintain the connectivity of the intermediate graphs) has length exactly $\frac{\delta(G, H)}{2} - mnc(G, H)$.*

Note that Theorem 2 indeed provides a lower bound for a transformation of SHORTEST CONNECTED GRAPH TRANSFORMATION.

2.2 Basic facts concerning flips

Let $G = (V, E)$ be an unoriented graph and $v \in V(G)$. The set $N_G(v)$ of *neighbours* of v in G is the set of vertices u such that $uv \in E(G)$. Let D be a directed graph and $v \in V(D)$. The set $N_D^-(v)$ of *in-neighbours* of v in D is the set of vertices u such that uv is an arc of D , and the set $N_D^+(v)$ of *out-neighbours* of v in D is the set of vertices u such that vu is an arc of D . When G and D are obvious from the context we will simply write $N(v), N^-(v), N^+(v)$.

The *inverse* σ^{-1} of a flip σ is the flip such that $\sigma \circ \sigma^{-1} = id$, i.e. applying σ and then σ^{-1} leaves the initial graph. The *opposite* $-\sigma$ of a flip σ is the unique other flip that can be applied to the two edges of σ . If we consider a flip $\sigma = (ab, cd) \rightarrow (ac, bd)$, then $\sigma^{-1} = (ac, bd) \rightarrow (ab, cd)$ and $-\sigma = (ab, cd) \rightarrow (ad, bc)$. Note that $-\sigma$ is a flip deleting the same edges as σ while σ^{-1} cancels the flip σ . When we transform a graph G into another graph H , we can flip the edges of G or the edges of H . Indeed, applying the sequence of flips $(\sigma_1, \dots, \sigma_i)$ to transform G into a graph K , and the sequence of flips (τ_1, \dots, τ_j) to transform H into K is equivalent to applying the sequence $(\sigma_1, \dots, \sigma_i, \tau_j^{-1}, \dots, \tau_1^{-1})$ to transform G into H .

Let $G = (V, E)$ be a connected graph and let H be a graph with the degree sequence of G . A flip is *good* if it flips bad edges and creates at least one good edge. It is *bad* otherwise. A *connected flip* is a flip such that its resulting graph is connected. Otherwise, it is *disconnected*. A *path* from $a \in V$ to $b \in V$ is a sequence of vertices (v_1, \dots, v_k) such that $a = v_1, b = v_k$, for every integer $i \in [k-1]$, $v_i v_{i+1} \in E(G)$ and there is no repetition of vertices. Similarly, a path from e to f with $e, f \in E(G)$ is a path from an endpoint of e to an endpoint of f that does not contain the other endpoint of e and of f . A path between x and y (vertices or edges) is a path from x to y or a path from y to x . The *content* of a path is its set of vertices. We say that an edge e *belongs to* (or *is on*) a path P if both endpoints of e appear consecutively in P . The *intersection* $P_1 \cap P_2$ of two paths P_1 and P_2 is the intersection of their contents. The vertices of a sequence (v_1, \dots, v_k) are *aligned* in G if there exists a path P which is the concatenation of $k-1$ paths $P_1 P_2 \dots P_{k-1}$ where P_i is a path from v_i to v_{i+1} for $i \in [k-1]$. Note that we might have $v_i = v_{i+1}$ and then $P_i = v_i$.

Note that, for every connected graph G , if $ab, cd \in E(G)$, $ab \neq cd$, then (a, b, c, d) , (a, b, d, c) , (b, a, c, d) , or (b, a, d, c) are aligned. Moreover, if G is a tree, exactly one of them is aligned. Let G be a connected graph and $a, b, c, d \in V(G)$ such that (a, b, c, d) are aligned. The *in-area* of the two edges ab and cd is the connected component of $G \setminus \{ab, cd\}$ containing the vertices b and c . The other components are called *out-areas*. The following lemma links the connectivity of a flip and the alignment of its vertices:

Lemma 1. *Let G be a connected graph and $ab, cd \in E(G)$ where a, b, c and d are pairwise distinct vertices of G . If (a, b, c, d) or (b, a, d, c) are aligned in G , then the flip $(ab, cd) \rightarrow (ac, bd)$ is connected. If G is a tree, then it is also a necessary condition.*

Proof. The deletion of ab and cd leaves at most three connected components. Let us assume that (a, b, c, d) are aligned, the other case being symmetrical. Let $G_{b,c}$ be the in-area of ab, cd . Let G_a (resp. G_d) be the connected component containing a (resp. d). Note that some of these components might be identical. The addition of ac and bd connects $G_a, G_{b,c}$ and G_d back again. Thus, $(ab, cd) \rightarrow (ac, bd)$ is connected.

Assume now that G is a tree. Supposed that (a, b, c, d) and (b, a, d, c) are not aligned. Then, (a, b, d, c) or (b, a, c, d) are. Thus, the deletion of ab and cd splits G into exactly three components $G_a, G_{b,d}$ and G_c , or $G_b, G_{a,c}$ and G_d . In both cases, when we create ac and bd , we create an edge in the in-area of ab and cd . The resulting graph then contains a cycle, and thus cannot be connected since the total number of edges is still $|V| - 1$. \square

Lemma 1 ensures that, for trees, exactly one of the two flips σ and $-\sigma$ is connected. For paths, we have seen that applying a connected flip is equivalent to reversing the portion of the path between the two involved edges. A similar statement holds for trees:

Remark 1. Let T be a tree. Let e_1, f_1, e_2, f_2 be four pairwise distinct edges of T , and let σ_1 be the flip of e_1, f_1 such that the resulting tree T' is connected. Let P_1 be the path from e_1 to f_1 in T , P_2 be the path from e_2 to f_2 in T , and P'_2 be the path from e_2 to f_2 in T' .

- If both e_2 and f_2 are in the in-area of e_1 and f_1 , $P_2 = P'_2$.
- If both e_2 and f_2 are in the out-areas of e_1 and f_1 , the contents of P_2 and P'_2 are the same, but the order of the portion of the path that corresponds to P_1 is reversed (if it exists).
- If e_2 is in the in-area of e_1 and f_1 , and f_2 is in the out-areas (or the converse), the content of P'_2 is distinct from the content of P_2 . Indeed, the edges that belong to $P_1 \cap P_2$ are changed for the edges of $P_1 \setminus ((P_1 \cap P_2) \cup e_2)$. (See Figure 1 for an illustration of this case).

We can also remark the following:

Remark 2. Let T be a tree and e_1, f_1, e be three pairwise distinct edges. The edge e is in the in-area of e_1, f_1 if and only if it is in the in-area of e'_1, f'_1 where e'_1, f'_1 are the edges created by the unique connected flip on e_1 and f_1 . Moreover e is on the path between e_1 and f_1 in T if and only if e is on the path between e'_1 and f'_1 in the resulting tree.




Fig. 1: The consequences of a connected flip in a tree. The blue thick path goes from an edge of the in-area of ab and cd to an edge of an out-area before the flip, and links the two same edges afterwards.

Let e and f be two vertex-disjoint edges of a tree T , and let σ_2 be a flip in T that does not flip e nor f . The flip σ_2 *depends on* e and f if applying the connected flip on e and f changes the connectivity of σ_2 . By abuse of notation, for any two flips σ_1 and σ_2 on pairwise disjoint edges, σ_2 *depends on* σ_1 if σ_2 depends on the edges of σ_1 . The flip σ_1 *sees* σ_2 if exactly one of the edges of σ_2 is on the path linking the two edges of σ_1 in G .

The following lemma links the dependency of two flips and the position of their edges in a tree:

Lemma 2. Let T be a tree and σ_1 and σ_2 be two flips on T , whose edges are pairwise distinct. The three following points are equivalent:

1. σ_2 depends on σ_1 ,
2. σ_1 depends on σ_2 ,
3. σ_2 sees σ_1 and σ_1 sees σ_2 .

Proof. Since T is a tree, Lemma 1 ensures that exactly one flip amongst σ_1 and $-\sigma_1$ is connected. Moreover this connected flip modifies the connectivity of σ_2 : $(ab, cd) \rightarrow (ac, bd)$ if and only if σ_1 modifies the alignment of a, b, c and d from (a, b, c, d) or (b, a, d, c) to (a, b, d, c) or (b, a, c, d) (or conversely). Equivalently the orientation of one of the two edges ab and cd is modified relatively to the other. Equivalently, by Remark 2, one of the edges ab and cd belongs to the path between the two edges e_1 and f_1 of σ_1 in T , and the other is in an out-area of e_1 and f_1 . Let us call this property $(1')$. We thus have $(1 \Leftrightarrow 1')$. Let us now show that $(1' \Leftrightarrow 3)$. It will indeed give $(1 \Leftrightarrow 3)$ and, by symmetry, $(2 \Leftrightarrow 3)$.

(1' \Rightarrow 3). If one of the edges ab and cd belongs to the path from e_1 to f_1 and the other is in a out-area of e_1 and f_1 , then in particular, one edge is in the in-area of e_1 and f_1 and the other is in an out-area. Thus, exactly one edge of σ_1 is on the path from ab to cd , and σ_2 sees σ_1 . Moreover, one of the edges ab and cd belongs to the path from e_1 to f_1 and the other does not, so that σ_1 sees σ_2 .

(3 \Rightarrow 1'). Since σ_1 sees σ_2 , exactly one edge of σ_2 is on the path from e_1 to f_1 . We can assume without loss of generality that ab is, and cd is not. Moreover, since σ_2 sees σ_1 , exactly one edge of σ_1 is on the path from ab to cd , which means that one is in the in-area of e_1 and f_1 , and the other is in an out-area. Since ab is on the path from e_1 to f_1 , ab is in the in-area of e_1, f_1 . And thus bc is in one out-area of e_1, f_1 . \square

We now give two consequences of applying a connected flip.

Lemma 3. *Let T be a tree and σ_1 and σ_2 be two flips on T with pairwise disjoint edges, where σ_1 is connected. Let T' be the tree obtained after applying σ_1 to T . The flip σ_1^{-1} sees σ_2 in T' if and only if σ_1 sees σ_2 in T . And σ_2 sees σ_1^{-1} in T' if and only if σ_2 sees σ_1 in T .*

Proof. The flip σ_1 sees σ_2 in T whenever exactly one edge of σ_2 is on the path between the edges of σ_1 in T . By Remark 2, the number of edges of $\{e_2, f_2\}$ between the edges of σ_1 is equal to the number of edges of $\{e_2, f_2\}$ between the edges of σ_1^{-1} in T' . Thus σ_1 sees σ_2 in T if and only if σ_1^{-1} sees σ_2 in T' .

On the other hand, σ_2 sees σ_1 if and only if exactly one edge of σ_2 is in the in-area of the edges of σ_1 , and the other is in an out-area. By Remark 2, the same holds in T' for the edges of σ_1^{-1} , and thus σ_2 sees σ_1^{-1} in T' if and only if σ_2 sees σ_1 in T . \square

Lemma 4. *Let T be a tree and σ_1, σ_2 and σ_3 be three flips on T whose edges are pairwise disjoint and such that σ_1 sees σ_2 , σ_2 sees σ_3 , and σ_2 is connected. Let T' be the tree obtained by applying the flip σ_2 to T . The flip σ_1 sees σ_3 in T if and only if σ_1 does not see σ_3 in T' .*

Proof. Let e_1 and f_1 (resp. e_2, f_2 and e_3, f_3) be the edges of σ_1 (resp. σ_2 and σ_3). Let P_1 be the path from e_1 to f_1 in T , P'_1 be the path from e_1 to f_1 in T' , and P_2 be the path from e_2 to f_2 in T .

Since σ_1 sees σ_2 in T , exactly one edge of σ_2 is on the path P_1 . Thus, one edge of σ_1 is in the in-area of e_2 and f_2 , and the other is in an out-area. We can assume without loss of generality that f_1 is in the in-area. Thus, as described in Remark 1, the edges that belong to P'_1 differ from the ones that belong to P_1 in the following way: the portion $P_1 \cap P_2$ is replaced by the portion $P_2 \setminus ((P_1 \cap P_2) \cup f_1)$.

Now, since σ_2 sees σ_3 , exactly one edge of σ_3 is on the path P_2 in T . Thus, exactly one edge of σ_3 is either on $P_1 \cap P_2$ or on $P_2 \setminus ((P_1 \cap P_2) \cup f_1)$. Therefore, in T' , P'_1 has either exactly one edge of σ_3 which is added or removed compared to P_1 .

Thus, if exactly one edge of σ_3 belongs to P_1 , either both or none of the edges of σ_3 belong to P'_1 , and if both or none of the edges of σ_3 belong to P_1 , exactly one edge of σ_3 belongs to P'_1 . This concludes the proof. \square

3 Upper bound

Let us first give a short proof of a result of Bousquet and Mary [3].

Lemma 5. *Let G, H be two connected graphs with the same degree sequence. There exists a sequence of at most two flips that decreases $\delta(G, H)$ by at least 2. Moreover, if there is an alternating C_4 in $\Delta(G, H)$, it can be removed in at most 2 steps, without modifying the rest of the graph.*

Proof. Let \mathcal{C} be a partition of $\Delta(G, H)$ into alternating cycles, and u, v, w, x, y be five consecutive vertices of a cycle C of \mathcal{C} , with $uv, wx \in E(G)$ and $vw, xy \in E(H)$. Note that we may have $y = u$, if C is a C_4 . At least one of the two flips $\sigma_1 : (uv, wx) \rightarrow (uw, vx)$ and $-\sigma_1 : (uv, wx) \rightarrow (ux, vw)$ is connected in G . If $-\sigma_1$ is connected, then we can apply it. Since $vw \in E(H)$, $\delta(G, H)$ decreases by at least 2 (resp. 4 if C is a C_4). Similarly, at least one of the two flips $\sigma_2 : (vw, xy) \rightarrow (vx, wy)$ and $-\sigma_2 : (vw, xy) \rightarrow (vy, wx)$ is connected in H . If $-\sigma_2$ is connected then we can apply it and $\delta(G, H)$ decreases by at least 2 (resp. 4 if C is a C_4). Thus,

we can assume that σ_1 and σ_2 are the only flips that are connected. We apply σ_1 to G and σ_2 to H , and reduce $\delta(G, H)$ by 2, since both flips create the edge vx (resp. 4 if C is a C_4 since both flips also create the edge uw). \square

It immediately implies the following since, in an optimal solution, the size of the symmetric difference decreases by at most 4 at each step.

Corollary 1. *SHORTEST CONNECTED GRAPH TRANSFORMATION admits a polynomial time 4-approximation algorithm.*

The goal of the rest of this section is to improve the approximation ratio. The crucial lemma is the following:

Lemma 6. *Let G, H be two trees with the same degree sequence. There exists a sequence of at most 3 flips that decreases $\delta(G, H)$ by at least 4. Moreover, this sequence only flips bad edges.*

Proof. Let G' be the graph whose vertices are the connected components of $G \cap H$ and where two vertices S_1 and S_2 of G' are incident if there exists an edge in G between a vertex of S_1 and a vertex of S_2 . In other words, G' is obtained from G by contracting every connected component of $G \cap H$ into a single vertex. Note that the edges of G' are the edges of $G - H$. Moreover, as G is a tree, G' also is. We can similarly define H' . Note that G' and H' have the same degree sequence.

Let S_1 be a leaf of G' and S_2 be its parent in G' . Let us show that S_2 is not a leaf of G' . Indeed, otherwise G' would be reduced to a single edge. In particular, $E(G - H)$ would contain only one edge. Since the degree sequence of $G - H$ and $H - G$ are the same, the edge of $H - G$ would have to be the same, a contradiction. Thus, we can assume that S_2 is not a leaf. Let u_1u_2 be the edge of $G - H$ between $u_1 \in S_1$ and $u_2 \in S_2$. Since $G - H$ and $H - G$ have the same degree sequence and S_1 is a leaf of G' , there exists a unique vertex v_1 such that $u_1v_1 \in E(H - G)$. Moreover there exists a vertex v_2 such that $u_2v_2 \in E(H - G)$.

Let us first assume that $v_1 = v_2$. Then there exists a vertex w distinct from u_1 and u_2 such that $v_1w \in E(G - H)$ since v_1 has degree at least 2 in $H - G$. Since S_1 is a leaf of G' , $w \notin S_1$ and either (u_1, u_2, v_1, w) or (u_1, u_2, w, v_1) are aligned in G . If (u_1, u_2, v_1, w) are aligned then the flip $(u_1u_2, v_1w) \rightarrow (u_1v_1, u_2w)$ in G is connected and creates the edge u_1v_1 . If (u_1, u_2, w, v_1) are aligned then $(u_1u_2, v_1w) \rightarrow (u_1w, u_2v_1)$ is connected and creates the edge $u_2v_1 = u_2v_2$. In both cases, we reduce the size of the symmetric difference by at least 2 in one flip, and we can conclude with Lemma 5.

From now on, we assume that $v_1 \neq v_2$. We focus on the alignment of u_1, v_1, u_2 and v_2 in H . Since S_1 is a leaf of G' , it is also a leaf of H' . Thus, v_1 is on the path from u_1 to u_2 and either (u_1, v_1, u_2, v_2) or (u_1, v_1, v_2, u_2) are aligned. If (u_1, v_1, u_2, v_2) are aligned then Lemma 1 ensures that $(u_1v_1, u_2v_2) \rightarrow (u_1u_2, v_1v_2)$ is connected in H and reduces the size of the symmetric difference by at least 2. We can conclude with Lemma 5. Thus, we can assume that (u_1, v_1, v_2, u_2) are aligned in H (see Figure 2 for an illustration).

Let us first remark that if u_2 has degree at least 2 in $H - G$ (or equivalently in $G - H$), then we are done. Indeed, if there exists $w \neq v_2$ such that $u_2w \in E(H - G)$ then, since (u_1, v_1, v_2, u_2) are aligned, (u_1, v_1, u_2, w) have to be aligned. Indeed, v_2u_2 is the only edge of $H - G$ on the path from v_1 to u_2 incident to u_2 . Thus the flip $(u_1v_1, u_2w) \rightarrow (u_1u_2, v_1w)$ is connected in H . Since it reduces $\delta(G, H)$ by at least 2, we can conclude with Lemma 5.

From now on, we will assume that u_2 has degree 1 in $H - G$. Let H_3 (resp. H_4) be the connected component of v_1 and v_2 (resp. u_2) in $H \setminus \{u_1v_1, u_2v_2\}$, which exists since (u_1, v_1, v_2, u_2) are aligned. Note that the third component of $H \setminus \{u_1v_1, u_2v_2\}$ is reduced to S_1 . By definition, H_3 is the in-area of u_1v_1 and u_2v_2 .

We now show that there exists an edge $u_3u_4 \in E(G - H)$, with $u_3 \in H_3$, $u_4 \in H_4$, and such that the connected component S_4 of $G \cap H$ containing u_4 is not a leaf of G' . Indeed, since G is connected, there exists a path P from v_1 to u_2 in G . Since u_1u_2 is the only edge of $G - H$ that has an endpoint in S_1 , this path does not contain any vertex of S_1 . Thus, it necessarily contains an edge u_3u_4 between a vertex u_3 of H_3 and a vertex u_4 of H_4 . Since H_3 and H_4 are anticomplete in $G \cap H$, $u_3u_4 \in E(G - H)$. Moreover, the connected component S_4 of $G \cap H$ containing u_4 is not a leaf of G' , as it is either S_2 which is not a leaf, or P has to leave S_4 at some point with an edge of $G - H$ since P ends in $u_2 \in S_2$.

Since u_3 and u_4 have the same degree in $G - H$ and $H - G$, there exist v_3, v_4 such that $u_3v_3, u_4v_4 \in E(H - G)$. Moreover, since S_4 is not a leaf of G' (and thus of H'), there exists an edge of $H - G$ between a vertex $u_5 \in S_4$ and a vertex $v_5 \in V \setminus S_4$ where $u_5v_5 \neq u_4v_4$.

Let us prove that u_3, v_3, u_4 and v_4 are pairwise distinct. By definition, we have $u_3 \neq v_3, u_4 \neq v_4$ and $u_3 \neq u_4$. Moreover, since $u_3u_4 \in E(G - H)$, $u_3 \neq v_4$ and $u_4 \neq v_3$. Thus, the only vertices that can be identical are v_3 and v_4 . If $v_3 = v_4$, since $u_3 \in H_3, u_4 \in H_4$, and v_2u_2 is the only edge of $H - G$ from H_3 to H_4 , then either $v_3 = v_4 = v_2$ or $v_3 = v_4 = u_2$. In the first case, $u_4 = u_2$ since v_2u_2 is the only edge of $H - G$ from H_3 to H_4 . Thus, u_2 is the endpoint of both u_1u_2 and u_2u_3 in $G - H$. In the second case, u_2 is the endpoint of both u_2u_3 and u_2u_4 in $H - G$. Thus, in both cases, u_2 has degree at least 2 in $H - G$, a contradiction.

We now focus on the alignment of u_3, u_4, v_3 and v_4 in H . If (v_3, u_3, v_4, u_4) or (u_3, v_3, u_4, v_4) are aligned, then the flip $(u_3v_3, u_4v_4) \rightarrow (u_3u_4, v_3v_4)$ is connected in H and reduces the size of the symmetric difference by at least 2, since $u_3u_4 \in E(G - H)$. Note that the flip is well-defined since all the vertices are distinct. Thus, we can conclude with Lemma 5. Therefore, we can assume that (u_3, v_3, v_4, u_4) or (v_3, u_3, u_4, v_4) are aligned in H .

We give, in each case, a sequence of three flips that decreases the size of the symmetric difference by at least 4. We first state the three flips that reduce the symmetric difference in each case and then prove that these sequences of flips can be applied.

Case 1. (u_3, v_3, v_4, u_4) are aligned. (See Figure 2 for an illustration).

We successively apply the flips $\sigma_1 : (u_2v_2, u_5v_5) \rightarrow (u_2v_5, u_5v_2)$, $\sigma_2 : (u_3v_3, u_4v_4) \rightarrow (u_3u_4, v_3v_4)$ where $x = u_5$ if $u_3 = v_2$ and $v_3 = u_2$, and $x = v_3$ otherwise, and $\sigma_3 : (u_1v_1, u_2v_5) \rightarrow (u_1u_2, v_1v_5)$ in H . Since $u_1u_2, u_3u_4 \in E(G - H)$, this sequence of flips indeed reduces $\delta(G, H)$ by at least 4.

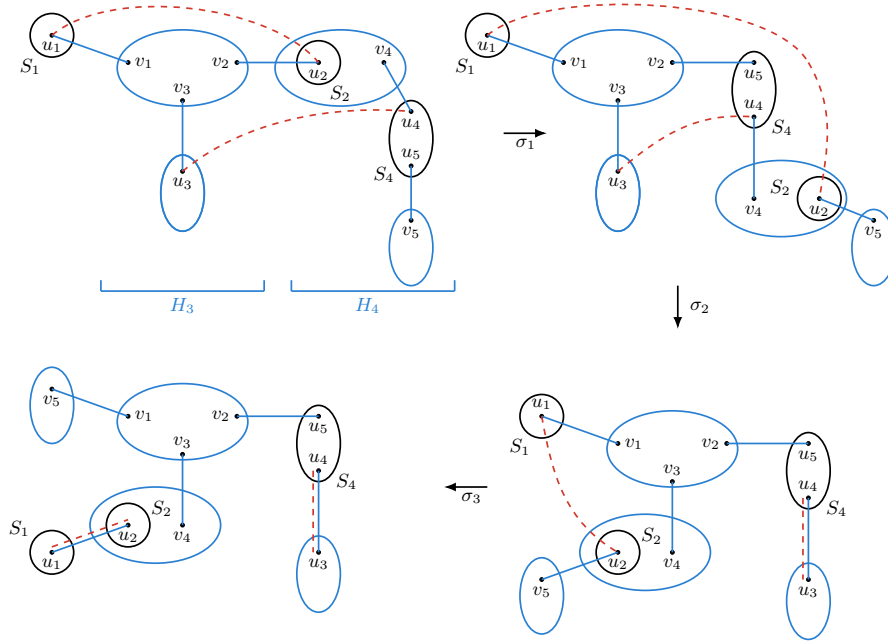


Fig. 2: The three flips $\sigma_1 : (u_2v_2, u_5v_5) \rightarrow (u_2v_5, u_5v_2)$, $\sigma_2 : (u_3v_3, u_4v_4) \rightarrow (u_3u_4, v_3v_4)$ and $\sigma_3 : (u_1v_1, u_2v_5) \rightarrow (u_1u_2, v_1v_5)$ applied to the graph H where (u_3, v_3, v_4, u_4) are aligned. The blue full edges are in $E(H - G)$ and the red dashed edges are in $E(G - H)$.

Let us now show that this sequence of flips can be applied. We first prove that the flip $\sigma_1 : (u_2v_2, u_5v_5) \rightarrow (u_2v_5, u_5v_2)$ is well-defined since the vertices are pairwise distinct. Indeed, by definition, $u_5 \neq v_5$ and $u_2 \neq$

v_2 . Since $u_5 \in H_4$ and $v_2 \notin H_4$, we have $u_5 \neq v_2$. Similarly, let us show that $v_5 \in H_4$, and thus $v_5 \neq v_2$. Firstly, since $u_4 \neq u_2$ (otherwise the degree of u_2 in $G - H$ is at least 2, a contradiction), we have $v_4 \neq v_2$ and thus $v_4 \in H_4$. Moreover, by hypothesis, (u_3, v_4, u_4) are aligned in H , and since $u_4, u_5 \in S_4$ and $v_4, v_5 \notin S_4$, (u_3, v_4, u_5, v_5) are aligned in H . But $u_3 \in H_3$ and $v_4 \in H_4$ where $v_2 u_2$ is the only edge from H_3 to H_4 . This ensures that (v_2, u_2, u_5, v_5) are aligned in H , and since $v_2 u_2$ is the only edge from H_3 to H_4 , $v_5 \in H_4$. Thus we can only have $u_2 = u_5$ or $u_2 = v_5$. But, in both cases, u_2 would have degree at least 2 in $H - G$, a contradiction.

We have shown that (v_2, u_2, u_5, v_5) are aligned in H . Thus, Lemma 1 ensures that σ_1 is connected.

Let H_{σ_1} be the graph obtained after applying σ_1 to H (which is connected). We apply the flip $\sigma_2 : (u_3 x, u_4 v_4) \rightarrow (u_3 u_4, x v_4)$ in H_{σ_1} , where $x = u_5$ if $u_3 = v_2$ and $v_3 = u_2$ and where $x = v_3$ otherwise.

Let us first prove that σ_2 is well-defined. The vertices of σ_2 are pairwise distinct. Indeed, we have previously shown that the vertices u_3, v_3, u_4 and v_4 are pairwise distinct which gives the conclusion in the second case. When $x = u_5$, since $v_4 \notin S_4$ and $u_3 \notin H_4$, we also have $u_5 \neq v_4$ and $u_5 \neq u_3$. Moreover, if $x = u_5$ and $u_5 = u_4$, by definition of x we have $v_2 = u_3$ and σ_1 created the edge $v_2 u_5 = u_3 u_4 \in G - H$, so that we can conclude with Lemma 5. Therefore, all the vertices of σ_2 are distinct.

Let us now show that its two edges, $u_3 x$ and $u_4 v_4$, exist in H_{σ_1} . In order to do it, we have to prove that these edges are not the edges of σ_1 . By definition, we first have $u_4 v_4 \neq u_5 v_5$. Moreover, if $u_4 v_4 = u_2 v_2$, since $v_2 \in H_3$ and $u_4 \notin H_3$, we then have $v_2 = v_4$ and $u_2 = u_4$. Thus, u_2 is the endpoint of both $u_1 u_2$ and $u_2 u_3$ in $G - H$, a contradiction with its degree assumption. Thus we can assume that $u_4 v_4 \neq u_2 v_2$ and $u_4 v_4$ is not equal to any of the edges of σ_1 . Since $u_4 v_4$ is in H , it is in H_{σ_1} . If $x = v_3$ then $u_3 v_3 \neq u_2 v_2$ by definition of x . And $u_3 v_3$ with both endpoints in H_3 is distinct from $u_5 v_5$ which has both endpoints in H_4 . If $x = u_5$, then $u_3 = v_2$ and $v_3 = u_2$. But in this case, $u_3 x = v_2 u_5$ was created by σ_1 and thus is in H_{σ_1} . So both edges of σ_2 exist in H_{σ_1} and then σ_2 can be applied.

We now show that σ_2 is connected in H_{σ_1} . By hypothesis, (u_3, v_3, v_4, u_4) are aligned in H . Moreover, as $u_4, u_5 \in S_4$ and $v_4, v_5 \notin S_4$, (v_4, u_4, u_5, v_5) are aligned. Therefore, if $u_3 = v_2$ and $v_3 = u_2$, $(u_3, v_3, v_4, u_4, u_5, v_5)$ are aligned in H and $(u_3, u_5, u_4, v_4, v_3, v_5)$ are aligned in H_{σ_1} . In particular, since $x = u_5$ in this case, (u_3, x, u_4, v_4) are aligned and σ_2 is connected. Otherwise, $(u_3, v_3, v_2, u_2, v_4, u_4, u_5, v_5)$ are aligned in H and $(u_3, v_3, v_2, u_5, u_4, v_4, u_2, v_5)$ are aligned in H_{σ_1} . In particular, since $x = v_3$ in this case, (u_3, x, u_4, v_4) are aligned and σ_2 is also connected.

Let H_{σ_2} be the graph obtained after applying σ_2 to H_{σ_1} . We want to apply the flip $\sigma_3 : (u_1 v_1, u_2 v_5) \rightarrow (u_1 u_2, v_1 v_5)$ in H_{σ_2} . Let us first prove that it is well-defined. By definition, we have $u_1 \neq u_2$ and $u_1 \neq v_1$. Since $v_1 \in H_3$ and $u_2 \notin H_3$, $v_1 \neq u_2$. Since $v_5 \in H_4$ while $u_1, v_1 \notin H_4$, we have $v_5 \neq u_1$ and $v_5 \neq v_1$. Finally, $v_5 \neq u_2$ was proven before applying σ_1 . So the vertices of σ_3 are pairwise distinct. Let us now prove that both $u_1 v_1$ and $u_2 v_5$ exist in H_{σ_2} . Since u_1 is the only vertex of S_1 defined in our construction and u_1 does not appear as an endpoint in σ_1 and σ_2 , $u_1 v_1$ exists in H_{σ_2} . The edge $u_2 v_5$ is created by σ_1 , so $u_2 v_5 \in E(H_{\sigma_1})$. Since $u_2, v_5 \in H_4$ and $u_3 \notin H_4$, we have $u_2 v_5 \neq u_3 x$. Moreover, $v_5 \notin S_4$ and $v_5 \neq v_4$ and then $u_2 v_5 \neq u_4 v_4$. Thus $u_2 v_5$ is not an edge of σ_2 , and $u_2 v_5 \in E(H_{\sigma_2})$.

In order to prove that σ_3 is connected, we will use Lemma 2. Let us first prove that σ_3 is connected in H_{σ_1} . In H , (u_1, v_1, v_2, u_2) and (v_2, u_2, u_5, v_5) are aligned. Thus, in H_{σ_1} , $(u_1, v_1, v_2, u_5, u_2, v_5)$ are aligned. In particular, (u_1, v_1, u_2, v_5) are aligned and σ_3 is connected in H_{σ_1} .

Finally, we prove that in H_{σ_1} , σ_3 does not depend on σ_2 . We claim that σ_2 does not see σ_3 , as none of its two edges are on the path from $u_3 x$ to $u_4 v_4$. Since S_1 is a leaf of $G \cap H$, $u_1 v_1$ is not on it. If $x = v_3$, since $(u_3, v_3, v_2, u_5, u_4, v_4, u_2, v_5)$ are aligned in H_{σ_1} , $u_2 v_5$ is not on it either, and if $x = u_5$, $(u_3, u_5, u_4, v_4, v_3, v_5)$ are aligned in H_{σ_1} but since $u_2 = v_3$ in this case, $u_2 v_5$ is not on the path either. Thus, by Lemma 2, σ_2 and σ_3 are independent. Therefore, σ_3 is still connected in H_{σ_2} .

Case 2. (v_3, u_3, u_4, v_4) are aligned.

We apply $\sigma_1 : (u_2 v_2, u_4 v_4) \rightarrow (u_2 v_4, u_4 v_2)$, $\sigma_2 : (u_3 v_3, u_4 v_2) \rightarrow (u_3 u_4, v_2 v_3)$ then $\sigma_3 : (u_1 v_1, u_2 v_4) \rightarrow (u_1 u_2, v_1 v_4)$ to H . Again, $u_1 u_2, u_3 u_4 \in E(G - H)$ and it reduces $\delta(G, H)$ by at least 4.

Let us first prove that we can apply this sequence of flips. We first prove that the vertices of σ_1 are pairwise distinct. Since, by hypothesis, (v_3, u_3, u_4, v_4) are aligned in H , with $u_3 \in H_3$ and $u_4 \in H_4$, we have that (v_2, u_2, u_4, v_4) are aligned in H . By definition, $v_2 \neq u_2$ and $u_4 \neq v_4$. Thus, the only vertices that might

be identical are u_2 and u_4 . But if $u_2 = u_4$, then u_2 is both the endpoint of u_1u_2 and u_2u_3 in $G - H$. Thus, u_2 has degree at least 2 in $G - H$, a contradiction.

Moreover, since (v_2, u_2, u_4, v_4) are aligned, σ_1 is connected.

Let H_{σ_1} be the graph obtained after applying σ_1 to H . We apply in H_{σ_1} the flip $\sigma_2 : (u_3v_3, u_4v_2) \rightarrow (u_3u_4, v_2v_3)$.

Let us show that its vertices are pairwise distinct. We have previously shown that $u_3, v_3 \neq u_4$. By definition, we have $u_3 \neq v_3$. Since $u_4 \in S_4$ and $v_2 \notin S_4$, $u_4 \neq v_2$. Now, if $v_2 = u_3$, then σ_1 created the edge $u_3u_4 \in E(G - H)$ and we can conclude with Lemma 5. Finally, since (v_3, u_3, u_4, v_4) are aligned in H with $v_3, u_3 \in H_3$ and $v_4, u_4 \in H_4$ and since v_2u_2 is the only edge of H from H_3 to H_4 , we know that $(v_3, u_3, v_2, u_2, u_4, v_4)$ are aligned in H . As $v_3 \neq u_3$, this gives $v_3 \neq v_2$.

Let us now show that the two edges of σ_2 , u_3v_3 and u_4v_2 , exist in H_{σ_1} . We know that u_4v_2 is created by σ_1 , so that $u_4v_2 \in E(H_{\sigma_1})$. Moreover, $u_3v_3 \in E(H)$. Thus, we only have to show that u_3v_3 is not an edge of σ_1 . Since $u_3, v_3 \neq u_4$ and $u_3, v_3 \neq v_2$, it is straightforward.

Finally, let us show that σ_2 is connected in H_{σ_1} . We have seen that $(v_3, u_3, v_2, u_2, u_4, v_4)$ are aligned in H . Thus, in H_{σ_1} , $(v_3, u_3, v_2, u_4, u_2, v_4)$ are aligned. In particular, (v_3, u_3, v_2, u_4) are aligned and, by Lemma 1 σ_2 is connected.

Let H_{σ_2} be the graph obtained after applying σ_1 to H_{σ_1} . We apply in H_{σ_2} the flip $\sigma_3 : (u_1v_1, u_2v_4) \rightarrow (u_1u_2, v_1v_4)$. Let us prove that the endpoints of its edges are all distinct, that both its edges exist in H_{σ_2} , and that it is connected.

We have seen that (u_1, v_1, v_2, u_2) are aligned in H , and since $v_4 \in H_4$, (u_1, v_1, u_2, v_4) are aligned in H . Moreover, by definition, $u_1 \neq v_1$, we have previously shown that $u_2 \neq v_4$, and since $v_1 \in H_3$ and $u_2 \notin H_3$, $u_2 \neq v_1$. Therefore, the vertices of σ_3 are all pairwise distinct.

Let us now prove that its edges exist in H_{σ_2} . Since u_1 is the only vertex of S_1 we considered, we know that it is distinct from all the other vertices and thus u_1v_1 is distinct from all the other edges. Therefore, it is not an edge of σ_1 nor σ_2 and since $u_1v_1 \in E(H)$, $u_1v_1 \in E(H_{\sigma_2})$. Since u_2v_4 is created by σ_1 , we have $u_2v_4 \in E(H_{\sigma_1})$. Since $u_2, v_4 \in H_4$ and $u_3, v_2 \notin H_4$, we have $u_2v_4 \neq u_3v_3$ and $u_2v_4 \neq u_4v_2$. Thus, u_2v_4 is not identical to any edge of σ_2 , and $u_2v_4 \in E(H_{\sigma_2})$.

We now show that σ_3 is connected in H_{σ_2} . Since (u_1, v_1, v_2, u_2) and (v_2, u_2, u_4, v_4) are aligned in H , $(u_1, v_1, v_2, u_2, u_4, v_4)$ are aligned in H and $(u_1, v_1, v_2, u_4, u_2, v_4)$ are aligned in H_{σ_1} . In particular, (u_1, v_1, u_2, v_4) are aligned and σ_3 is connected in H_{σ_1} .

Let us prove that in H_{σ_1} , σ_3 does not depend on σ_2 . We claim that σ_2 does not see σ_3 . Since S_1 is a leaf of $G \cap H$, u_1v_1 is not on the path from u_3v_3 to v_2u_4 in H_{σ_1} . Moreover, since $(v_3, u_3, v_2, u_4, u_2, v_4)$ are aligned in H_{σ_1} , u_2v_4 is not on it either. Thus, by Lemma 2, σ_2 and σ_3 are independent. Therefore, σ_3 is still connected in H_{σ_2} .

Therefore, in all the cases, we have found a sequence of three flips whose edges are in the symmetric difference and that reduce $\delta(G, H)$ by at least 4. Moreover, the proof immediately provides a polynomial time algorithm to find such a sequence. \square

Note that Lemma 6 allows to obtain a 3-approximation algorithm for SHORTEST CONNECTED GRAPH TRANSFORMATION. Indeed, as shown in the proof of Lemma 1 in [3], as long as there exists an edge of the symmetric difference in a cycle of G , one can reduce the size of the symmetric difference by 2 in one step. Afterwards, we can assume that the remaining graphs $G - H$ and $H - G$ are trees. By Lemma 6, in three flips, the symmetric difference of the optimal solution decreases by at most 12 while our algorithm decreases it by at least 4. (Note that free to try all the flips, finding these flips is indeed polynomial). But we can actually improve the approximation ratio. The idea consists in treating differently short cycles. A *short* cycle is a C_4 , a *long* cycle is a cycle of length at least 6. We now give the main result of this section.

Theorem 3. SHORTEST CONNECTED GRAPH TRANSFORMATION admits a $5/2$ -approximation algorithm running in polynomial time. It becomes a $9/4$ -approximation algorithm if $\Delta(G, H)$ does not contain any short cycle.

Proof. Let \mathcal{C} be an optimal partition of $\Delta(G, H)$ into alternating cycles, i.e. a partition with $mnc(G, H)$ cycles. Let c be the number of short cycles in \mathcal{C} . Bereg and Ito [1] provide a polynomial time algorithm to find a partition of $\Delta(G, H)$ into alternating cycles having at least $\frac{c}{2}$ short cycles. Lemma 5 ensures that we can

remove their $2c$ edges from the symmetric difference in at most c flips. If an edge of the symmetric difference is in a cycle of G or H , then in one step we can reduce the symmetric difference by 2 [3]. Otherwise, by Lemma 6, we can remove the remaining $\delta(G, H) - 2c$ edges using at most $\frac{3(\delta(G, H) - 2c)}{4}$ flips in polynomial time. Therefore, we can transform G into H with at most $c + \frac{3(\delta(G, H) - 2c)}{4}$ flips.

Let us now provide a lower bound on the length of a shortest transformation from G to H . By definition, \mathcal{C} contains c short cycles. Theorem 2 ensures that we need at least c steps to remove the short cycles, plus $\ell - 1$ flips to remove each cycle of length 2ℓ . Therefore, we need at least $\frac{\delta(G, H) - 4c}{3}$ flips to remove the $\delta(G, H) - 4c$ remaining edges from the symmetric difference.

The ratio between the upper bound and the lower bound is

$$f(c) := \frac{c + \frac{3\delta(G, H) - 6c}{4}}{c + \frac{\delta(G, H) - 4c}{3}} = \frac{3(3\delta(G, H) - 2c)}{4(\delta(G, H) - c)}.$$

The function f being increasing and since the number of short cycles in \mathcal{C} cannot exceed $\frac{\delta(G, H)}{4}$, we have $f(c) \leq f(\frac{\delta(G, H)}{4}) = \frac{5}{2}$. It gives a $\frac{5}{2}$ -approximation in polynomial time. Moreover, when there is no alternating short cycle in $\Delta(G, H)$, $c = 0$. Since $f(0) = \frac{9}{4}$, we obtain a $\frac{9}{4}$ -approximation. \square

4 Discussion on the tightness of the lower bound

In this section, we discuss the quality of the lower bound of Theorem 2. We first prove that if we only flip bad edges of the same cycle of the symmetric difference then the length of a shortest transformation can be almost twice longer than the one given by the lower bound of Theorem 2. In order to prove it, we generalize several techniques and results of Christie [5], proved for the SORTING BY REVERSALS problem.

Note that the result of Hannenhalli and Pevzner [9] actually proves that in the case of paths, when the symmetric difference only contains vertex-disjoint cycles, it is not necessarily optimal to only flip edges of the same cycle. However, studying this restriction gives us a better understanding of the general problem.

We also prove that, if we only flip bad edges (which are not necessarily in the same cycle of the symmetric difference), then the length of a shortest transformation can be almost $3/2$ times longer than the one given by the lower bound. Note that all the existing approximation algorithms for SORTING BY REVERSALS and SHORTEST CONNECTED GRAPH TRANSFORMATION only flip bad edges. But again no formal proof guarantees that there always exists a shortest transformation where we only flip bad edges.

Both results are obtained with the same graphs G_k and H_k represented in Figure 3 for $k = 4$. For any $k \geq 2$, let $G_k = (V_k, E(G_k))$ and $H_k = (V_k, E(H_k))$ be the graphs with $V_k = \{v_{i,j}, 1 \leq i \leq k, 1 \leq j \leq 4\} \cup \{c\}$, $E(G_k) = \bigcup_{i \in [k]} \{cv_{i,1}, v_{i,1}v_{i,2}, v_{i,2}v_{i,3}, v_{i,3}v_{i,4}\}$, and $E(H_k) = \bigcup_{i \in [k]} \{cv_{i,1}, v_{i,1}v_{i+1,3}, v_{i,2}v_{i,3}, v_{i,2}v_{i+1,4}\}$, where the additions are defined modulo k . One can easily check that, in this construction, both G_k and H_k are the subdivisions of a star where each branch has 4 vertices. Note that $\Delta(G, H)$ is the disjoint union of k short cycles. Moreover, the partition of $\Delta(G, H)$ into alternating cycles is unique.

4.1 Flipping bad edges of the same cycle

Let G and H be two trees with the same degree sequence. The *digraph of flips* $\mathcal{F}(G, H)$ of G and H is the labelled directed graph whose vertices are the good flips in $G - H$ (i.e. the flips that create at least one edge of $G \cap H$, regardless of the fact that they maintain the connectivity of G or not). Every vertex σ is labelled as a connected or non-connected flip. And (σ_1, σ_2) is an arc of $\mathcal{F}(G, H)$ if and only if σ_1 sees σ_2 . Note that every vertex of the digraph of flips corresponds to a good flip σ in $G - H$ and thus corresponds to a pair of edges of $G - H$. Since there exists a connected flip between any pair of edges in a tree, if σ is disconnected, then $-\sigma$ is connected and thus any vertex of $\mathcal{F}(G, H)$ can be associated to a connected flip, either itself or its opposite.

If G and H are paths, if exactly one of the two edges of a flip σ_1 is on the path between the two edges of a flip σ_2 , then exactly one of the two edges of σ_2 is on the path between the two edges of σ_1 . Thus, for paths,




Fig. 3: The graphs G_4 and H_4 . The black thick edges are in $E(G_4 \cap H_4)$, the blue thin edges are in $E(G_4 - H_4)$ and the red dashed edges are in $E(H_4 - G_4)$.




Fig. 4: The digraph of flips of G_4 and H_4 , where $\sigma_{i,j} := (v_{i,1}v_{i,2}, v_{j,3}v_{j,4}) \rightarrow (v_{i,1}v_{j,3}, v_{i,2}v_{j,4})$ for any i and j , and the label D stands for disconnected.

the digraph of flips is a non-oriented graph, and it corresponds to the reversal graph introduced by Christie [5]. The reversal graph is related to the interleaving graph introduced in [9] for the sorting by reversals of signed permutations (which corresponds to the case where G and H are paths and the partition \mathcal{C} of $\Delta(G, H)$ into alternating cycles is unique). The vertices of the interleaving graph are the cycles of \mathcal{C} , labeled as connected if there exists a connected good flip between two edges of \mathcal{C} and disconnected otherwise, and there is an edge between two cycles C_1 and C_2 if there exists a connected flip between two edges of C_1 that changes the connectivity of a flip between two edges of C_2 . Again, for paths, the converse is also true and the graph is therefore non-directed.

For paths, Christie gives in [5] a characterization of the resulting reversal graph when we apply a flip (of bad edges). Unfortunately, his proof cannot be extended easily to the case of trees for the digraph of flips. Indeed, the arcs and the labels above are not enough to determine the connectivity of the flips in the resulting graph. However, when we restrict to the case where the partition \mathcal{C} of $\Delta(G, H)$ into alternating cycles is unique and only contains short cycles, the model becomes simple enough to be understood. The first part of this section consists in proving that, under these assumptions, we can characterize the resulting digraph of flips when we apply a good flip (or, if the flip is disconnected, its opposite) on the digraph of flips.

Two graphs G and H with the same degree sequence are *close* if $\Delta(G, H)$ admits a unique partition into alternating cycles \mathcal{C} and all the cycles of \mathcal{C} are short. Note that G_k and H_k are close. All along the proofs of the section, we will implicitly use the following remarks:

Remark 3. Let G, H be two close trees. Let \mathcal{C} be the unique partition of $\Delta(G, H)$ into alternating cycles. The digraph of flips contains $|\mathcal{C}|$ flips and each vertex corresponds to the unique good flip that removes some cycle C of the symmetric difference. Moreover, if we flip edges on the same cycle of the symmetric difference, then we either use the flips of the digraph of flips or their opposite.

Remark 4. Let G, H be two close trees. If we apply any connected flip between edges of the same cycle of $G - H$, then, the resulting graph G' and H are still close.

Proof. Since G and H are close, the decomposition \mathcal{C} of $\Delta(G, H)$ into alternating cycles is unique and only contains short cycles. Thus, if we apply a connected flip $(ab, cd) \rightarrow (ac, bd)$, with ab and cd in the same cycle of \mathcal{C} , then (a, b, d, c, a) or (a, b, c, d, a) is an alternating C_4 in $\Delta(G, H)$. In the first case, σ is good and the corresponding C_4 disappears from $\Delta(G, H)$ and in the second case, it is replaced by (a, c, b, d, a) . Since the edges of the other cycles are unchanged, in both cases, the partition of $\Delta(G', H)$ into alternating cycles remains unique and still contains only short cycles (where G' is the resulting graph).

A vertex σ' of a digraph of flips \mathcal{F}' replaces a vertex σ of a digraph of flips \mathcal{F} if $V(\mathcal{F}') = (V(\mathcal{F}) \setminus \sigma) \cup \sigma'$, and the label, the in-neighbours and the out-neighbours of σ' in \mathcal{F}' are exactly the label, the in-neighbours and the out-neighbours of σ in \mathcal{F} .

Lemma 7. Let G and H be two close trees and let $\mathcal{F} := \mathcal{F}(G, H)$. Let $\sigma \in V(\mathcal{F})$ and G' be the graph obtained from G by applying σ if it is connected and by applying $-\sigma$ otherwise. The graph $\mathcal{F}' := \mathcal{F}(G', H)$ is characterized as follows:

1. If σ is connected, then $V(\mathcal{F}') = V(\mathcal{F}) \setminus \{\sigma\}$. If σ is disconnected, then $-((- \sigma)^{-1})$ replaces σ .
2. For every $\sigma_1 \notin N_{\mathcal{F}}^-(\sigma) \cap N_{\mathcal{F}}^+(\sigma)$, σ_1 is connected in \mathcal{F} if and only if σ_1 is connected in \mathcal{F}' .
3. For every $\sigma_1 \in N_{\mathcal{F}}^-(\sigma) \cap N_{\mathcal{F}}^+(\sigma)$, σ_1 is connected in \mathcal{F} if and only if σ_1 is disconnected in \mathcal{F}' .
4. For every $\sigma_1, \sigma_2 \in V(\mathcal{F}')$, with $\sigma_1 \notin N_{\mathcal{F}}^-(\sigma)$ or $\sigma_2 \notin N_{\mathcal{F}}^+(\sigma)$, $\sigma_1 \sigma_2 \in E(\mathcal{F})$ if and only if $\sigma_1 \sigma_2 \in E(\mathcal{F}')$.
5. For every $\sigma_1 \in N_{\mathcal{F}}^-(\sigma)$ and every $\sigma_2 \in N_{\mathcal{F}}^+(\sigma)$ such that $\sigma_1 \neq \sigma_2$, $\sigma_1 \sigma_2 \in E(\mathcal{F})$ if and only if $\sigma_1 \sigma_2 \notin E(\mathcal{F}')$.

Proof. Let a, b, c and d be the vertices of G such that $\sigma = (ab, cd) \rightarrow (ac, bd)$. First, notice that since the partition of $\Delta(G, H)$ is unique and only contains short cycles, and since we only apply good flips, or opposites of good flips, all the different flips we consider here are on disjoint edges.

Proof of (1).

First note that after applying the flip, all the flips of the graph of flips distinct from σ still exist. Indeed, all the cycles of \mathcal{C} distinct from the one of σ are still in \mathcal{C}' . Thus, by Remark 3, the set of vertices $V(\mathcal{F}) \setminus \sigma$ is in $V(\mathcal{F}')$. Moreover, all the cycles of \mathcal{C}' distinct from the one of the edges created by σ are also in \mathcal{C} . Thus, by Remark 3, $V(\mathcal{F}') \setminus \{((- \sigma)^{-1}), -((- \sigma)^{-1})\}$ is in $V(\mathcal{F})$.

If σ is connected, then the number of cycles in the partition decreases by one and the vertex is removed. The vertex corresponding to σ disappears but all the other vertices still exist.

Let us now show that if σ is disconnected, then $-((- \sigma)^{-1})$ replaces σ . Firstly, if σ is disconnected in G , then $-\sigma$ is applied to G . Since the partition of $\Delta(G, H)$ into alternating cycles is unique, each vertex of $\Delta(G, H)$ is incident to exactly one edge of $G - H$ and one edge of $H - G$, and as σ is a good flip, $-\sigma$ is not. Thus, the edges created by $-\sigma$ are in $\Delta(G', H)$. That being said, $(-\sigma)^{-1}$ is not a good flip of \mathcal{F}' . Indeed, $(-\sigma)^{-1} = (ad, bc) \rightarrow (ab, cd)$, and we know that ab and cd are in $\Delta(G, H)$. On the other hand, the flip $-((- \sigma)^{-1}) = (ad, bc) \rightarrow (ac, bd)$ is good, as it creates the same edges as σ . Therefore, $-((- \sigma)^{-1}) \in V(\mathcal{F}')$.

Moreover, $-((- \sigma)^{-1})$ is disconnected. Indeed otherwise σ would be connected in \mathcal{F} since σ and $-((- \sigma)^{-1})$ create the same edges. Thus, $-((- \sigma)^{-1})$ has in \mathcal{F}' the label of σ in \mathcal{F} .

Finally, by Lemma 3, $(-\sigma)^{-1}$ has the same in and out neighbourhoods as $-\sigma$. Thus, as $-((- \sigma)^{-1})$ is flipping the same edges as $(-\sigma)^{-1}$, and $-\sigma$ is flipping the same edges as σ , $-((- \sigma)^{-1})$ and σ have the same in and out neighbourhoods.

Proof of (2) and (3).

The points 2 and 3 are a direct consequence of Lemma 2: the label of σ is considering the fact that the edges, and thus the in-neighbours and out-neighbours, of σ and $-\sigma$ are the same.

Proof of (4).

By Remark 1, applying a flip on the edges e and f to a tree T can modify the content of a path P only if one endpoint of P is in the in-area of e and f and the other is in their out-area. When it changes, the portion $P_1 \cap P$ of P is replaced by $P_1 \setminus (P_1 \cap P)$, where P_1 is the path from e to f in T .

Suppose that an arc $\sigma_1\sigma_2$ is in \mathcal{F} but not in \mathcal{F}' , or conversely, and let e_1 and f_1 be the edges of σ_1 , and e_2 and f_2 be the edges of σ_2 . The content of the path P_1 from e_1 to f_1 in G is different from the content of the path P'_1 from e_1 to f_1 in G' , as either P_1 contains exactly one edge of σ_2 and P'_1 contains both or neither edges of σ_2 , or conversely. Thus, one of the edges of σ_1 is in the in-area of ab and cd , and the other is in an out-area. Therefore, σ_1 sees σ in G .

Moreover, in P'_1 , exactly one edge of σ_2 is either added or removed compared to the content of P_1 . Thus, either one edge is on the portion of the path from ab to cd that is common to P_1 and none were on the other portion, or the opposite. Thus, exactly one edge of σ_2 is on the path from ab to cd in G , and σ sees σ_2 in G .

Proof of (5).

Finally, point 5 is a consequence of Lemma 4. \square

Note that Lemma 7 generalizes the results of [5] when G, H are paths. Indeed, in that case, the graph is non-directed and then the subgraph induced by the neighbourhood of σ in \mathcal{F} is complemented after the flip.

Lemma 8. *Let G and H be two close trees. Every disconnected flip of $\mathcal{F}(G, H)$ belongs to an oriented cycle in $\mathcal{F}(G, H)$.*

Proof. Let $\mathcal{F} := \mathcal{F}(G, H)$. Assume by contradiction that there exists a disconnected flip $\sigma \in V(\mathcal{F})$ such that σ does not belong to any oriented cycle of \mathcal{F} .

Since the the partition of $\Delta(G, H)$ into alternating cycles only contains short cycles, the proof of Lemma 5 ensures that there exists a sequence of flips transforming G into H that only flips edges that are in the same short cycle. In other words, there always exists a sequence of flips using flips of \mathcal{F} that transforms G into H . By Remark 4, all the intermediate graphs and H are close, so Lemma 7 holds at any step. Moreover, during such a transformation, every vertex of \mathcal{F} has to be removed at some point (since $\mathcal{F}(H, H)$ is empty). By Lemma 7.1, a vertex σ_2 can be removed only if σ_2 is connected and we apply σ_2 . Thus, the label of σ_2 has to change at some point. Lemma 7 ensures that for if the label of σ_2 changes then σ_2 is the in- and out-neighbourhood of one flip, and thus is in a oriented cycle of \mathcal{F} (of size 2).

Let \mathcal{F}_1 be the last step where σ_2 is not in a cycle of the digraph of flips. Assume that we apply a flip σ_1 and let \mathcal{F}'_1 be the new digraph of flips. Let us prove by contradiction that σ_2 was in a cycle of \mathcal{F}_1 .

Let C' be a cycle of \mathcal{F}'_1 containing σ_2 . If σ_2 does not belong to any oriented cycle in \mathcal{F}_1 , since no vertices have been added to \mathcal{F}_1 , there exists an arc $\sigma_3\sigma_4$ of C' that is not in \mathcal{F}_1 . By Lemma 7, $\sigma_3\sigma_1, \sigma_1\sigma_4 \in E(\mathcal{F}_1)$. By replacing every arc $\sigma_3\sigma_4$ of C' that is not in \mathcal{F}_1 by the two arcs $\sigma_3\sigma_1$ and $\sigma_1\sigma_4$, we obtain a union of oriented cycles of \mathcal{F}_1 , one of them containing σ_2 . Therefore, σ_2 belongs to an oriented cycle in \mathcal{F}_1 , a contradiction. \square

Note that Lemma 8 generalizes a result of Christie [5] which ensures that when G, H are paths, no disconnected flip is isolated in the graph of flips.

Lemma 9. *Let G and H be two close trees. Let $\mathcal{F} := \mathcal{F}(G, H)$ and \mathcal{C} be the unique partition of $\Delta(G, H)$ into alternating cycles. If we only flip pairs of edges that are in the same cycle of \mathcal{C} , then the shortest transformation from G to H has length at least $|V(\mathcal{F})| + \gamma(\mathcal{F})$, where $\gamma(\mathcal{F})$ is defined as follows:*

- If there is no oriented cycle in \mathcal{F} , or if there exists an oriented cycle in \mathcal{F} that only contains connected flips, $\gamma(\mathcal{F}) = 0$.
- Otherwise, $\gamma(\mathcal{F}) = nd(\mathcal{F}) - 1$, where $nd(\mathcal{F})$ is the minimum number of disconnected flip in any oriented cycle of \mathcal{F} .

Proof. We prove Lemma 9 by induction on $(|V(\mathcal{F})| + \gamma(\mathcal{F}))$. If $(|V(\mathcal{F})| + \gamma(\mathcal{F})) = 0$, then \mathcal{F} is the empty graph and thus $G = H$. Assume now that $(|V(\mathcal{F})| + \gamma(\mathcal{F})) \geq 1$. Let σ be a good flip between two edges

that are in the same cycle of \mathcal{C} . Either σ or $-\sigma$ is connected. Let G' be the graph obtained after applying the connected flip, σ or $-\sigma$, and $\mathcal{F}' := \mathcal{F}(G', H)$. We will show that $(|V(\mathcal{F}')| + \gamma(\mathcal{F}')) - (|V(\mathcal{F})| + \gamma(\mathcal{F})) \geq -1$. Since we only flip edges of the same short cycle, Remark 4 ensures that G' and H are close and the proof immediately follows.

By Lemma 7.1, we have $|V(\mathcal{F}')| - |V(\mathcal{F})| = -1$ if σ is connected and 0 if σ is disconnected. We call this property the property (*). Let us now bound the quantity $\gamma(\mathcal{F}') - \gamma(\mathcal{F})$.

Assume that there exists $\sigma_1 \in V(\mathcal{F})$ such that $\sigma_1 \in N_{\mathcal{F}}^-(\sigma) \cap N_{\mathcal{F}}^+(\sigma)$. Note that σ, σ_1 forms a cycle of size 2. If σ or σ_1 are connected then $\gamma(\mathcal{F}) = 0$. So we have $\gamma(\mathcal{F}') - \gamma(\mathcal{F}) \geq 0$. If both σ and σ_1 are disconnected then $\gamma(\mathcal{F}) \leq 1$. And then $\gamma(\mathcal{F}') - \gamma(\mathcal{F}) \geq -1$. Thus, in both cases, the property (*) gives $(|V(\mathcal{F}')| + \gamma(\mathcal{F}')) - (|V(\mathcal{F})| + \gamma(\mathcal{F})) \geq -1$ and the result is proven.

So we can assume that no vertex σ_1 of \mathcal{F} satisfies $\sigma_1 \in N_{\mathcal{F}}^-(\sigma) \cap N_{\mathcal{F}}^+(\sigma)$. By Lemma 7.2 and 7.3, the labels of the vertices of \mathcal{F}' and \mathcal{F} are the same. Thus, if $\gamma(\mathcal{F}') < \gamma(\mathcal{F})$, it is because \mathcal{F}' has no oriented cycle, or because at least one cycle of \mathcal{F}' is not in \mathcal{F} . Let us consider both cases.

Suppose that \mathcal{F}' has no oriented cycle. Lemma 8 ensures that every flip of \mathcal{F}' is connected. Since the labels are the same in \mathcal{F} and \mathcal{F}' , only σ can be disconnected in \mathcal{F} if it has been removed in \mathcal{F}' , but by Lemma 7, if σ has been removed, σ is connected in \mathcal{F} . Thus, \mathcal{F} only contains connected flips, and either there are no oriented cycles in \mathcal{F} , or the only oriented cycles contain only connected flips. Thus, combining $\gamma(\mathcal{F}) = \gamma(\mathcal{F}') = 0$ with property (*), we have $(|V(\mathcal{F}')| + \gamma(\mathcal{F}')) - (|V(\mathcal{F})| + \gamma(\mathcal{F})) \geq -1$.

So we can assume that at least one oriented cycle C' of \mathcal{F}' is not in \mathcal{F} . Since $V(\mathcal{F}') \subseteq V(\mathcal{F})$ and the labels are the same in \mathcal{F} and \mathcal{F}' , at least one arc of C' is not in \mathcal{F} . Let us prove that there exists a cycle C in \mathcal{F} such that $V(C) \subseteq V(C') \cup \{\sigma\}$.

If exactly one arc $\sigma_1\sigma_2$ of C' is not in \mathcal{F} then, by Lemma 7.4 and 7.5, $\sigma_1 \in N_{\mathcal{F}}^-(\sigma)$ and $\sigma_2 \in N_{\mathcal{F}}^+(\sigma)$. Moreover, the path P of C' from σ_2 to σ_1 is also in \mathcal{F} , so that in \mathcal{F} , P plus $\sigma_1\sigma$ and $\sigma\sigma_2$ form an oriented cycle C in \mathcal{F} , with $V(C) \subseteq V(C') \cup \{\sigma\}$. If at least two distinct arcs of C' are not in \mathcal{F} , let $\sigma_1\sigma_2$ and $\sigma_3\sigma_4$ be two such arcs. We can choose $\sigma_1\sigma_2$ and $\sigma_3\sigma_4$ so that the oriented path P from σ_2 to σ_3 in C' only contains arcs that are in \mathcal{F} . Lemma 7 ensures that $\sigma_1, \sigma_3 \in N_{\mathcal{F}}^-(\sigma)$ and $\sigma_2, \sigma_4 \in N_{\mathcal{F}}^+(\sigma)$. If $\sigma_2 = \sigma_3$ or $\sigma_4 = \sigma_1$, $N_{\mathcal{F}}^-(\sigma) \cap N_{\mathcal{F}}^+(\sigma)$ is not empty, a contradiction with the assumptions. Thus, $\sigma_1\sigma_2$ and $\sigma_3\sigma_4$ are not consecutive in C' . Since $\sigma_3 \in N_{\mathcal{F}}^-(\sigma)$ and $\sigma_2 \in N_{\mathcal{F}}^+(\sigma)$, in \mathcal{F} , P plus $\sigma_3\sigma$ and $\sigma\sigma_2$ forms an oriented cycle C in \mathcal{F} , with $V(C) \subseteq V(C') \cup \{\sigma\}$.

Therefore, in both cases, there exists an oriented cycle C in \mathcal{F} such that $V(C) \subseteq V(C') \cup \{\sigma\}$. Since all the vertices have the same label in \mathcal{F} and \mathcal{F}' , the minimum number of disconnected flips in an oriented cycle of \mathcal{F} is therefore at most the number of disconnected flips in C if σ is connected, and the number of disconnected flips in C +1 if σ is disconnected. Thus, if σ is connected, $\gamma(\mathcal{F}') - \gamma(\mathcal{F}) \geq 0$, and if σ is disconnected, $\gamma(\mathcal{F}') - \gamma(\mathcal{F}) \geq -1$. Again, in both cases, property (*) ensures that $(|V(\mathcal{F}')| + \gamma(\mathcal{F}')) - (|V(\mathcal{F})| + \gamma(\mathcal{F})) \geq -1$. \square

Note that the lower bound given by Lemma 9 corresponds to the upper bound given by Christie [5] for paths when the graph of flips is connected. Indeed, Christie gives an algorithm to transform any path G into another one H by using $|V(\mathcal{F}(G, H))| + s$ good flips, where s is the number of connected components of $\mathcal{F}(G, H)$ that only have disconnected flips. Thus, if the graph of flips is connected, s is equal to 0 if there exists a connected flip in it, and 1 otherwise. As the graph is unoriented in this case, s is thus equal to $\gamma(\mathcal{F}(G, H))$.

In our case, the lower bound given by Lemma 9 is not necessarily tight when we only flip bad edges of the same cycle. Indeed, let us consider for example the graphs G'_k and H'_k obtained from G_k and H_k by adding a connected and a disconnected C_4 , that see each other, on the same branch of the original subdivided star (see Figure 3).

We claim that a proof similar to the one of Lemma 11 can be adapted to prove that the shortest transformation from G'_k to H'_k has length at least $\frac{3k}{2}$. On the other hand, the addition of the two cycles on a leaf of a branch created in $\mathcal{F}(G'_k, H'_k)$ an oriented cycle of length 2 with a connected and a disconnected vertex. Thus, Lemma 9 gives the lower bound $k + 2$.

We can now apply Lemma 9 to prove the following. Recall that G_k and H_k were defined at the beginning of the section.




Fig. 5: G'_4 and H'_4 . The black thick edges are in $E(G'_4 \cap H'_4)$, the blue thin edges are in $E(G'_4 - H'_4)$ and the red dashed edges are in $E(H'_4 - G'_4)$.

Lemma 10. *If we only flip pairs of edges that are in the same cycle of \mathcal{C} , then the shortest transformation from G_k to H_k has length at least $2k - 1$.*

Proof. Let us first show that G_k and H_k are close. We thus need to prove that $\Delta(G_k, H_k)$ has a unique partition into alternating cycles, containing only short cycles. In the proof all the indices have to be read modulo k . By construction, the vertex c is not incident to any edge of $\Delta(G_k, H_k)$ and all the other vertices are incident to one edge of $G_k - H_k$ and one edge of $H_k - G_k$. For every i , $v_{i,1}$ is incident to an edge of $G_k - H_k$, namely $v_{i,1}v_{i,2}$. The vertex $v_{i,2}$ is incident to an edge of $H_k - G_k$, namely $v_{i,2}v_{i+1,4}$, which in turn is incident to an edge of $G_k - H_k$, $v_{i+1,4}v_{i+1,3}$. And $v_{i+1,3}$ is incident to an edge of $H_k - G_k$, $v_{i+1,3}v_{i,1}$. This set of four edges create a short cycle, denoted by C_i . Since this property holds for every i , $\Delta(G_k, H_k)$ can be partitioned into k short cycles. The uniqueness of the edges of $G_k - H_k$ and $H_k - G_k$ incident to a vertex $v_{i,j}$ ensures the uniqueness of the decomposition into cycles.

Let us now prove that $\mathcal{F} := \mathcal{F}(G_k, H_k)$ only contains disconnected flips. Since there are $4k$ edges in $\Delta(G_k, H_k)$, there are therefore k vertices in \mathcal{F} . Let us prove that, for every i , the good flip $\sigma_i \in \mathcal{F}(G_k, H_k)$ for C_i is disconnected. Since $v_{i,2}v_{i+1,4}$ and $v_{i+1,3}v_{i,1}$ are the edges of $H_k - G_k$ in C_i , the good flip of \mathcal{F} is $\sigma_i : (v_{i,1}v_{i,2}, v_{i+1,4}v_{i+1,3}) \rightarrow (v_{i,1}v_{i+1,3}, v_{i,2}v_{i+1,4})$. By construction, $(v_{i,2}, v_{i,1}, v_{i+1,3}, v_{i+1,4})$ are aligned in G_k and then Lemma 1 ensures that σ_i is disconnected.

Let us finally show that \mathcal{F} is an oriented cycle of length k . For every i , the path from $v_{i,1}v_{i,2}$ to $v_{i+1,4}v_{i+1,3}$ in G_k is $P_i := (v_{i,1}, c, v_{i+1,1}, v_{i+1,2}, v_{i+1,3})$. The only edge of $G_k - H_k$ that belongs to P_i is $v_{i+1,1}v_{i+1,2}$, which is an edge of σ_{i+1} . So σ_i only sees σ_{i+1} .

Therefore, \mathcal{F} is an oriented cycle of length k containing only disconnected flips.

Since G_k, H_k are close, Lemma 9 ensures that a shortest transformation when we only flip pairs of edges that are in the same cycle of \mathcal{C} has length at least $|V(\mathcal{F})| + \gamma(\mathcal{F}) = 2k - 1$. Indeed $|V(\mathcal{F})| = k$ and $\gamma(\mathcal{F}) = k - 1$. \square

We claim that the lower bound given by Lemma 10 is also an upper bound. Indeed, by Lemma 5, every short cycle can be removed within two steps. We can apply this strategy for $k - 1$ cycles which needs $2k - 2$ steps. When there only remains one cycle in the symmetric difference, then it is easy to check that the good flip indeed keeps the connectivity of the graph. Thus the last short cycle can be removed in one step.

From Lemma 10, we can finally deduce the following corollary:

Corollary 2. *There exist some connected graphs G and H with the same degree sequence for which, if we only flip edges of the same cycle of any decomposition of $\Delta(G, H)$ into alternating cycles, the shortest connected transformation from G to H has length at least $2(\frac{\delta(G, H)}{2} - \text{mnc}(G, H)) - 1$.*

4.2 Flipping bad edges

The restriction of flipping only edges that are in the same cycle of the partition of $\Delta(G_k, H_k)$ into alternating cycles might seem strong.

That being said, we have also studied the transformation from G_k to H_k under a weaker assumption, which is the one of only flipping bad edges.

Lemma 11. *If at any time, we only flip pairs of bad edges, then the shortest transformation from G_k to H_k has length at least $\lceil \frac{3k}{2} \rceil - 1$.*

Proof. Let G_t be a graph obtained after applying t arbitrary flips, whose edges are in $\Delta(G_k, H_k)$, to G_k . The branch B_i of G_t is the unique path from $v_{i,1}$ to a leaf of G_t (that is therefore identified as the leaf of the branch) that does not contain the vertex c . Note that since we only flip bad edges, in G_t , one of the two edges incident to $v_{i,1}$ is $cv_{i,1}$. A core of a branch B_i is an edge $v_{j,2}v_{j,3}$ that belongs to the branch B_i . Note that a branch might contain no core. An edge of $G_t - H$ is *external* if it is incident to a leaf of G_t , and *internal* otherwise. An *inversion* of B_i is a flip whose edges are both on B_i . A *displacement* between two branches B_i and B_j is a flip between one edge of B_i and one edge of B_j . Note that all the flips are either inversions or displacements.

Let S be a sequence of flips that transforms G_k into H_k using only bad edges, and let us show that S has length at least $\lceil \frac{3k}{2} \rceil - 1$.

Let us first show that S contains at least $k - 2$ displacements if k is even, and $k - 1$ if k is odd. First note that the only way to change the leaf of the branch B_i consists in making a displacement between B_i and another branch B_j . Indeed, an inversion flips two edges of the same branch, and therefore does not change its content. On the other hand, a displacement between the branches B_i and B_j permutes the leaves of the two branches.

Since the leaf of the branch B_i is $v_{i,4}$ in G_k , and $v_{i+2,4}$ in H_k (the addition being modulo k), the leaves associated to the branches (B_1, \dots, B_k) must be changed from $(v_{1,4}, v_{2,4}, \dots, v_{k-2,4}, v_{k-1,4}, v_{k,4})$ to $(v_{3,4}, v_{4,4}, \dots, v_{k,4}, v_{1,4}, v_{2,4})$, only using displacements, i.e. transpositions of the leaves. The canonical notation of the permutation (i.e partition into cycles) from $(1, 2, \dots, k-2, k-1, k)$ to $(3, 4, \dots, k, 1, 2)$ is either $(1, 3, \dots, k-1)(2, 4, \dots, k)$ if k is even, or $(1, 3, \dots, k, 2, 4, \dots, k-1)$ if k is odd. Thus, it is partitioned into 2 orbits if k is even, and 1 orbit if k is odd, and therefore its decomposition into transpositions contains $k - 2$ transpositions if k is even, and $k - 1$ if k is odd. Therefore, in order to put the leaves $v_{i+2,4}$ on the branches B_i for any i , at least $k - 2$ transpositions are needed if k is even, and at least $k - 1$ transpositions are needed if k is odd. Therefore, if ℓ is the number of inversions in S , then S has length at least $\ell + k - 2$ if k is even, and $\ell + k - 1$ if k is odd.

Let us now prove that S has length at least $2k - \ell$. Assume that S contains ℓ inversions. First note that, in both G_k and H_k , the vertices $v_{i,2}v_{i,3}$ appear in the same branch, but in G_k , $(c, v_{i,2}, v_{i,3})$ are aligned and in H_k , $(c, v_{i,3}, v_{i,2})$ are aligned. Thus, the order of $v_{i,2}v_{i,3}$ in the branch has to change during the transformation. The only way to change the order of $v_{i,2}$ and $v_{i,3}$ in a branch is to make an inversion of a subpath containing the edge $v_{i,2}v_{i,3}$. Since S contains only ℓ inversions, it means that there exist at least $k - \ell$ indices j for which the edge $v_{j,1}v_{j,2}$ has to belong to a flip before its inversion. Moreover, after each inversion, at most one core belongs to its final branch. Thus, there exist at least $k - \ell$ indices j such that the bad edge incident to $v_{j,2}$ has to belong to a flip after the inversion of $v_{j,1}v_{j,2}$ in order to connect it with $v_{j-1,1}$. So $2k - 2\ell$ internal edges have to be flipped during displacements.

Let us now focus on the leaves. Since S contains ℓ inversions, at most ℓ indices j satisfy that $v_{j,3}$ is incident to a leaf just before the inversion of $v_{j,2}v_{j,3}$. Since all the edges $v_{i,2}v_{i,3}$ have to be inversed during the transformation and since in G_k , every vertex $v_{i,3}$ is incident to a leaf, at least $k - \ell$ external edges have to belong to a flip before the inversion of the core they are incident to in G_k . Similarly, at most ℓ indices j satisfy that $v_{j,2}$ is incident to a leaf just after the inversion of $v_{j,2}v_{j,3}$. Since in H_k , every vertex $v_{i,2}$ is incident to a leaf, at least $k - \ell$ external edges have to belong to a flip after the inversion of the core they are incident to in H_k . So at least $2k - 2\ell$ external edges have to be flipped during displacements.

Therefore, in total, we need to flip at least $4k - 4\ell$ edges during displacements. So at least $2k - 2\ell$ displacements are needed in addition to the ℓ inversions, and the total number of flips in S is at least $2k - \ell$.

Thus, if k is even, S has length at least $\max(\ell + k - 2, 2k - \ell)$ and if k is odd, S has length at least $\max(\ell + k - 1, 2k - \ell)$. In both cases, the two lower bounds meet for $\ell = \lfloor \frac{k}{2} \rfloor + 1$ at the value $\lceil \frac{3k}{2} \rceil - 1$. Therefore, S has length at least $\lceil \frac{3k}{2} \rceil - 1$. \square

From Lemma 11, we can deduce the following:

Corollary 3. *There exist some connected graphs G and H with the same degree sequence for which, if we only flip edges of $\Delta(G, H)$, the shortest connected transformation from G to H has length at least $\frac{3}{2}(\frac{\delta(G, H)}{2} - \text{mnc}(G, H)) - 1$.*

Conjecture 1. The shortest transformation from G_k to H_k has length $2k - 1$.

References

1. S. Bereg and H. Ito. Transforming graphs with the same graphic sequence. *Journal of Information Processing*, 25:627–633, 2017.
2. P. Berman, S. Hannenhalli, and M. Karpinski. 1.375-Approximation Algorithm for Sorting by Reversals. In R. Möhring and R. Raman, editors, *Algorithms — ESA 2002*, page 200–210, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
3. N. Bousquet and A. Mary. Reconfiguration of graphs with connectivity constraints. In *Approximation and Online Algorithms - WAOA 2018*, pages 295–309, 2018.
4. A. Caprara. Sorting by reversals is difficult. In *Proceedings of the First Annual International Conference on Computational Molecular Biology, RECOMB '97*, pages 75–83. ACM, 1997.
5. D. A. Christie. A 3/2-approximation algorithm for sorting by reversals. In *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '98*, pages 244–252, 1998.
6. R. Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, Heidelberg, third edition, 2005.
7. S. L. Hakimi. On realizability of a set of integers as degrees of the vertices of a linear graph. i. *Journal of the Society for Industrial and Applied Mathematics*, 10(3):496–506, 1962.
8. S. L. Hakimi. On realizability of a set of integers as degrees of the vertices of a linear graph ii. uniqueness. *Journal of the Society for Industrial and Applied Mathematics*, 11(1):135–147, 1963.
9. S. Hannenhalli and P. A. Pevzner. Transforming cabbage into turnip: polynomial algorithm for sorting signed permutations by reversals. *Journal of the ACM (JACM)*, 46(1):1–27, 1999.
10. J. D. Kececioğlu and D. Sankoff. Exact and approximation algorithms for sorting by reversals, with application to genome rearrangement. *Algorithmica*, 13(1/2):180–210, 1995.
11. N. Nishimura. Introduction to reconfiguration. *preprint*, 2017.
12. J. Senior. Partitions and their representative graphs. *American Journal of Mathematics*, 73(3):663–689, 1951.
13. R. Taylor. Constrained switchings in graphs. In K. L. McAvaney, editor, *Combinatorial Mathematics VIII*, pages 314–336, Berlin, Heidelberg, 1981. Springer Berlin Heidelberg.
14. T. G. Will. Switching distance between graphs with the same degrees. *SIAM Journal on Discrete Mathematics*, 12(3):298–306, 1999.