KERNELIZATION OF WHITNEY SWITCHES[*]

FEDOR V. FOMIN[†] AND PETR A. GOLOVACH[†]

Abstract. A fundamental theorem of Whitney from 1933 asserts that 2-connected graphs G and H are 2-isomorphic, or equivalently, their cycle matroids are isomorphic if and only if G can be transformed into H by a series of operations called Whitney switches. In this paper we consider the quantitative question arising from Whitney's theorem: Given two 2-isomorphic graphs, can we transform one into another by applying at most k Whitney switches? This problem is already NP-complete for cycles, and we investigate its parameterized complexity. We show that the problem admits a kernel of size $\mathcal{O}(k)$ and thus is fixed-parameter tractable when parameterized by k.

Key words. Whitney switch, 2-isomorphism, parameterized complexity, kernelization

AMS subject classifications. 05C85, 68R10

DOI. 10.1137/20M1367519

1. Introduction. A fundamental result of Whitney from 1933 [36] asserts that every 2-connected graph is completely characterized, up to a series of Whitney switches (also known as 2-switches), by its edge set and cycles. This theorem is one of the cornerstones of matroid theory, since it provides an exact characterization of two graphs having isomorphic cycle matroids [33]. In graph drawing and graph embeddings, this theorem (applied to dual graphs) is used to characterize all drawings of a planar graph on the plane [8].

A *Whitney switch* is an operation that from a 2-connected graph G constructs graph G' as follows. Let $\{u, v\}$ be two vertices of G whose removal separates G into two disjoint subgraphs G_1 and G_2. The graph G' is obtained by flipping the neighbors of u and v in the set of vertices of G_2. In other words, for every vertex $w \in V(G_2)$, if w was adjacent to u in G, in graph G' edge uw is replaced by vw. Similarly, if w was adjacent to v in G, we replace vw by uw. See Figure 1 for an example.

If we view the graph G as a graph with labeled edges, then a Whitney switch transforms G into a graph G' with the same set of labeled edges; however, graphs G and G' are not necessarily isomorphic. On other hand, graphs G and G' have the same set of cycles in the following sense: a set of (labeled) edges forms a cycle in G if and only if it forms a cycle in G'. (In other words, the cycle matroids of G and G' are isomorphic.) Whitney's theorem says that the opposite is also true: if there is a cycle-preserving mapping between graphs G and G', then one graph can be transformed into another by a sequence of Whitney switches. To state the theorem of Whitney more precisely, we need to define 2-isomorphisms.

We say that 2-connected graphs G and H are *2-isomorphic* if there is a bijection $\varphi \colon E(G) \to E(H)$ such that φ and φ^{-1} preserve cycles, that is, for every cycle C of G, C is mapped to a cycle of H by φ and, symmetrically, every cycle of H is mapped to a cycle of G by φ^{-1}. We refer to φ as a *2-isomorphism* from G to H.

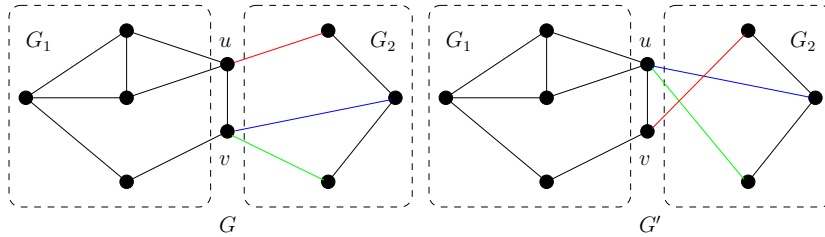[†]Department of Informatics, University of Bergen, Bergen, 5020 Norway (Fedor.Fomin@uib.no, Petr.Golovach@uib.no).

FIG. 1. *G′ is obtained from G by the Whitney switch with respect to the partition of $G - \{u, v\}$ into G_1 and G_2.*
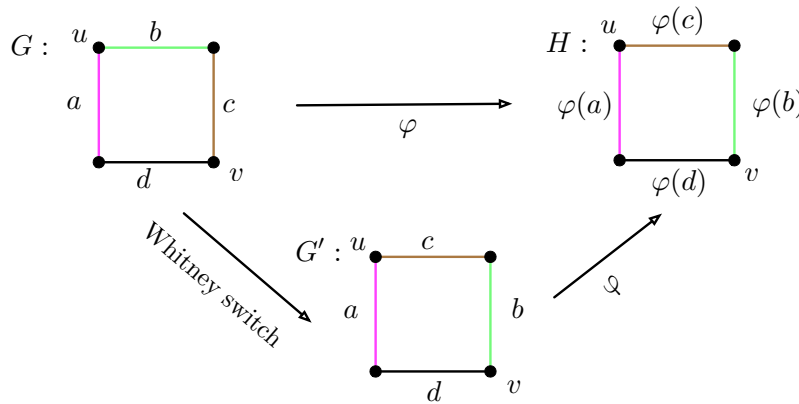

FIG. 2. *Graph G is not φ-isomorphic to H but its Whitney switch G′ is.*

An isomorphism $\psi \colon V(G) \to V(H)$ is a *φ-isomorphism* if for every edge $uv \in E(G)$, $\varphi(uv) = \psi(u)\psi(v)$, and G and H are *φ-isomorphic* if there is an isomorphism G to H that is a φ-isomorphism. Let us note that for 3-connected graphs the notions of 2-isomorphism and φ-isomorphism coincide. More precisely, if G is 3-connected and 2-isomorphic to H under φ, then G and H are φ-isomorphic [30, Lemma 1]. But for 2-connected graphs this is not true. For example, the graphs in Figure 1 are not isomorphic but are 2-isomorphic. But even isomorphic graphs with 2-isomorphism φ do not always have a φ-isomorphism. For example, for the 2-isomorphism φ (Figure 2) mapping a cycle G into another cycle H (we view these cycles as labeled graphs), there is no φ-isomorphism. (For every φ-isomorphism edges $\varphi(a)$ and $\varphi(b)$ should have an endpoint in common.) On the other hand, graph G' obtained from G by Whitney switch (for vertices u and v) is φ-isomorphic to H.

THEOREM 1 (Whitney's theorem [36]). *If there is a 2-isomorphism φ from graph G to graph H, then G can be transformed by a sequence of Whitney switches to a graph G' which is φ-isomorphic to H.*

However, Whitney's theorem does not provide an answer to the following computational question: Given a 2-isomorphism φ from graph G to graph H, *what is the minimum number of Whitney switches required to transform G to a graph φ-isomorphic to H?* Truemper in [30] proved that $n - 2$ switches always suffices, where n is the number of vertices in G. He also proved that this upper bound it tight, that is, there are graphs G and H for which $n - 2$ switches are necessary. In this paper we study the algorithmic complexity of the following problem about Whitney switches.

WHITNEY SWITCHES

Input:	2-Isomorphic n-vertex graphs G and H with a 2-isomorphism $\varphi\colon E(G) \to E(H)$ and a nonnegative integer k.
Task:	Decide whether it is possible to obtain from G a graph G' that is φ-isomorphic to H by at most k Whitney switches.

The departure point for our study of WHITNEY SWITCHES is an easy reduction (Theorem 3) from SORTING BY REVERSALS that establishes NP-completeness of WHITNEY SWITCHES even when input graphs G and H are cycles. Our main algorithmic result is the following theorem (we postpone the definition of a kernel till section 2). Informally, it means that the instance of the problem can be compressed in polynomial time to an equivalent instance with two graphs on $\mathcal{O}(k)$ vertices. It also implies that WHITNEY SWITCHES is fixed-parameter tractable parameterized by k.

THEOREM 2. WHITNEY SWITCHES *admits a kernel with* $\mathcal{O}(k)$ *vertices and is solvable in* $2^{\mathcal{O}(k \log k)} \cdot n^{\mathcal{O}(1)}$ *time.*

While Theorem 2 is not restricted to planar graphs, pipelined with the well-known connection of planar embeddings and Whitney switches, it can be used to obtain interesting algorithmic consequences about the distance between planar embeddings of a graph. Recall that graphs G and G^* are called abstractly dual if there is a bijection $\pi\colon E(G) \to E(G^*)$ such that edge set $E \subseteq E(G)$ forms a cycle in G if and only if $\pi(E)$ is a minimal edge-cut in G^*. By another classical theorem of Whitney [35], a graph G has a dual graph if and only if G is planar. Moreover, an embedding of a planar graph into a sphere is uniquely defined by the planar graph G and edges of the faces, or equivalently, its dual graph G^*. While every 3-connected planar graph has a unique embedding into the sphere, a 2-connected graph can have several nonequivalent embeddings, and hence several nonisomorphic dual graphs. If G_1^* and G_2^* are dual graphs of graph G, then G_1^* is 2-isomorphic to G_2^*. Then by Theorem 1, by a sequence of Whitney switches G_1^* can be transformed into G_2^*, or equivalently, the embedding of G corresponding to G_1^* can be transformed into an embedding of G corresponding to G_2^*. We refer to the survey of Thomassen [29, section 2.2] for more details. By Theorem 2, we have that given two planar embeddings of a (labeled) 2-connected graph G, deciding whether one embedding can be transformed into another by making use of at most k Whitney switches admits a kernel of size $\mathcal{O}(k)$ and is fixed-parameter tractable.

Related work. Whitney's theorem had a strong impact on the development of modern graph and matroid theories. While the original proof is long, a number of simpler proofs are known in the literature. The most relevant to our work is the proof of Truemper [30], whose proof is on the application of Tutte decomposition [31, 32].

WHITNEY SWITCHES can be seen as an example of reconfiguration problems. The study of reconfiguration problems becomes a popular trend in parameterized complexity (see, e.g., [22, 25]).

The well-studied problem, which is similar in spirit to WHITNEY SWITCHES, is the problem of computing the flip distance for triangulations of a set of points. (As we mentioned above, WHITNEY SWITCHES for planar graphs is equivalent to the problem of computing the Whitney switch distance between planar embeddings.) The parameterized complexity of this problem was studied in [10, 24]. We refer to the survey of Bose and Hurtado [4] for the discussion of the relations between

geometric and graph variants. The problem is known to be NP-complete [23] and FPT parameterized by the number of flips [19]. For the special case when the set of points defines a convex polygon, the problem of computing the flip distance between triangulations is equivalent to computing the rotation distance between two binary trees. For that case linear kernels are known [10, 24], but for the general case the existence of a polynomial kernel is open.

Overview of the proof of Theorem 2. The main tool in the construction of the kernel is the classical Tutte decompositions [31, 32]. We postpone the formal definition till section 2; informally, the Tutte decomposition of a 2-connected graph represents the vertex separators of size two in a tree-like structure. Each node of this tree represents a part of the graph (or *bag*) that is either a 3-connected graph or a cycle, and each edge corresponds to a separator of size two. Then a 2-isomorphism of G and H allows us to establish an isomorphism of the trees representing the Tutte decompositions of the input graphs. After that, potential Whitney switches can be divided into two types: the switches with respect to separators corresponding to the edges of the trees and the switches with respect to separators formed by nonadjacent vertices of a cycle-bag. The switches of the first type are relatively easy to analyze and we can identify necessary switches of this type. The "troublemakers" that make the problem hard are switches of the second type. To deal with them, we use the structural results about sorting of permutations by reversals of Hannenhalli and Pevzner [16] adapted for our purposes. This allows us to identify a set of vertices of size $\mathcal{O}(k)$ that potentially can be used for Whitney switches transforming G to H. Given such a set of crucial vertices, we simplify the structure of the input graphs and then reduce their size.

Organization of the paper. In section 2, we give basic definitions. In section 3, we discuss the SORTING BY REVERSALS problem for permutations that is closely related to WHITNEY SWITCHES. Section 4 contains structural results used by our kernelization algorithm, and in section 5, we give the algorithm itself. We conclude in section 6 by discussing further directions of research.

2. Preliminaries.

Graphs. All graphs considered in this paper are finite undirected graphs without loops or multiple edges, unless it is specified explicitly that we consider directed graphs (in section 6 we deal with tournaments). We follow the standard graph theoretic notation and terminology (see, e.g., [13]). For each of the graph problems considered in this paper, we let $n = |V(G)|$ and $m = |E(G)|$ denote the number of vertices and edges, respectively, of the input graph G if it does not create confusion. For a graph G and a subset $X \subseteq V(G)$ of vertices, we write $G[X]$ to denote the subgraph of G induced by X. For a set of vertices S, $G - S$ denotes the graph obtained by deleting the vertices of S, that is, $G - S = G[V(G) \setminus S]$; for a vertex v, we write $G - v$ instead of $G - \{v\}$. Similarly, for a set of edges A (an edge e, respectively), $G - A$ ($G - e$, respectively) denotes the graph obtained by the deletion of the edges of A (an edge e, respectively). For a set of edges A, we use $V(A)$ to denote the set of end-vertices of the edges of A; for an edge e, we write $V(e)$ instead of $V(\{e\})$. For a vertex v, we denote by $N_G(v)$ the *(open) neighborhood* of v, i.e., the set of vertices that are adjacent to v in G and we use $E_G(v)$ to denote the set of edges incident to v. We use $N_G[v]$ to denote the *closed neighborhood*, that is, $N_G(v) \cup \{v\}$. For $S \subseteq V(G)$, $N_G[S] = \bigcup_{v \in S} N_G[v]$ and $N_G(S) = N_G[S] \setminus S$. We write $N_G^2(v) = N_G(N_G[v])$ for a vertex v to denote the *second neighborhood*. A vertex v is *simplicial* if $N_G(v)$ is a *clique*, that is, a set of pairwise adjacent vertices. A pair (A, B), where $A, B \subseteq V(G)$,

is a *separation* of G if $A \cup B = V(G)$, $A \setminus B \neq \emptyset$, $B \setminus A \neq \emptyset$, and G has no edge uv with $u \in A \setminus B$ and $v \in B \setminus A$; $|A \cap B|$ is the *order* of the separation. If the order is 2, then we say that (A, B) is a *Whitney separation*. A set $S \subseteq V(G)$ is a *separator* of G if there is a separation (A, B) of G with $S = A \cap B$. For a positive integer k, a graph G is *k-connected* if G is a connected graph with at least $k + 1$ vertices without a separator of size at most $k - 1$. In particular, G is 2-connected if $G - v$ is connected for every $v \in V(G)$.

Isomorphisms. Graphs G and H are *isomorphic* if there is bijection $\eta \colon V(G) \to V(H)$, called *isomorphism*, preserving edges, that is, $uv \in E(G)$ if and only if $\eta(u)\eta(v) \in E(H)$. We say that 2-connected graphs G and H are *2-isomorphic* if there is a bijection $\varphi \colon E(G) \to E(H)$ such that φ and φ^{-1} preserve cycles, that is, for every cycle C of G, C is mapped to a cycle of H by φ and, symmetrically, every cycle of H is mapped to a cycle of G by φ^{-1}. We refer to φ as a *2-isomorphism* from G to H. An isomorphism $\psi \colon V(G) \to V(H)$ is a *φ-isomorphism* if for every edge $uv \in E(G)$, $\varphi(uv) = \psi(u)\psi(v)$, and G and H are φ-isomorphic if there is an isomorphism G to H that is a φ-isomorphism.

Whitney switches. Let G be a 2-connected graph. Let also (A, B) be a Whitney separation of G with $A \cap B = \{u, v\}$. The Whitney switch operation with respect to (A, B) transforms G as follows: for every vertex $w \in B \setminus A$ that is adjacent to u, replace wu by wv, and symmetrically, for every vertex $w \in B$ that is adjacent to v, replace wv by wu. Equivalently, we take $G[A]$ and $G[B]$ and identify the vertex u of $G[A]$ with the vertex v of $G[B]$ and, symmetrically, v of $G[A]$ with u of $G[B]$; if u and v are adjacent in G, then the edges uv of $G[A]$ and $G[B]$ are identified as well. Let G' be the obtained graph. We define the mapping $\sigma_{(A,B)} \colon E(G) \to E(G')$ that maps the edges of $G[A]$ and $G[B]$, respectively, to themselves. It is easy to see that $\sigma_{(A,B)}$ is a 2-isomorphism of G to G'. Therefore, if φ is a 2-isomorphism of G to a graph H, then $\varphi \circ \sigma_{(A,B)}^{-1}$ is a 2-isomorphism of G' to H. To simplify notation, we assume, if it does not create confusion, that the sets of edges of G and G' are identical and we only change incidences by switching. In particular, under this assumption, we have that $\varphi \circ \sigma_{(A,B)}^{-1} = \varphi$. Note that the graphs G and G' have the same sets of vertices.

By definition, the graphs obtained from G by the Whitney switches with respect to (A, B) and (B, A) are isomorphic. However, these two switches are not symmetric with respect to the vertex notation. Moreover, the end-vertices of an edge may change by a switch. To deal with this issue, sometimes it is convenient for us to consider edge Whitney separations instead of vertex separations. We say that a pair (L, R), where $L, R \subseteq E(G)$ and $L \cup R = E(G)$, is an *edge Whitney separation* if $(V(L), V(R))$ is a Whitney separation. Then the Whitney switch with respect to (L, R) is defined as the Whitney switch with respect to $(V(L), V(R))$. It is straightforward to see that if (A, B) is a Whitney separation, then the Whitney switch with respect to (A, B) is equivalent to the Whitney switch with respect to the edge Whitney separation (L, R), where $L = E(G[A])$ and $R = E(G[B])$.

Tutte decomposition. Our kernelization algorithm for WHITNEY SWITCHES is based on the classical result of Tutte [31, 32] about decomposing 2-connected graphs via separators of size two. Following Courcelle [11], we define Tutte decompositions in terms of tree decompositions.

A *tree decomposition* of a graph G is a pair $\mathcal{T} = (T, \{X_t\}_{t \in V(T)})$, where T is a tree whose every node t is assigned a vertex subset $X_t \subseteq V(G)$, called a *bag*, such that the following three conditions hold:

(T1) $\bigcup_{t \in V(T)} X_t = V(G)$, that is, every vertex of G is in at least one bag,

(T2) for every $uv \in E(G)$, there exists a node t of T such that the bag X_t contains both u and v,

(T3) for every $v \in V(G)$, the set $T_v = \{t \in V(T) \mid v \in X_t\}$, i.e., the set of nodes whose corresponding bags contain v, induces a connected subtree of T.

To distinguish between the vertices of the decomposition tree T and the vertices of the graph G, we will refer to the vertices of T as *nodes*.

Let $\mathcal{T} = (T, \{X_t\}_{t \in V(T)})$ be a tree decomposition of G. The *torso* of X_t for $t \in V(T)$ is the graph obtained from $G[X_t]$ by additionally making adjacent every two distinct vertices $u, v \in X_t$ such that there is $t' \in V(T)$ adjacent to t with $u, v \in X_t \cap X_{t'}$. For adjacent $t, t' \in V(T)$, $X_t \cap X_{t'}$ is the *adhesion set* of the bags X_t and $X_{t'}$ and $|X_t \cap X_{t'}|$ is the *adhesion* of the bags. The maximum adhesion of adjacent bags is called the *adhesion* of the tree decomposition.

Let G be a 2-connected graph. A tree decomposition $\mathcal{T} = (T, \{X_t\}_{t \in V(T)})$ is said to be a *Tutte decomposition* if $\mathcal{T}$ is a tree decomposition of adhesion 2 such that there is a partition $(W_2, W_{\geq 3})$ of $V(T)$ such that the following holds:

(T4) $|X_t| = 2$ for $t \in W_2$ and $|X_t| \geq 3$ for $t \in W_{\geq 3}$,

(T5) the torso of X_t is either a 3-connected graph or a cycle for every $t \in W_{\geq 3}$,

(T6) for every $t \in W_2$, $d_T(t) \geq 2$ and $t' \in W_{\geq 3}$ for each neighbor t' of t,

(T7) for every $t \in W_{\geq 3}$, $t' \in W_2$ for each neighbor t' of t,

(T8) if $t \in W_2$ and $d_T(t) = 2$, then for the neighbors t' and t'' of t, either the torso of t' or the torso of t'' is a 3-connected graph or the vertices of X_t are adjacent in G.

Notice that the bags X_t for $t \in W_2$ are distinct separators of G of size two, and $X_t \subseteq X_{t'}$ for $t \in W_2$ and $t' \in N_T(t)$. Observe also that if $\{u, v\}$ is a separator of G of size two, then either $\{u, v\} = X_t$ for some $t \in W_2$ or $u, v \in X_t$ for $t \in W_{\geq 3}$ such that the torso of X_t is a cycle and u and v are nonadjacent vertices of the torso.

Combining the results of Tutte [31, 32] and of Hopcroft and Tarjan [18], we state the following proposition.

PROPOSITION 1 ([31, 32, 18]). *A 2-connected graph G has a unique Tutte decomposition that can be constructed in linear time.*

Parameterized complexity and kernelization. We refer to the books [12, 14, 15] for a detailed introduction to the field. Here we only give the most basic definitions. In parameterized complexity theory, the computational complexity is measured as a function of the input size n of a problem and an integer *parameter* k associated with the input. A parameterized problem is said to be *fixed parameter tractable* (or FPT) if it can be solved in time $f(k) \cdot n^{\mathcal{O}(1)}$ for some function f. A *kernelization* algorithm for a parameterized problem Π is a polynomial algorithm that maps each instance (I, k) of Π to an instance (I', k') of Π such that

(i) (I, k) is a yes-instance of Π if and only if (I', k') is a yes-instance of Π and

(ii) $|I'| + k'$ is bounded by $f(k)$ for a computable function f.

Respectively, (I', k') is a *kernel* and f is its *size*. A kernel is *polynomial* if f is polynomial. It is common to present a kernelization algorithm as a series of *reduction rules*. A reduction rule for a parameterized problem is an algorithm that takes an instance of the problem and computes in polynomial time another instance that is "simpler" in a certain way. A reduction rule is *safe* if the computed instance is equivalent to the input instance.

3. Sorting by reversals. Sorting by reversals is the classical problem with many applications including bioinformatics. We refer to the book of Pevzner [26] for a detailed survey of results and applications of this problem. This problem is also strongly related to WHITNEY SWITCHES—solving the problem for two cycles is basically the same as sorting circular permutations by reversals. First we use this relation to observe the NP-completeness. But we also need to establish some structural properties of sorting by reversals which will be used in a kernelization algorithm.

Let $\pi = (\pi_1, \ldots, \pi_n)$ be a permutation of $\{1, \ldots, n\}$, that is, a bijective mapping of $\{1, \ldots, n\}$ to itself. Throughout this section, all considered permutations are permutations of $\{1, \ldots, n\}$. For two permutations π and π', $\pi' \circ \pi$ denotes the permutation such that $(\pi' \circ \pi)(i) = \pi'(\pi(i))$ for every $i \in \{1, \ldots, n\}$. For $1 \leq i \leq j \leq n$, the *reversal* $\rho(i, j)$ reverses the order of elements $\pi_i, \ldots, \pi_j$ and transforms π into

$$\rho(i, j) \circ \pi = (\pi_1, \ldots, \pi_{i-1}, \pi_j, \pi_{j-1}, \ldots, \pi_i, \pi_{j+1}, \ldots, \pi_n).$$

The *reversal distance* $d(\pi, \sigma)$ between two permutations π and σ is the minimum number of reversals needed to transform π to σ. For a permutation π, $d(\pi) = d(\pi, \iota)$, where ι is the identity permutation; note that $d(\pi, \sigma) = d(\sigma^{-1} \circ \pi, \iota)$ and this means that computing the reversal distance can be reduced to sorting a permutation by the minimum number of reversals.

These definitions can be extended for circular permutations (further, we may refer to usual permutations as *linear* to avoid confusion). We say that $\pi^c = (\pi_1, \ldots, \pi_n)$ is a *circular* permutation if π^c is the class of the permutations that can be obtained from the linear permutation $(\pi_1, \ldots, \pi_n)$ by *rotations* and *reflections*, that is, π^c consists of the permutations

$$(\pi_1, \ldots, \pi_n), (\pi_n, \pi_1, \ldots, \pi_{n-1}), \ldots, (\pi_2, \ldots, \pi_n, \pi_1)$$

and

$$(\pi_n, \ldots, \pi_1), (\pi_1, \pi_n, \ldots, \pi_2), \ldots, (\pi_{n-1}, \ldots, \pi_1, \pi_n).$$

To simplify notation, we use one representative from this class of permutations π^c to denote it and do not distinguish distinct representatives when discussing circular permutations. For $i, j \in \{1, \ldots, n\}$, the *circular reversal* $\rho^c(i, j)$ for $\pi^c = (\pi_1, \ldots, \pi_n)$ is defined in the same way as $\rho(i, j)$ if $i \leq j$ and for $i > j$, $\rho^c(i, j)$ transforms π^c into

$$\rho^c(i, j) \circ \pi^c = (\pi_n, \pi_{n-1}, \ldots, \pi_i, \pi_{j+1}, \ldots, \pi_{i-1}, \pi_j, \pi_{j-1} \ldots, \pi_1).$$

The *circular reversal distances* $d^c(\pi^c, \sigma^c)$ and $d^c(\pi^c)$ are defined in the same way as for linear permutations.

To see the connection between Whitney switches and circular reversals of permutations, consider a cycle G with the vertices $v_1, \ldots, v_n$ for $n \geq 4$ taken in the cycle order and the edges $e_i = v_{i-1}v_i$ for $i \in \{1, \ldots, r\}$ assuming that $v_0 = v_n$. Let $1 \leq i < j \leq n$ be such that v_i and v_j are not adjacent. Then the Whitney switch with respect to (A, B), where $A = \{v_1, \ldots, v_i\} \cup \{v_j, \ldots, v_n\}$ and $B = \{v_i, \ldots, v_j\}$, is equivalent to applying the reversal $\rho^c(i+1, j)$ to the circular permutation $(e_1, \ldots, e_n)$ of the edges of G. Moreover, let H be a cycle with n vertices and denote by $e'_1, \ldots, e'_n$ its edges in the cycle order. Notice that every bijection $\varphi \colon E(G) \to E(H)$ is a 2-isomorphism of G to H, and G and H are φ-isomorphic if and only if the circular permutation $\pi^c = (\varphi^{-1}(e'_1), \ldots, \varphi^{-1}(e'_n))$ is the same as $\sigma^c = (e_1, \ldots, e_n)$. We can assume that π^c is a permutation of $\{1, \ldots, n\}$ and σ^c is the identity permutation.
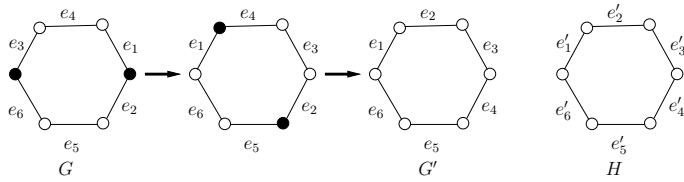
FIG. 3. *The construction of G' that is φ-isomorphic to H by the Whitney switches corresponding to the sorting by reversals $(3,4,1,2,5,6) \to (1,4,3,2,5,6) \to (1,2,3,4,5,6)$; $\varphi(e_i) = e_i'$ for $i \in \{1,\ldots,6\}$, and the vertices of the separators for the switches are shown in black.*

Then G can be transformed to a graph G' φ-isomorphic to H by at most k Whitney switches if and only if $d^c(\pi^c) \leq k$. An example is shown in Figure 3.

In particular, the above observation implies the hardness of WHITNEY SWITCHES, because the computing of the reversal distances is known to be NP-hard. For linear permutations, it was shown by Caprara in [7]. The following result for circular permutations was obtained by Solomon, Sutcliffe, and Lister [28].

PROPOSITION 2 ([28]). *It is NP-complete to decide, given a circular permutation π^c and a nonnegative integer k, whether $d^c(\pi^c) \leq k$.*

This brings us to the following theorem.

THEOREM 3. WHITNEY SWITCHES *is NP-complete even when restricted to cycles.*

For our kernelization algorithm, we need some further structural results about reversals in an optimal sorting sequence.

Let $\pi = (\pi_1, \ldots, \pi_n)$ be a linear permutation. For $1 \leq i \leq j \leq n$, we say that $(\pi_i, \ldots, \pi_j)$ is an *interval* of π. An interval $(\pi_i, \ldots, \pi_j)$ is called a *block* if either $i = j$ or $i < j$ and for every $h \in \{i+1, \ldots, j\}$, $|\pi_{h-1} - \pi_h| = 1$, that is, a block is formed by consecutive integers in π in either the ascending or descending order. An inclusion maximal block is called a *strip*. In other words, a strip is an inclusion maximal interval that has no *breakpoint*, that is, a pair of elements π_{h-1}, π_h with $|\pi_{h-1} - \pi_h| \geq 2$. It is said that a reversal $\rho(p,q)$ *cuts* a strip $(\pi_i, \ldots, \pi_j)$ if either $i < p \leq j$ or $i \leq q < j$, that is, the reversals separate elements that are consecutive in the identity permutation.

It is known that there are cases when every optimal sorting by reversal requires a reversal that cuts a strip. For example, as was pointed out by Hannenhalli and Pevzner in [16], the permutation $(3,4,1,2)$ requires three reversals that do not cut strips, but the sorting can be done by two reversals:[1]

$$(3,4,1,2) \to (1,4,3,2) \to (1,2,3,4).$$

Nevertheless, it was conjectured by Kececioglu and Sankoff [21] that there is an optimal sorting that does not cut strips other than at their first or last elements. This conjecture was proved by Hannenhalli and Pevzner in [16]. More precisely, they proved that there is an optimal sorting that does not cut strips of length at least three.

It is common for bioinformatics applications to consider *signed permutations* (see, e.g., [26]). In a signed permutation $\overrightarrow{\pi} = (\pi_1, \ldots, \pi_n)$, each element π_i has its *sign* "−" or "+." Then for $i, j \in \{1, \ldots, n\}$, the reversal reverses the sign of each element $\pi_i, \ldots, \pi_j$ besides reversing their order. We generalize this notion and define *partially signed linear permutations*, where each element has one of the signs: "−," "+," or "no

---

[1]This example can be extended for circular permutations: $(3,4,1,2,5,6) \to (1,4,3,2,5,6) \to (1,2,3,4,5,6)$.

sign." Formally, a partially signed linear permutation is $\overrightarrow{\pi} = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle)$ with $s_i \in \{-1, +1, 0\}$ for $i \in \{1, \ldots, n\}$. Then for $1 \leq i \leq j \leq n$, the reversal

$$\overrightarrow{\rho}(i, j) \circ \overrightarrow{\pi} = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_{i-1}, s_{i-1} \rangle,$$
$$\langle \pi_j, -s_j \rangle, \ldots, \langle \pi_i, -s_i \rangle, \langle \pi_{j+1}, s_{j+1} \rangle, \ldots, \langle \pi_n, s_n \rangle).$$

We say that $\overrightarrow{\pi} = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle)$ is *signed* if $s_i = -1$ or $s_i + 1$ for each $i \in \{1, \ldots, n\}$, and $\overrightarrow{\pi}$ is *unsigned* if $s_i = 0$ for every $i \in \{1, \ldots, n\}$. We define the *signed linear identity permutation* as $\overrightarrow{\iota} = (\langle 1, +1 \rangle, \ldots, \langle n, +1 \rangle)$.

We say that a partially signed linear permutation $\overrightarrow{\pi} = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle)$ *agrees in signs* with a signed linear permutation $\overrightarrow{\pi}' = (\langle \pi_1, s_1' \rangle, \ldots, \langle \pi_n, s_n' \rangle)$ if $s_i = s_i'$ for $i \in \{1, \ldots, n\}$ such that $s_i \neq 0$, that is, the zero signs are replaced by either -1 or $+1$ in the signed permutation. For a partially signed linear permutation $\overrightarrow{\pi}$, $\Sigma(\overrightarrow{\pi})$ denotes the set of all signed linear permutations $\overrightarrow{\pi}'$ that agree in signs with $\overrightarrow{\pi}$. The *reversal distance* $\overrightarrow{d}(\overrightarrow{\pi}, \overrightarrow{\sigma})$ between a partially signed linear permutation $\overrightarrow{\pi}$ and a signed linear permutation $\overrightarrow{\sigma}$ is the minimum number or reversal needed to obtain from $\overrightarrow{\pi}$ a partially signed linear permutation $\overrightarrow{\pi}'$ that agrees in signs with $\overrightarrow{\sigma}$; $\overrightarrow{d}(\overrightarrow{\pi}) = \overrightarrow{d}(\overrightarrow{\pi}, \overrightarrow{\iota})$. We say that a sequence of reversals of minimum length that result in a partially signed linear permutation that agrees in signs with $\overrightarrow{\iota}$ is an *optimal sorting* sequence. It is straightforward to observe the following.

Observation 1. For every partially signed linear permutation

$$\overrightarrow{d}(\overrightarrow{\pi}) = \min\{\overrightarrow{d}(\overrightarrow{\pi}') \mid \overrightarrow{\pi}' \in \Sigma(\overrightarrow{\pi})\}.$$

We generalize the results of Hannenhalli and Pevzner in [16] for partially signed linear permutations. Let $\overrightarrow{\pi} = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle)$. For $1 \leq i \leq j \leq n$, $(\langle \pi_i, s_i \rangle, \ldots, \langle \pi_j, s_j \rangle)$ is an *interval* of $\overrightarrow{\pi}$. An interval $(\langle \pi_i, s_i \rangle, \ldots, \langle \pi_j, s_j \rangle)$ is a *signed block* if $i = j$ or $i < j$ and the following holds:

  (i) for every $h \in \{i+1, \ldots, j\}$, $|\pi_{h-1} - \pi_h| = 1$,
  (ii) the block is *canonically signed*, that is, $s_h \in \{0, +1\}$ if $\pi_i < \cdots < \pi_j$ and
      $s_h \in \{0, -1\}$ if $\pi_i > \cdots > \pi_j$.

Similarly to unsigned permutations, an inclusion maximal signed block is called a *signed strip*. A reversal $\overrightarrow{\rho}(p, q)$ *cuts* a signed strip $(\langle \pi_i, s_i \rangle, \ldots, \langle \pi_j, s_j \rangle)$ if either $i < p \leq j$ or $i \leq q < j$.

Let $\overrightarrow{\pi} = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle)$ and $\overrightarrow{\pi}' = (\langle \pi_1, s_1' \rangle, \ldots, \langle \pi_n, s_n' \rangle)$ be signed linear permutations that may differ only in signs and let $\sigma = (\langle \pi_i, s_i \rangle, \ldots, \langle \pi_j, s_j \rangle)$ be a signed strip of $\overrightarrow{\pi}$. It is said that $\overrightarrow{\pi}$ and $\overrightarrow{\pi}'$ are *twins* with respect to σ if $s_h = s_h'$ for all $h \in \{1, \ldots, i-1\} \cup \{j+1, \ldots, n\}$, that is, the signs may be only different for elements of σ. The crucial nontrivial claim of Hannenhalli and Pevzner that was used to show that sorting of unsigned permutations can be done without cutting strips of length at least three is [16, Lemma 3.2].

LEMMA 1 ([16]). *Let $\overrightarrow{\pi}$ and $\overrightarrow{\pi}'$ be signed linear permutations that are twins with respect to a signed strip σ of $\overrightarrow{\pi}$ with $|\sigma| \geq 3$. Then $\overrightarrow{d}(\overrightarrow{\pi}) \leq \overrightarrow{d}(\overrightarrow{\pi}')$.*

Further, Hannenhalli and Pevzner used the result of Kececioglu and Sankoff [20] that for signed permutations, it is always possible to avoid cutting strips.

PROPOSITION 3 ([20]). *For a signed linear permutation $\overrightarrow{\pi}$, there is an optimal sorting sequence such that no reversal cuts a signed strip.*

Then the result of Hannenhalli and Pevzner [16] is obtained by combining Observation 1, Lemma 1, and Proposition 3. We use the same arguments for partially

signed linear permutations and the proof of the following lemma essentially repeats the proof of [16, Theorem 3.1] and we give it here for completeness.

LEMMA 2. *For a partially signed linear permutation $\overrightarrow{\pi}$, there is an optimal sorting sequence such that no reversal cuts a signed strip of length at least three.*

Proof. Let $\overrightarrow{\pi}$ be a partially signed linear permutation. The lemma is proved by the induction on $d = \overrightarrow{d}(\overrightarrow{\pi})$. The claim is straightforward if $d \leq 1$. Assume that $d \geq 2$ and the claim holds for the lesser values. By Observation 1, there is a signed permutation $\pi' \in \Sigma(\overrightarrow{\pi})$ such that $\overrightarrow{d}(\overrightarrow{\pi}) = \overrightarrow{d}(\overrightarrow{\pi}')$. By Lemma 1, we can assume that every signed strip σ of $\overrightarrow{\pi}$ of length at least 3 is a signed strip of $\overrightarrow{\pi}'$, i.e., σ remains canonically ordered when zero signs in $\overrightarrow{\pi}$ are replaced by -1 or $+1$ to construct $\overrightarrow{\pi}'$. Then, by Proposition 3, there is an optimal sorting sequence for $\overrightarrow{\pi}'$ such that no reversal cuts a signed strip of this permutation. Let $\overrightarrow{\rho}(i, j)$ be the first reversal in this sorting sequence. We apply it for $\overrightarrow{\pi}$ and denote $\overrightarrow{\pi}^* = \overrightarrow{\rho}(i, j) \circ \overrightarrow{\pi}$. Note that $\overrightarrow{\rho}(i, j)$ does not cut signed strips of $\overrightarrow{\pi}$ of length at least three. We also have that $\overrightarrow{d}(\overrightarrow{\pi}^*) \leq \overrightarrow{d}(\overrightarrow{\rho}(i, j) \circ \pi') = d - 1$. By induction, there is an optimal sorting sequence for $\overrightarrow{\pi}^*$ such that no reversal cuts a signed strip of length at least three. This completes the proof. □

In our study of Whitney switches, we are interested in circular permutations and, therefore, we extend Lemma 2 for such permutations. For this, we define a *partially signed circular permutation* $\overrightarrow{\pi}^c = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle)$, where $(\pi_1, \ldots, \pi_n)$ is a linear permutation and $s_i \in \{-1, +1, 0\}$ for $i \in \{1, \ldots, m\}$, as the class of the linear permutations that can be obtained from $(\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle)$ by rotations and reflections such that every reflection reverse signs. In other words, the linear permutations

$$(\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle),$$
$$(\langle \pi_n, s_n \rangle, \langle \pi_1, s_1 \rangle \ldots, \langle \pi_{n-1}, s_{n-1} \rangle), \ldots, (\langle \pi_2, s_2 \rangle, \ldots, \langle \pi_n, s_n \rangle, \langle \pi_1, s_1 \rangle)$$

and

$$(\langle \pi_n, -s_n \rangle, \ldots, \langle \pi_1, -s_1 \rangle),$$
$$(\langle \pi_1, -s_1 \rangle, \langle \pi_n, -s_n \rangle \ldots, \langle \pi_1, -s_2 \rangle), \ldots, (\langle \pi_{n-1}, -s_2 \rangle, \ldots, \langle \pi_1, -s_1 \rangle, \langle \pi_n, -s_n \rangle)$$

represent the same circular permutation. For $i, j \in \{1, \ldots, n\}$, the reversal

$$\overrightarrow{\rho}^c(i, j) \circ \overrightarrow{\pi}^c = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_{i-1}, s_{i-1} \rangle,$$
$$\langle \pi_j, -s_j \rangle, \ldots, \langle \pi_i, -s_i \rangle, \langle \pi_{j+1}, s_{j+1} \rangle \ldots, \langle \pi_n, s_n \rangle)$$

if $i \leq j$, and

$$\overrightarrow{\rho}^c(i, j) \circ \overrightarrow{\pi}^c = (\langle \pi_n, -s_n \rangle, \ldots, \langle \pi_i, -s_i \rangle,$$
$$\langle \pi_{j+1}, s_{j+1} \rangle, \ldots, \langle \pi_{i-1}, s_{i-1} \rangle, \langle \pi_j, -s_j \rangle \ldots, \langle \pi_1, -s_1 \rangle)$$

otherwise.

In the same way as with partially signed linear permutations, $\overrightarrow{\pi}^c$ is *signed* if each s_i is either -1 or $+1$ and the *signed circular identity permutation* is $\overrightarrow{\iota}^c = (\langle 1, +1 \rangle, \ldots, \langle n, +1 \rangle)$. Also a partially signed circular permutation $\overrightarrow{\pi}^c = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle)$ *agrees in signs* with a signed circular permutation $\overrightarrow{\pi}'^c = (\langle \pi_1, s_1' \rangle, \ldots, \langle \pi_n, s_n' \rangle)$ if $s_i = s_i'$ for $i \in \{1, \ldots, n\}$ such that $s_i \neq 0$, that is, the zero

signs are replaced by either -1 or $+1$ in the signed permutation, and $\Sigma(\overrightarrow{\pi}^c)$ is used to denote the set of all signed circular permutations $\overrightarrow{\pi}'^c$ that agree in signs with $\overrightarrow{\pi}^c$. Then *reversal distance* $\overrightarrow{d}^c(\overrightarrow{\pi}^c, \sigma^c)$, where $\overrightarrow{\sigma}^c$ is a signed circular permutation, is the minimum number or reversals needed to obtain from $\overrightarrow{\pi}^c$ a partially signed circular permutation $\overrightarrow{\pi}'^c$ that agrees in signs with $\overrightarrow{\sigma}^c$, and $\overrightarrow{d}^c(\overrightarrow{\pi}^c) = \overrightarrow{d}^c(\overrightarrow{\pi}^c, \overrightarrow{\iota}^c)$. A sequence of reversals of minimum length that result in a partially signed circular permutation that agrees in signs with $\overrightarrow{\iota}^c$ is an *optimal sorting* sequence.

We exploit the following properties of partially signed permutations. To state them, we need some auxiliary notation. For a partially signed linear permutation $\overrightarrow{\pi} = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle)$, we define the negation $-\overrightarrow{\pi} = (\langle \pi_n, -s_n \rangle, \ldots, \langle \pi_1, -s_1 \rangle)$. For an integer h, we denote $\overrightarrow{\pi} \oplus h = (\langle \pi_{1+h}, s_{1+h} \rangle, \ldots, \langle \pi_{n+h}, s_{n+h} \rangle)$, where it is assumed that $\pi_0 = \pi_n$, $s_0 = s_n$ and the other indices are taken modulo n. The negation corresponds to the reflection and $\oplus$ defines rotations.

LEMMA 3. *Let $\overrightarrow{\pi}$ be partially signed linear permutation, $\overrightarrow{\sigma}$ be a signed permutation, and h be an integer. Then*

$$\min\{\overrightarrow{d}(\overrightarrow{\pi}, \overrightarrow{\sigma}), \overrightarrow{d}(\overrightarrow{\pi}, -\overrightarrow{\sigma})\} = \min\{\overrightarrow{d}(\overrightarrow{\pi} \oplus h, \overrightarrow{\sigma} \oplus h), \overrightarrow{d}(\overrightarrow{\pi} \oplus h, -(\overrightarrow{\sigma} \oplus h))\}.$$

Proof. We show that

$$\min\{\overrightarrow{d}(\overrightarrow{\pi}, \overrightarrow{\sigma}), \overrightarrow{d}(\overrightarrow{\pi}, -\overrightarrow{\sigma})\} \geq \min\{\overrightarrow{d}(\overrightarrow{\pi} \oplus h, \overrightarrow{\sigma} \oplus h), \overrightarrow{d}(\overrightarrow{\pi} \oplus h, -(\overrightarrow{\sigma} \oplus h))\}.$$

The proof of the opposite inequality is symmetric and is done by replacing h by $-h$.

The proof is by induction on the distance between permutations. Let $\overrightarrow{\pi} = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle)$ and $\overrightarrow{\sigma}$ be arbitrary partially signed and signed linear permutations, respectively, with $d = \min\{\overrightarrow{d}(\overrightarrow{\pi}, \overrightarrow{\sigma}), \overrightarrow{d}(\overrightarrow{\pi}, -\overrightarrow{\sigma})\}$. The claim is trivial for $d = 0$. Let $d \geq 1$ and assume that the claim holds for every two permutations at reversal distance at most $d-1$. We assume without loss for generality that $d = \overrightarrow{d}(\overrightarrow{\pi}, \overrightarrow{\sigma})$, as the other case is symmetric.

Consider the corresponding sequence of reversals of length d and assume that $\overrightarrow{\rho}(i,j)$ for $1 \leq i \leq j \leq n$ is the first reversal in the sequence. Recall that

$$\overrightarrow{\pi}' = \overrightarrow{\rho}(i,j) \circ \overrightarrow{\pi} = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_{i-1}, s_{i-1} \rangle,$$
$$\langle \pi_j, -s_j \rangle, \ldots, \langle \pi_i, -s_i \rangle, \langle \pi_{j+1}, s_{j+1} \rangle \ldots, \langle \pi_n, s_n \rangle).$$

Note that either $i \neq 1$ or $j \neq n$, because $d \geq \overrightarrow{d}(\overrightarrow{\pi}, -\overrightarrow{\sigma})$. Let $i' = (i + h) \mod n$ and $j' = (j + h) \mod n$ assuming that $n \mod n = n$. Let $\overrightarrow{\pi}^* = \overrightarrow{\pi} \oplus h$.

Suppose that $i' \leq j'$. Then $\overrightarrow{\rho}(i',j') \circ \overrightarrow{\pi}^* = \overrightarrow{\pi}' \oplus h$. By the inductive assumption,

$$d - 1 \geq \min\{\overrightarrow{d}(\overrightarrow{\pi}', \overrightarrow{\sigma}), \overrightarrow{d}(\overrightarrow{\pi}', -\overrightarrow{\sigma})\}$$
$$\geq \min\{\overrightarrow{d}(\overrightarrow{\pi}' \oplus h, \overrightarrow{\sigma} \oplus h), \overrightarrow{d}(\overrightarrow{\pi}' \oplus h, -(\overrightarrow{\sigma} \oplus h))\}$$
$$\geq \min\{\overrightarrow{d}(\overrightarrow{\pi} \oplus h, \overrightarrow{\sigma} \oplus h), \overrightarrow{d}(\overrightarrow{\pi} \oplus h, -(\overrightarrow{\sigma} \oplus h))\} - 1,$$

and the claim follows.

Assume that $i' > j'$. Let $i'' = j' + 1$ and $j'' = i' - 1$. Notice that since $i \neq 1$ or $j \neq n$, $i'' \leq j''$. Then $\overrightarrow{\rho}(i'',j'') \circ \overrightarrow{\pi}^* = -(\overrightarrow{\pi}' \oplus h)$. Using the inductive assumption

we obtain that

$$
\begin{aligned}
d - 1 &\geq \min\{\overrightarrow{d}(\overrightarrow{\pi}', \overrightarrow{\sigma}), \overrightarrow{d}(\overrightarrow{\pi}', -\overrightarrow{\sigma})\} \\
&\geq \min\{\overrightarrow{d}(\overrightarrow{\pi}' \oplus h, \overrightarrow{\sigma} \oplus h), \overrightarrow{d}(\overrightarrow{\pi}' \oplus h, -(\overrightarrow{\sigma} \oplus h))\} \\
&\geq \min\{\overrightarrow{d}(-(\overrightarrow{\pi}' \oplus h), \overrightarrow{\sigma} \oplus h), \overrightarrow{d}(-(\overrightarrow{\pi}' \oplus h), -(\overrightarrow{\sigma} \oplus h))\} \\
&\geq \min\{\overrightarrow{d}(\overrightarrow{\pi} \oplus h, \overrightarrow{\sigma} \oplus h), \overrightarrow{d}(\overrightarrow{\pi} \oplus h, -(\overrightarrow{\sigma} \oplus h))\} - 1.
\end{aligned}
$$

This completes the proof. $\qquad\square$

LEMMA 4. *Let $\overrightarrow{\pi}$ be a partially signed circular permutation. Then*

$$
\overrightarrow{d}^{\,c}(\overrightarrow{\pi}^c) = \min\{\overrightarrow{d}(\overrightarrow{\sigma}) \mid \overrightarrow{\sigma} \in \overrightarrow{\pi}^c\}.
$$

Proof. Clearly, for every $\overrightarrow{\sigma} \in \overrightarrow{\pi}^c$, $\overrightarrow{d}^{\,c}(\overrightarrow{\pi}^c) \leq \overrightarrow{d}(\overrightarrow{\sigma})$. Therefore, we have to show that there is $\overrightarrow{\sigma} \in \overrightarrow{\pi}^c$ such that $\overrightarrow{d}^{\,c}(\overrightarrow{\pi}^c) \geq \overrightarrow{d}(\overrightarrow{\sigma})$. Let $\overrightarrow{\pi}^c = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle)$ and let $d = \overrightarrow{d}^{\,c}(\overrightarrow{\pi}^c)$.

We claim that there is an integer h such that for the partially signed linear permutation $\overrightarrow{\pi} = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle)$, $\min\{\overrightarrow{d}(\overrightarrow{\pi}, \overrightarrow{\iota} \oplus h), \overrightarrow{d}(\overrightarrow{\pi}, -(\overrightarrow{\iota} \oplus h))\} \leq d$.

The proof is by the induction on d. The claim is trivial if $d = 0$. Let $d \geq 1$ and assume that the claim holds for all partially signed circular permutations $\overrightarrow{\pi}'$ with $\overrightarrow{d}^{\,c}(\overrightarrow{\pi}') \leq d - 1$. Consider an optimal sorting sequence for $\overrightarrow{\pi}^c$ and let $\rho^c(i, j)$ be the first reversal in the sequence. Let $\overrightarrow{\pi}'^c = \overrightarrow{\rho}^c(i, j) \circ \overrightarrow{\pi}^c$.

Suppose that $i \leq j$. Let $\overrightarrow{\pi}' = \overrightarrow{\rho}(i, j) \circ \overrightarrow{\pi}$. By the inductive assumption, there is h such that $\min\{\overrightarrow{d}(\overrightarrow{\pi}', \overrightarrow{\iota} \oplus h), \overrightarrow{d}(\overrightarrow{\pi}', -(\overrightarrow{\iota} \oplus h))\} \leq d - 1$. Therefore, $\min\{\overrightarrow{d}(\overrightarrow{\pi}, \overrightarrow{\iota} \oplus h), \overrightarrow{d}(\overrightarrow{\pi}, -(\overrightarrow{\iota} \oplus h))\} \leq d$.

Let $i > j$. If $(j+1) - i = 0 \mod n$, that is, the indices j and i are consecutive in the cycle ordering, then $\overrightarrow{\rho}^c(i, j)$ just reflects $\overrightarrow{\pi}^c$ contradicting the optimality of the chosen sorting sequence. Hence, for $i' = j + 1$ and $j' = i - 1$, we have that $i' \leq j'$. Let $\overrightarrow{\pi}' = \overrightarrow{\rho}(i', j') \circ \overrightarrow{\pi}$. By induction, there is an integer h such that $\min\{\overrightarrow{d}(-\overrightarrow{\pi}', \overrightarrow{\iota} \oplus h), \overrightarrow{d}(-\overrightarrow{\pi}', -(\overrightarrow{\iota} \oplus h))\} \leq d - 1$. We have that $\min\{\overrightarrow{d}(-\overrightarrow{\pi}', \overrightarrow{\iota} \oplus h), \overrightarrow{d}(-\overrightarrow{\pi}', -(\overrightarrow{\iota} \oplus h))\} = \min\{\overrightarrow{d}(\overrightarrow{\pi}', \overrightarrow{\iota} \oplus h), \overrightarrow{d}(\overrightarrow{\pi}', -(\overrightarrow{\iota} \oplus h))\}$ and, therefore, $\min\{\overrightarrow{d}(\overrightarrow{\pi}, \overrightarrow{\iota} \oplus h), \overrightarrow{d}(\overrightarrow{\pi}, -(\overrightarrow{\iota} \oplus h))\} \leq d$. This competes the proof of the auxiliary claim.

To prove the lemma, observe that by Lemma 3, we obtain that

$$
\min\{\overrightarrow{d}(\overrightarrow{\pi} \oplus (-h), \overrightarrow{\iota}), \overrightarrow{d}(\overrightarrow{\pi} \oplus (-h), -\overrightarrow{\iota})\} = \min\{\overrightarrow{d}(\overrightarrow{\pi}, \overrightarrow{\iota} \oplus h), \overrightarrow{d}(\overrightarrow{\pi}, -(\overrightarrow{\iota} \oplus h))\} \leq d.
$$

If $\overrightarrow{d}(\overrightarrow{\pi} \oplus (-h), \overrightarrow{\iota}) \leq \overrightarrow{d}(\overrightarrow{\pi} \oplus (-h), -\overrightarrow{\iota})$, we set $\overrightarrow{\sigma} = \overrightarrow{\pi} \oplus (-h)$, and $\overrightarrow{\sigma} = -(\overrightarrow{\pi} \oplus (-h))$ otherwise. It is straightforward to see that $\overrightarrow{\sigma} \in \overrightarrow{\pi}^c$ and this completes the proof. $\quad\square$

The notion of signed strips can be extended for partially signed circular permutations in a natural way. More formally, this is done as follows. Let $\overrightarrow{\pi}^c = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_n, s_n \rangle)$ be a partially signed circular permutation. For $1 \leq i \leq j \leq n$, we say that $(\langle \pi_i, s_i \rangle, \ldots, \langle \pi_j, s_j \rangle)$ and $(\langle \pi_{j+1}, s_{j+1} \rangle, \ldots, \langle \pi_n, s_n \rangle, \langle \pi_1, s_1 \rangle, \ldots, \langle \pi_i, s_i \rangle)$ are *intervals* of $\overrightarrow{\pi}^c$. An interval is a *signed block* if it either has size one or for every two consecutive elements $\langle \pi_{i-1}, s_{i-1} \rangle, \langle \pi_i, s_i \rangle$, $|\pi_{i-1} - \pi_i| \leq 1$ and, moreover, if the elements of the interval are in increasing order, then all the signs $s_i \in \{0, +1\}$, and if they are in decreasing order, then all the signs $s_i \in \{0, -1\}$. A *signed strip* is an inclusion maximal signed block. A reversal $\overrightarrow{\rho}^c(p, q)$ *cuts* an interval if the reversed

part includes at least one element of the interval and excludes at least one element of the interval.

We conjecture that the result of Hannenhalli and Pevzner [16] can be extended for partially signed circular permutations in the same way as for the linear case in Lemma 2. However, it seems that for this, the variant of Lemma 1 for circular permutations should be proved. This can be done by following and adjusting the arguments from [16]. The proof of Lemma 1 is nontrivial and is based on the deep duality theorem of Hannenhalli and Pevzner [17], which is also is stated for linear permutations. Hence, proving the circular analogue of Lemma 1 would demand a lot of technical work and this goes beyond of the scope of our paper. Therefore, we show the simplified claim that can be derived from Lemma 2 and is sufficient for our purposes.

LEMMA 5. *For a signed circular permutation $\overrightarrow{\pi}^c$, there is an optimal sorting sequence such that no reversal in the sequence cuts the interval formed by a signed strip of $\overrightarrow{\pi}^c$ of length at least 5.*

Notice that we do not claim that no reversal cuts a strip of length at least 5 that is obtained by performing the previous reversals; only the long strips of the initial permutation $\overrightarrow{\pi}^c$ are not cut by any reversal in the sorting sequence.

Proof. Let $\overrightarrow{\pi}^c$ be a partially signed circular permutation. By Lemma 4, there is a partially signed linear permutation $\overrightarrow{\sigma} \in \overrightarrow{\pi}^c$ such that $d = \overrightarrow{d}(\sigma) = \overrightarrow{d}^c(\overrightarrow{\pi}^c)$. Let $\overrightarrow{\sigma} = (\langle \sigma_1, s_1 \rangle, \ldots, \langle \sigma_n, s_n \rangle)$. Note that, by definition, we can write that $\overrightarrow{\pi}^c = (\langle \sigma_1, s_1 \rangle, \ldots, \langle \sigma_n, s_n \rangle)$. We assume that $\overrightarrow{d}^c(\overrightarrow{\pi}^c) \geq 1$. We consider three cases.

Case 1. Every signed strip of length at least 5 of $\overrightarrow{\pi}^c$ is a signed strip of the linear permutation $\overrightarrow{\sigma}$. Consider an optimal sorting sequence for $\overrightarrow{\sigma}$ that does not cut strips of length at least 5 that exists by Lemma 2. Then this sequence is an optimal sorting sequence for $\overrightarrow{\pi}^c$ satisfying the claim.

Case 2. There is a unique signed strip $\omega = (\langle \sigma_i, s_i \rangle, \ldots, \langle \sigma_j, s_j \rangle)$ for $i \leq i < j \leq n$ of $\overrightarrow{\pi}^c$ with length at least 5 that is not a signed strip of σ. Then

$$(3.1) \qquad \omega = (\langle p, s_i \rangle, \ldots, \langle n, s_{n-p} \rangle, \langle 1, s_{n-p+1} \rangle, \ldots, \langle ((p+j-i) \mod n), s_j \rangle)$$

for some $p \geq n - (j - i) + 1$ or, symmetrically,

$$\omega = (\langle p, s_i \rangle, \ldots, \langle 1, s_{n-p} \rangle, \langle n, s_{n-p+1} \rangle, \ldots, \langle n - (j-i) + p, s_j \rangle).$$

Using symmetry, we assume without loss of generality that ω is of form (3.1) and write that $\omega = \omega'\omega''$, where $\omega' = (\langle p, s_i \rangle, \ldots, \langle n, s_{n-p} \rangle)$ and $\omega'' = (\langle 1, s_{n-p+1} \rangle \ldots, \langle ((p+j-i) \mod n), s_j \rangle)$. Since $j - i \geq 4$, either $|\omega'| \geq 3$ or $|\omega''| \geq 3$. Assume that $|\omega'| \geq 3$ as the other case is completely symmetric. By Lemma 2, there is an optimal sorting sequence $\mathcal{S}$ for $\overrightarrow{\sigma}$ that does not cut strips of length at least 3. In particular, ω' is not cut by any reversal in the sequence. We modify $\mathcal{S}$ as follows for every reversal:

- exclude the elements of ω'' from the reversed interval and its complement,
- if the reversed interval includes either w' or $-w'$, then replace w' by w,
- if the complement of the reversed interval contains either w' or $-w'$, then replace w' by w.

In other words, whenever we reverse w', we reverse it together with w'', and if we do not reverse w', we keep w'' together with w' and do not reverse the elements of w''. Let $\mathcal{S}'$ be the obtained sequence. It is straightforward to verify that every step of $\mathcal{S}'$ is indeed a reversal and no reversal cuts a strip of $\overrightarrow{\pi}^c$ of length at least 5. Moreover,

after performing all the reversals of $\mathcal{S}'$ we obtain the partially signed permutation that agrees in signs with $\overrightarrow{\iota} \oplus ((p + j - i) \mod n)$ that is in $\overrightarrow{\iota}^c$. This means that $\mathcal{S}'$ is a sorting sequence for $\overrightarrow{\pi}^c$ of length d.

Case 3. There are $1 \le i < j \le n$ such that

$$\omega = (\langle \sigma_j, s_j \rangle, \ldots, \langle \sigma_n, s_n \rangle, \langle \sigma_1, s_1 \rangle, \ldots, \langle \sigma_i, s_i \rangle)$$

is a strip of $\overrightarrow{\pi}^c$. Let $\overrightarrow{\sigma}' = \sigma \oplus (-i) = (\langle \sigma_1', s_1' \rangle, \ldots, \langle \sigma_n', s_n' \rangle)$. By Lemma 3,

$$d = \min\{\overrightarrow{d}(\overrightarrow{\sigma}', \overrightarrow{\iota} \oplus (-i)), \overrightarrow{d}(\overrightarrow{\sigma}, -(\overrightarrow{\iota} \oplus (-i)))\}.$$

Since the cases are symmetric, assume without loss of generality that $\overrightarrow{d}(\overrightarrow{\sigma}', \overrightarrow{\iota} \oplus (-i)) = d$. Consider $\overrightarrow{\sigma}'' = \sigma \oplus (-i) = (\langle \sigma_1'', s_1' \rangle, \ldots, \langle \sigma_n'', s_n' \rangle)$, where $\sigma_i'' = (\sigma_i' + i) \mod n$ (assuming that $n \mod n = n$). We have that $\overrightarrow{d}(\overrightarrow{\sigma}'') = d$ and sorting of $\overrightarrow{\sigma}''$ is equivalent to computing the minimum sequence of reversals needed to transform σ' to a partially signed permutation that agrees in signs with $\overrightarrow{\iota} \oplus (-i)$. Note that sorting of the circular partially signed permutation $\overrightarrow{\sigma}''^c$ is equivalent to sorting $\overrightarrow{\rho}^c$ and $\overrightarrow{\sigma}''^c$ has no strips including $\langle \sigma'', s_n' \rangle$ and $\langle \sigma_1'', s_1 \rangle$. Finally, observe that either every signed strip of length at least 5 of $\overrightarrow{\sigma}''^c$ is a signed strip of the linear permutation $\overrightarrow{\sigma}''$ and we are in Case 1 or there is a unique signed strip $\omega'' = (\langle \sigma_i'', s_i' \rangle, \ldots, \langle \sigma_j'', s_j'' \rangle)$ for $i \le i < j \le n$ of $\overrightarrow{\sigma}''^c$ with length at least 5 that is not a signed strip of σ'' and we are in Case 2. □

We conclude the section by observing that if the elements of a partially signed circular permutation are ordered, then the sorting can be done easily. We say that the reversal $\overrightarrow{\rho}^c(i, j)$ is *trivial* if $i = j$.

LEMMA 6. *Let* $\overrightarrow{\pi}^c = (\langle 1, s_1 \rangle, \ldots, \langle n, s_n \rangle)$. *Then* $\overrightarrow{d}^c(\overrightarrow{\pi}^c) = |I|$, *where* $I = \{i \mid 1 \le i \le n, s_i = -1\}$ *and the reversals* $\overrightarrow{\rho}(i, i)$ *for* $i \in I$ *compose an optimal sorting sequence.*

Proof. Let $i \in I$. Let $\mathcal{S}$ be an optimal sorting sequence. We assume that $\mathcal{S}$ does not contain reversals $\overrightarrow{\rho}^c(i - 2, i)$ for $i \in \{1, \ldots, n\}$ (as before, we take the values modulo n assuming that $n \mod n = n$), because they are equivalent to $\overrightarrow{\rho}^c(i, i)$. Observe that if $i \in I$, then the intervals $(\langle (i-1), s_{i-1} \rangle, \langle i, s_i \rangle)$ and $(\langle i, s_i \rangle, \langle (i+1), s_{i+1} \rangle)$ should be split by some reversals from $\mathcal{S}$. Moreover, we can observe the following for $i - 1, i \in I$. Assume that $\overrightarrow{\rho}^c(p, q)$ is the first reversal that splits $(\langle (i-1), s_{i-1} \rangle, \langle i, s_i \rangle)$ and assume that $i - 1$ keeps its sign s_{i-1}. Let $\sigma = (\langle (i-1), s_{i-1} \rangle, \langle j, s_j \rangle)$ be the interval composed by $(\langle (i-1), s_{i-1} \rangle$ and the next element after applying $\overrightarrow{\rho}^c(p, q)$. If the reversal is trivial, then $\sigma = (\langle (i-1), s_{i-1} \rangle, \langle i, -s_i \rangle)$ and σ should be split again. If $j \ne i$, then we have to split σ, because $j \ne i - 2$ and, therefore, $|j - (i-1)| > 1$. These observations imply that $\mathcal{S}$ contains at least $|I|$ reversals. Therefore, the sorting sequence formed by the reversals $\overrightarrow{\rho}(i, i)$ for $i \in I$ is optimal. □

4. Tutte decomposition and 2-isomorphisms. In this section we provide a number of auxiliary results about 2-isomorphisms and Tutte decompositions.

Recall that for two n-vertex 2-connected graphs G and H, a bijective mapping $\varphi \colon E(G) \to E(H)$ is a 2-isomorphism if φ and φ^{-1} preserve cycles. We also say that an isomorphism $\psi \colon V(G) \to V(H)$ is a φ-isomorphism if for every edge $uv \in E(G)$, $\varphi(uv) = \psi(u)\psi(v)$, and G and H are φ-isomorphic if there is an isomorphism G to H that is a φ-isomorphism. We need the following folklore observation about φ-isomorphisms that we prove for completeness. For this, we extend φ on sets of edges in standard way, that is, $\varphi(A) = \{\varphi(e) \mid e \in A\}$ and $\varphi(\emptyset) = \emptyset$.

LEMMA 7. *Let G and H be n-vertex 2-connected 2-isomorphic graphs with a 2-isomorphism φ. Then G and H are φ-isomorphic if and only if there is a bijective mapping $\psi\colon V(G) \to V(H)$ such that for every $v \in V(G)$, $\varphi(E_G(v)) = E_H(\psi(v))$. Moreover, G and H are φ-isomorphic if and only if φ bijectively maps the family of the sets of edges $\{E_G(v) \mid v \in V(G)\}$ to the family $\{E_H(v) \mid v \in V(H)\}$; furthermore, this property can be checked in polynomial time.*

Proof. If ψ is an φ-isomorphism of G to H, then $\varphi(E_G(v)) = E_H(\psi(v))$ for all $v \in V(G)$ by the definition. For the opposite direction, assume that $\psi\colon V(G) \to V(H)$ is a bijection such that $\varphi(E_G(v)) = E_H(\psi(v))$ for every $v \in V(G)$. Suppose that u and v are distinct vertices of G. We claim that u and v are adjacent in G if and only if $\psi(u)$ and $\psi(v)$ are adjacent in H. Suppose that u and v are adjacent in G. Then $E_G(u) \cap E_G(v) = \{uv\}$. Therefore, $E_H(\psi(u)) \cap E_H(\psi(v)) = \varphi(E_G(u)) \cap \varphi(E_G(v)) = \{\varphi(uv)\}$. This means that $\psi(u)$ and $\psi(v)$ are adjacent in H and $\psi(u)\psi(v) = \varphi(uv)$. If u and v are not adjacent, then $E_G(u) \cap E_G(v) = \emptyset$ and $E_H(\psi(u)) \cap E_H(\psi(v)) = \varphi(E_G(u)) \cap \varphi(E_G(v)) = \emptyset$, that is, $\psi(u)$ and $\psi(v)$ are not adjacent in H.

The second claim of the lemma immediately follows from the first. $\square$

By Lemma 7, we can restate the task of WHITNEY SWITCHES and ask whether it is possible to obtain a graph G' by performing at most k Whitney switches starting from G with the property that the extension of φ to the family of sets $\{E_{G'}(v) \mid v \in V(G')\}$ bijectively maps this family to $\{E_H(v) \mid v \in V(H)\}$.

We use Whitney's theorem [36] (see also [30]).

PROPOSITION 4 ([36]). *Let G and H be n-vertex graphs and let φ be a 2-isomorphism of G to H. Then there is a finite sequence of Whitney switches such that the graph G' obtained from G by these switches is φ-isomorphic to H.*

We also use the property of 3-connected graphs explicitly given by Truemper [30]. It also can be derived from Proposition 4.

PROPOSITION 5 ([30]). *Let G and H be 3-connected n-vertex graphs and let φ be a 2-isomorphism of G to H. Then G and H are φ-isomorphic.*

Throughout this section we assume that G and H are n-vertex 2-connected graphs and φ is a 2-isomorphism of G to H. Let also $\mathcal{T}^{(1)} = (T^{(1)}, \{X_t^{(1)}\}_{t \in V(T^{(1)})})$ and $\mathcal{T}^{(2)} = (T^{(2)}, \{X_t^{(2)}\}_{t \in V(T^{(2)})})$ be the Tutte decompositions of G and H, respectively, and denote by $(W_2^{(h)}, W_{\geq 3}^{(h)})$ the partition of $V(T^{(h)})$ satisfying (T4)–(T8) for $h = 1, 2$.

The following lemma is crucial for us.

LEMMA 8. *There is an isomorphism α of $T^{(1)}$ to $T^{(2)}$ such that*

(i) *for every $t \in V(T^{(1)})$, $|X_t^{(1)}| = |X_{\alpha(t)}^{(2)}|$, in particular, $t \in W_2^{(1)}$ ($t \in W_{\geq 3}^{(1)}$, respectively) if and only if $\alpha(t) \in W_2^{(2)}$ ($\alpha(t) \in W_{\geq 3}^{(2)}$, respectively),*

(ii) *for every $t \in W_{\geq 3}^{(1)}$, the torso of $X_t^{(1)}$ is a 3-connected graph (a cycle, respectively) if and only if the torso of $X_{\alpha(t)}^{(2)}$ is a 3-connected graph (a cycle, respectively),*

(iii) *for every $t \in V(T^{(1)})$, $\varphi(E(G[X_t^{(1)}]) = E(H[X_{\alpha(t)}^{(2)}])$.*

Proof. By Proposition 4, there is a finite sequence of Whitney switches such that the graph G' obtained from G by these switches is φ-isomorphic to H. We prove the lemma by induction on the number of switches. The claim is straightforward if this number is zero, because G and H are φ-isomorphic. It is sufficient to observe that the Tutte decomposition is unique by Proposition 1 and then use Lemma 7. Assume that

the sequence has length $\ell \geq 1$ and the claim of the lemma holds for the sequences of length at most $\ell - 1$.

Let (A, B) be a Whitney separation of G such that the first switch is done with respect to (A, B). Denote by $\{u, v\} = A \cap B$. Denote by G' the graph obtained from G by the Whitney switch with respect to (A, B). Recall that G' is constructed by replacing each edge $ux \in E(G)$ for $x \in B \setminus A$ by vx and by replacing each edge vx for $x \in B \setminus A$ by ux. Recall also that we denote by $\sigma_{(A,B)}$ the mapping of the edges of G to the edges of the graph G' obtained from G by the Whitney switch with respect to (A, B) that corresponds to the switch. We have that $\sigma_{(A,B)}$ is a 2-isomorphism of G to G' and, by our convention, the set of edges remains the same and we only modify incidences of some of them, that is, $\sigma_{(A,B)}$ is the identity mapping. Since H is obtained from G' by $\ell - 1$ switches, it is sufficient to show the claim for G'. We do it by constructing the Tutte decomposition of G' from the decomposition of G.

Suppose first that $A \cap B = X_t^{(1)}$ for some $t \in W_2^{(1)}$. By the definition of the Tutte decomposition, for each $s \in V(T)$, either $X_s^{(1)} \subseteq A$ or $X_s^{(1)} \subseteq B$. We construct the tree decomposition $\mathcal{T}' = (T^{(1)}, \{X_s'\}_{s \in V(T^{(1)})})$. For every $s \in V(T^{(1)})$ such that $X_s^{(1)} \subseteq A$, we define $X_s' = X_s^{(1)}$. Similarly, if $X_s^{(1)} \subseteq B$ and $u, v \notin X_s^{(1)}$, $X_s' = X_s^{(1)}$. For all $s \in V(T^{(1)})$ such that $X_s^{(1)} \subseteq B$, $s \neq t$, and $\{u, v\} \cap X_s^{(1)} \neq \emptyset$, we construct X_s' from $X_s^{(1)}$ as follows:

   (a) replace u by v if $u \in X_s^{(1)}$ and $v \notin X_s^{(1)}$,

   (b) replace v by u if $v \in X_s^{(1)}$ and $u \notin X_s^{(1)}$.

It is straightforward to verify that $\mathcal{T}'$ is the Tutte decomposition and α that maps the nodes of $T^{(1)}$ to themselves satisfies (i)–(iii).

Assume now that $A \cap B \neq X_t^{(1)}$ for all $t \in W_2^{(1)}$. By the definition of the Tutte decomposition, this means that $u, v \in X_t^{(1)}$ for some $t \in W_{\geq 3}^{(1)}$ such that the torso of $X_t^{(1)}$ is a cycle C and u, v are nonadjacent vertices of C. Notice that $Z_A = X_t^{(1)} \cap A$ and $Z_B = X_t^{(1)} \cap B$ induce distinct (u, v)-paths in C. We again construct the tree decomposition $\mathcal{T}' = (T^{(1)}, \{X_s'\}_{s \in V(T^{(1)})})$. Notice that for each $s \in V(T^{(1)})$ such that $s \neq t$, either $X_s^{(1)} \subseteq A$ or $X_s^{(1)} \subseteq B$. For all such s, we define X_s' in exactly the same way as in the previous case. We define $X_t' = X_t^{(1)}$. It is straightforward to verify that $\mathcal{T}'$ is a tree decomposition of G'. Since the torso of $X_t^{(1)}$ is a cycle composed by the paths with the vertices Z_A and Z_B, the torso of X_s' in $\mathcal{T}'$ is a cycle as well. This implies that $\mathcal{T}'$ is the Tutte decomposition of G'. Then α that maps the nodes of $T^{(1)}$ to themselves satisfies (i)–(iii). ☐

Let F be a 2-connected graph. Let also $\mathcal{T} = (T, \{X_t\}_{t \in V(T)})$ be the Tutte decomposition of F and let $(W_2, W_{\geq 3})$ be the partition of $V(T)$ satisfying (T4)–(T8). We denote by $\widehat{F}$ the graph obtained from F by making the vertices of X_t adjacent for every $t \in W_2$. We say that $\widehat{F}$ is the *enhancement* of F. Note that $\mathcal{T}$ is the Tutte decomposition of $\widehat{F}$ and the torso of each bag X_t is $\widehat{F}[X_t]$. Notice also that (A, B) is a Whitney separation of F if and only if (A, B) is a Whitney separation of $\widehat{F}$. We also say that F is *enhanced* if $F = \widehat{F}$.

To simplify the arguments in our proofs, it is convenient for us to switch from 2-isomorphisms of graphs to 2-isomorphisms of their enhancements. By Lemma 8, there is an isomorphism α of $T^{(1)}$ to $T^{(2)}$ satisfying conditions (i)—(ii) of the lemma. We

define the *enhanced* mapping $\widehat{\varphi}\colon E(\widehat{G}) \to E(\widehat{H})$ such that $\widehat{\varphi}(e) = \varphi(e)$ for $e \in E(G)$, and for each $e \in E(\widehat{G}) \setminus E(G)$ with its end-vertices in $X_t^{(1)}$ for some $t \in W_2^{(1)}$, we define $\widehat{\varphi}(e)$ be the edge with the end-vertices in $X_{\alpha(t)}^{(2)}$.

LEMMA 9. *The mapping $\widehat{\varphi}$ is a 2-isomorphism of $\widehat{G}$ to $\widehat{H}$. Moreover, a sequence of Whitney switches makes G φ-isomorphic to H if and only if the same sequence makes $\widehat{G}$ $\widehat{\varphi}$-isomorphic to $\widehat{H}$.*

Proof. Note that by Lemma 8, $\widehat{\varphi}$ is a bijection. It is sufficient to show the second claim of the lemma, because if a sequence of Whitney switches makes G φ-isomorphic to H, then φ is a 2-isomorphism of G to H. Notice that $\widehat{G}$ and $\widehat{H}$ have the same separators of size 2 as G and H, respectively, by the definition of the Tutte decomposition. Therefore, given a sequence of Whitney switches of G, the same sequence can be performed on $\widehat{G}$. Then Lemmas 7 and 8 imply that if a sequence of Whitney switches makes G φ-isomorphic to H, then the same sequence makes $\widehat{G}$ $\widehat{\varphi}$-isomorphic to $\widehat{H}$. Since it is straightforward to see that if a sequence of Whitney switches makes $\widehat{G}$ $\widehat{\varphi}$-isomorphic to $\widehat{H}$, then the same sequence makes G φ-isomorphic to H, and the second claim holds. $\square$

Lemma 9 allows us to consider an enhanced graph and this is useful, because we can strengthen the claim of Lemma 8.

LEMMA 10. *Let G and H be enhanced graphs. Then there is an isomorphism α of $T^{(1)}$ to $T^{(2)}$ such that conditions* (i)–(iii) *of Lemma 8 are fulfilled and, moreover,*

(iv) *for every $t \in V(T^{(1)})$, $G[X_t^{(1)}]$ is isomorphic to $H[X_{\alpha(t)}^{(2)}]$.*

Moreover, if $G[X_t^{(1)}]$ is 3-connected, then $G[X_t^{(1)}]$ is φ-isomorphic to $H[X_{\alpha(t)}^{(2)}]$.

Proof. We have that $G[X_t^{(1)}]$ is isomorphic to $H[X_{\alpha(t)}^{(2)}]$ for $t \in W_2^{(1)}$, because G and H are enhanced graphs. Let $t \in W_{\geq 3}^{(1)}$. By conditions (i) and (ii) of Lemma 8, $|X_t^{(1)}| = |X_{\alpha(t)}^{(2)}|$ and $G[X_t^{(1)}]$ is a 3-connected graph (a cycle, respectively) if and only if $H[X_{\alpha(t)}^{(2)}]$ is a 3-connected graph (a cycle, respectively). If $G[X_t^{(1)}]$ and $H[X_{\alpha(t)}^{(2)}]$ are cycles, then they are isomorphic. Assume that $G[X_t^{(1)}]$ and $H[X_{\alpha(t)}^{(2)}]$ are 3-connected. By (iii), $\varphi(E(G[X_t^{(1)}]) = E(H[X_{\alpha(t)}^{(2)}])$. This implies that φ is a 2-isomorphism of $G[X_t^{(1)}]$ to $H[X_{\alpha(t)}^{(2)}]$. By Proposition 5, $G[X_t^{(1)}]$ are $H[X_{\alpha(t)}^{(2)}]$ isomorphic and, moreover, φ-isomorphic. $\square$

For the remaining part of the sections, we assume that G and H are enhanced graphs and α is the isomorphism of $T^{(1)}$ to $T^{(2)}$ satisfying conditions (i)–(iv) of Lemmas 8 and 10.

Our next aim is to investigate properties of the sequences of Whitney switches that are used in solutions for WHITNEY SWITCHES. For a sequence $\mathcal{S}$ of Whitney switches such that the graph G' obtained from G by applying this sequence is φ-isomorphic to H, we say that $\mathcal{S}$ is an *H-sequence*. We also say that $\mathcal{S}$ is *minimum* if $\mathcal{S}$ has minimum length.

We show that Whitney switches in an H-sequence can be rearranged in a special way that simplifies the analysis. Recall that Whitney switches are not symmetric with respect to the names of vertices and the end-vertices of edges may change their names. In particular, this makes a rearrangement of switches complicated. To avoid additional technicalities, we consider Whitney switches with respect to edge Whitney separations. In the same way as above a sequence $\mathcal{S}_E$ of Whitney switches with

respect to edge Whitney separations such that the graph G' obtained from G by applying this sequence is φ-isomorphic to H is said to be an *H-sequence*, and $\mathcal{S}_E$ is *minimum* if $\mathcal{S}_E$ has minimum length. Observe that, given an H-sequence $\mathcal{S}$, we can construct the H-sequence $\mathcal{S}_E$ with respect to edge Whitney separations by replacing each switch with respect to separation (A, B) in $\mathcal{S}$ by the switch with respect to the edge separation (L, R) with $L = E(\hat{G}[A])$ and $R = E(\hat{G}[B])$, where $\hat{G}$ is the graph obtained from G by the previous switches. To simplify notation, we write $\mathcal{S}_E$ for the sequence constructed from $\mathcal{S}$ in this way. Symmetrically, given an H-sequence $\mathcal{S}_E$ with respect to edge Whitney separations, we can construct $\mathcal{S}$ by replacing each edge Whitney separation (L, R) by $(V(L), V(R))$.

The crucial observation exploited in the rest of this section and in section 5 is that, by Lemma 10, the edge bijection $\sigma_{(A,B)}$ corresponding to a Whitney switch with respect to (A, B) maps $S = E(G[X_t^{(1)}])$ to itself in such a way that $V(S)$ remains a bag of the Tutte decomposition corresponding to a node $t \in V(T^{(1)})$. Hence, we can use $T^{(1)}$ for the tree in the Tutte decompositions for all the graphs obtained from G by applying Whitney switches and can use $E(G[X_t^{(1)}])$ to denote the set of edges of the bag corresponding to t.

Let X be a set of vertices (a set of edges, respectively). We say that the Whitney switch with resect to a Whitney separation (A, B) (an edge Whitney separation (L, R)) *splits* X if $X \setminus A \neq \emptyset$ and $X \setminus B \neq \emptyset$ ($V(X) \setminus V(L) \neq \emptyset$ and $V(X) \setminus V(R) \neq \emptyset$, respectively).

Let $\mathcal{S}_E$ be an H-sequence. For $t \in V(T^{(1)})$, we say that an edge Whitney switch (L, R) in $\mathcal{S}_E$ is a *t-switch* if one of the following holds:

(i) $t \in W_2^{(1)}$ and $V(L) \cap V(R) = V(e)$, where e is the unique edge of $G[X_t^{(1)}]$, are consecutive in $\mathcal{S}_E$.

(ii) $t \in W_{\geq 3}^{(1)}$ and (L, R) splits $E(G[X_t^{(1)}])$.

Observe that every Whitney switch (L, R) is a t-switch for unique $t \in V(T^{(1)})$.

We say that an H-sequence $\mathcal{S}_E$ with respect to edge Whitney switches is *t-sorted* if the t-switches are consecutive in $\mathcal{S}_E$ and, furthermore, are the first switches in the sequence. An H-sequence $\mathcal{S}$ is *t-sorted* if $\mathcal{S}_E$ is sorted.

LEMMA 11. *Given a minimum H-sequence $\mathcal{S}_E$ and $t \in V(T^{(1)})$ such that $\mathcal{S}_E$ contains a t-switch, there is a minimum t-sorted H-sequence $\hat{\mathcal{S}}_E$ with exactly the same switches as in $\mathcal{S}_E$.*

Proof. The proof is by induction on the number of switches in $\mathcal{S}_E$. The claim is trivial in the base case when such a sequence contains a unique Whitney switch. Assume that the minimum number of switches in an H-sequence is at least two.

Let $\mathcal{S}'_E$ be an H-sequence with exactly the same switches as in $\mathcal{S}_E$ such that a t-switch has the minimum number in the sequence. We claim that the first switch in $\mathcal{S}'_E$ is a t-switch.

We are targeting toward a contradiction. Assume that the first t-switch σ in $\mathcal{S}'_E$ is performed with respect to an edge Whitney separation (L, R) and that it is not the first switch in the considered H-sequence. Let a switch σ' with respect to an edge Whitney separation (L', R') be the predecessor of σ in $\mathcal{S}_E$. Denote by $\hat{G}$ the graph obtained from G by the switches in $\mathcal{S}'_E$ that are prior to σ'. Suppose that $t \in W_2^{(1)}$, that is, $V(L) \cap V(R) = V(e)$, where e is the unique edge of $G[X_t^{(1)}]$. Then either $V(L') \subset V(L)$ or $V(R') \subset V(R)$ in $\hat{G}$. Let now $t \in W_{\geq 3}^{(1)}$, that is, (L, R) splits $E(G[X_t^{(1)}])$. Because (L', R') is not a t-switch, we again have that either

$V(L') \subset V(L)$ or $V(R') \subset V(R)$ in $\hat{G}$, i.e., this property holds in both cases. Consider graphs $\hat{G}_1$ and $\hat{G}_2$ obtained from $\hat{G}$ by preforming the switches σ, σ' and σ', σ. By the definition of switches, $\hat{G}_1$ and $\hat{G}_2$ are φ-isomorphic, and hence the sequence $\mathcal{S}''_E$ obtained from $\mathcal{S}'_E$ by swapping σ and σ' is an H-sequence. However, this contradicts the choice of $\mathcal{S}'_H$, because a t-switch occurs $\mathcal{S}''_E$ earlier than in $\mathcal{S}'_E$. This proves that the first switch in $\mathcal{S}'_E$ is a t-switch.

Denote by $\hat{G}$ the graph obtained from G by applying the first switch in $\mathcal{S}'_E$ and let $\mathcal{S}''_E$ be the sequence obtained from $\mathcal{S}'_E$ by deleting the first switch. We have that $\mathcal{S}''_E$ is an H-sequence for $\hat{G}$. If $\mathcal{S}''_E$ contains no t-switch, we set $\hat{\mathcal{S}}_E = \mathcal{S}'_E$ and this concludes the proof. Otherwise, we apply the inductive assumption using the fact that the length of $\mathcal{S}''_E$ is lesser that the length of $\mathcal{S}_E$. Then we can replace $\mathcal{S}''_E$ in $\mathcal{S}'_E$ by a minimum t-sorted H-sequence for $\hat{G}$ such that the t-switches in it are the first switches in the sequence. The obtained H sequence $\hat{\mathcal{S}}_E$ is a minimum t-sorted H-sequence $\hat{\mathcal{S}}_E$ with exactly the same switches as in $\mathcal{S}_E$. This concludes the proof. $\quad\square$

It is convenient to note the following property of minimum H-sequences.

LEMMA 12. *For every $t \in W_2^{(1)}$, any minimum H-sequence contains at most one t-switch.*

Proof. The proof is by contradiction. Assume that there is a minimum H-sequence $\mathcal{S}_E$ with at least two t-switches for some $t \in W_2^{(1)}$. By Lemma 11, we can assume that $\mathcal{S}_E$ is t-sorted. Then the first two switches σ and σ' are t-switches. Let G' be the graph obtained from G by applying σ and σ'. Let σ and σ' be switches with respect to edge Whitney separations (L, R) and (L', R'), respectively. We have that $V(L) \cap V(R) = V(L') \cap V(R') = V(e)$ for the unique edge e of $G[X_t^{(1)}]$, i.e, σ and σ' use the same separator $V(e) \subseteq V(G)$.

Let $\hat{L} = (L \cap L') \cup (R \cap R')$ and $\hat{R} = E(G) \setminus \hat{L}$. Notice that for every connected component C of $G - V(e)$, either $V(C) \subseteq V(L) \setminus V(e)$ or $V(C) \subseteq V(R) \setminus V(e)$ and either $V(C) \subseteq V(L') \setminus V(e)$ or $V(C) \subseteq V(R') \setminus V(e)$. Then by the definition of Whitney switches, we conclude that the graph obtained from G by the Whitney switch with respect to $(\hat{L}, \hat{R})$ is φ-isomorphic to G'. However, this contradicts the minimality of $\mathcal{S}_E$ as we have that σ and σ' can be replaced by a single switch. $\quad\square$

We show that we can further restrict the set of considered Whitney switches.

For $t \in W_{\geq 3}^{(1)}$, we say that $X_t^{(1)}$ is φ-*good* if $G[X_t^{(1)}]$ is φ-isomorphic to $H[X_{\alpha(t)}^{(1)}]$, and $X_t^{(1)}$ is φ-*bad* otherwise. Notice that if $G[X_t^{(1)}]$ is 3-connected, then $X_t^{(1)}$ is φ-good but this not always so if $G[X_t^{(1)}]$ is a cycle.

Let $t \in W_{\geq 3}^{(1)}$ such that $X_t^{(1)}$ is φ-bad. Clearly, $G[X_t^{(1)}]$ is a cycle with at least four vertices, because only such bags may be φ-bad. Let $\{t_1, \ldots, t_s\} = N_{T^{(1)}}^2(t)$ and denote $G_t = G[X_t^{(1)} \cup \bigcup_{i=1}^s X_{t_i}^{(1)}]$ and $H_{\alpha(t)} = H[X_{\alpha(t)}^{(2)} \cup \bigcup_{i=1}^s X_{\alpha(t_i)}^{(2)}]$. Let $P = v_0 \cdots v_r$ be a path in $G[X_t^{(1)}]$ and $e_i = v_{i-1}v_i$ for $i \in \{1, \ldots, r\}$. We say that P is a φ-*good segment* of $X_t^{(1)}$ if the following holds (see Figure 4 for an example):
  (i) the length of P is at least 5,
  (ii) there is a path $P' = u_0 \cdots u_r$ in $H[X_{\alpha(t)}^{(2)}]$ such that $u_{i-1}u_i = \varphi(e_i)$ for all $i \in \{1, \ldots, r\}$,
  (iii) for every $i \in \{1, \ldots, r\}$ and for every $t' \in W_{\geq 3}^{(1)}$ such that $X_t^{(1)} \cap X_{t'}^{(1)} = \{v_{i-1}, v_i\}$, $X_{t'}^{(1)}$ is φ-good,
  (iv) for every $i \in \{1, \ldots, r-1\}$, $\varphi(E_{G_t}(v_i)) = E_{H_{\alpha(t)}}(u_i)$.
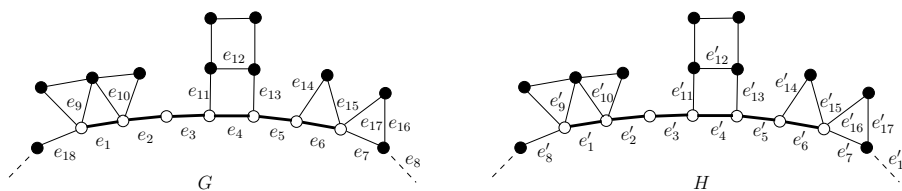
FIG. 4. *An example of a φ-good segment; $\varphi(e_i) = e'_i$ for $i \in \{1, \ldots, 18\}$, and the vertices of the segment are white.*
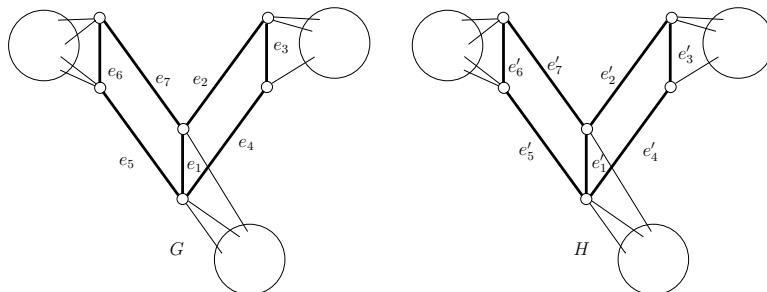

FIG. 5. *Mutually φ-good bags; $\varphi(e_i) = e'_i$ for $i \in \{1, \ldots, 7\}$, and the vertices of the mutually φ-good bags of G and the corresponding bags of H are white.*

For distinct $t_1, t_2 \in W_{\geq 3}^{(1)}$ with a common neighbor in $T^{(1)}$, we say that $X_{t_1}^{(1)}$ and $X_{t_2}^{(1)}$ are *mutually φ-good* (see Figure 5) if they are φ-good and $G[X_{t_1}^{(1)} \cup X_{t_2}^{(1)}]$ is φ-isomorphic to $H[X_{\alpha(t_1)}^{(2)} \cup X_{\alpha(t_2)}^{(2)}]$.

We say that an H-sequence $\mathcal{S}_E$ is *φ-good* if no Whitney switch of $\mathcal{S}_E$ splits (mutually) φ-good bags and segments. Formally,

  (i) for every switch $\sigma \in \mathcal{S}_E$, σ does not split $E(G[X_t^{(1)}])$ for for every φ-good bag $X_t^{(1)}$,

  (ii) for every switch $\sigma \in \mathcal{S}_E$, σ does not split $E(P)$ for every φ-good segment P,

  (iii) for every switch $\sigma \in \mathcal{S}_E$, σ does not split $E(G[X_{t_1}^{(1)} \cup X_{t_2}^{(1)}])$ for every two distinct mutually φ-good bags $X_{t_1}^{(1)}$ and $X_{t_2}^{(1)}$.

An H-sequence $\mathcal{S}$ is *φ-good* if $\mathcal{S}_E$ is φ-good.

We prove that it is sufficient to consider φ-good H-sequences.

LEMMA 13. *There is a minimum sorted H-sequence of Whitney switches $\mathcal{S}_E$ that is φ-good.*

Proof. First, we show that if $\mathcal{S}_E$ is a minimum sorted H-sequence of Whitney switches such that the number of switches that split φ-good bags is minimum, then condition (i) is fulfilled, that is, no switch in $\mathcal{S}_E$ splits a φ-good bag $X_t^{(1)}$.

The proof is by induction on the length of $\mathcal{S}_E$. For the base case, assume that $\mathcal{S}_E$ contains only t-switches for some $t \in V(T^{(1)})$. If $t \in W_2^{(1)}$, then the switches of $\mathcal{S}_E$ do not split any bag $X_{t'}^{(1)}$ for $t' \in V(T^{(1)})$ and the claim holds. Suppose that $t \in W_{\geq 3}^{(1)}$. Then the switches of $\mathcal{S}_E$ split $X_t^{(1)}$ and only this bag may be split by them. If $X_t^{(1)}$ is φ-bad, then the claim is fulfilled. Assume that $X_t^{(1)}$ is φ-good. If $G[X_t^{(1)}]$ is 3-connected or a cycle with 3 vertices, then $X_t^{(1)}$ cannot be split. Then $G'[X_t^{(1)}]$ is a cycle of length at least 4.

Denote by $v_1, \ldots, v_r$ the vertices of the cycle $G[X_t^{(1)}]$ (in the cycle order) and let $e_i = v_{i-1} v_i$ for $i \in \{1, \ldots, r\}$ assuming that $v_0 = v_r$ (i.e., the indices are taken modulo r). Notice that for each neighbor t' of t in $T^{(1)}$, $t' \in W_2^{(1)}$ and $X_{t'}^{(1)} = \{v_{i-1}, v_i\}$ for some $i \in \{1, \ldots, r\}$. Assume that $N_T(t) = \{t_1, \ldots, t_s\}$, where $X_{t_i}^{(1)} = \{v_{j_i-1}, v_{j_i}\}$ for $1 \leq j_1 < \ldots < j_s \leq r$. Denote by $T_1^{(1)}, \ldots, T_s^{(1)}$ the subtrees of $T^{(1)} - t$ containing $t_1, \ldots, t_s$, respectively, and let G_i be the subgraph of G induced by the vertices of $\bigcup_{h \in V(T_i^{(1)})} X_h^{(1)}$ for $i \in \{1, \ldots, s\}$. For $i \in \{1, \ldots, s\}$, let $T_1^{(2)}, \ldots, T_s^{(2)}$ be the subtrees of $T^{(2)} - \alpha(t)$ that contain $\alpha(t_1), \ldots, \alpha(t_s)$, respectively, and let H_i be the subgraph of H induced by the vertices of $\bigcup_{h \in V(T_i^{(2)})} X_h^{(2)}$ for $i \in \{1, \ldots, s\}$.

Since $X_t^{(1)}$ is φ-good, $\varphi(e_1), \ldots, \varphi(e_r)$ form a cycle of H in the given order. Assume that $\varphi(e_i) = u_{i-1} u_i$ for $i \in \{1, \ldots, r\}$ for $u_1, \ldots, u_r$ forming $X_{\alpha(t)}^{(2)}$ (assuming that $u_0 = u_r$). Observe that the graph φ-isomorphic to H is obtained from G by Whitney switches with respect to edge Whitney separations (L, R) splitting $X_t^{(1)}$, that is, $V(L) \cap V(R) \subseteq X_t^{(1)}$. This implies that G_i is φ-isomorphic to H_i for every $i \in \{1, \ldots, s\}$ as the switches do not affect these graphs. However, G and H are not φ-isomorphic by the minimality of $\mathcal{S}$. By Lemma 7, we obtain that there are $i \in \{1, \ldots, r\}$ such that $\varphi(E_G(v_i)) \neq E_H(u_i)$. More precisely, taking into account that every G_i is φ-isomorphic to H_i, we have that there is $i \in \{1, \ldots, s\}$ such that $\varphi(E_{G_i}(v_{j_i})) = E_{H_i}(u_{j_i-1})$ and $\varphi(E_{G_i}(v_{j_i-1})) = E_{H_i}(u_{j_i})$. Denote by $I \subseteq \{1, \ldots, s\}$ the set of all such indices $i \in \{1, \ldots, s\}$.

We define the partially signed circular permutation $\overrightarrow{\pi}^c = (\langle 1, s_1 \rangle, \ldots, \langle r, s_r \rangle)$ such that $s_{j_i} = -1$ for all $i \in I$, $s_{j_i} = +1$ for all $i \in \{1, \ldots, s\} \setminus I$, and $s_j = 0$ for $j \in \{1, \ldots, r\} \setminus \{j_1, \ldots, j_s\}$. The crucial observation is that obtaining the graph φ-isomorphic to H from G by Whitney switches is equivalent to sorting $\overrightarrow{\pi}^c$ by reversals. By Lemma 6, there is an optimal sorting sequence composed by trivial reversals $\overrightarrow{\rho}^c(j, j)$ for $s_j = -1$. This corresponds to performing the Whitney switches with respect to edge Whitney separations $(E(G_i), (E(G) \setminus E(G_i)) \cup \{e_i\})$ for all $i \in I$. This contradicts the choice of $\mathcal{S}$, because these switches do not split $X_t^{(1)}$. This completes the proof for the base case.

Suppose that $\mathcal{S}_E$ contains t-switches for at least two distinct $t \in V(T^{(1)})$. By Lemma 11, we can assume that $\mathcal{S}_E$ is t-sorted for some $t \in V(T^{(1)})$. Denote by $\mathcal{S}_E'$ the inclusion maximal subsequence of t-switches in $\mathcal{S}_E$. Let $\mathcal{S}_E''$ be the subsequence of the switches that are after the switches in $\mathcal{S}_E$ and let G' be the graph obtained from G by applying $\mathcal{S}_E'$. Note that $\mathcal{S}_E''$ is nonempty. Since G' is obtained from G by a minimum sorted G'-sequence of Whitney switches with minimum number of switches splitting φ-good bags, we obtain that the switches in $\mathcal{S}_E'$ do not split φ-good bags by the proved base case. Then we apply the inductive assumption for $\mathcal{S}_E''$ and G', because $\mathcal{S}_E''$ is a minimum H-sequence for G'. Then the switches $\mathcal{S}$ do not split φ-good bags.

By the next step, we show that there is a minimum sorted H-sequence of Whitney switches $\mathcal{S}_E$ such that conditions (i) and (ii) of the definition of φ-good sequences are fulfilled, that is, the switches of $\mathcal{S}_E$ do not split φ-good bags and φ-good segments. The proof is similar to the first part. Suppose that $\mathcal{S}_E$ is a minimum H-sequence of Whitney switches that satisfies (i) such that the number of switches splitting φ-good segments is minimum. We claim that $\mathcal{S}$ satisfies (ii).

The proof is by induction on the length of $\mathcal{S}_E$. As in the previous step, it is crucial to deal with a base case. Assume that $\mathcal{S}_E$ contains only t-switches for some $t \in W_{\geq 3}^{(1)}$, where the bag $X_t^{(1)}$ contains a φ-good segment and t'-switches for $t' \in N_{T^{(1)}}(t)$ such

that the unique edge of $E(X_t^{(1)})$ is in a φ-good segment of $X_t^{(1)}$. Then $G[X_t^{(1)}]$ is a cycle and any switch of $\mathcal{S}_E$ can only split φ-good segments in $G[X_t^{(1)}]$. Denote by $P_1, \ldots, P_\ell$ the family of inclusion maximal φ-good segments in $G[X_t^{(1)}]$. If no switch of $\mathcal{S}_E$ splits these φ-good segments, then the claim holds. Assume that this is not the case.

Denote by $v_1, \ldots, v_r$ the vertices of the cycle $G[X_t^{(1)}]$ (in the cycle order) and let $e_i = v_{i-1}v_i$ for $i \in \{1, \ldots, r\}$ assuming that $v_0 = v_r$ (i.e., the indices are taken modulo r). Notice that for each neighbor t' of t in $T^{(1)}$, $t' \in W_2^{(1)}$ and $X_{t'}(1) = \{v_{i-1}, v_i\}$ for some $i \in \{1, \ldots, r\}$. Assume that $N_{T^{(1)}}(t) = \{t_1, \ldots, t_s\}$, where $X_{t_i}^{(1)} = \{v_{j_i-1}, v_{j_i}\}$ for $1 \le j_1 < \ldots < j_s \le r$. Denote by $T_1^{(1)}, \ldots, T_s^{(1)}$ the subtrees of $T^{(1)} - t$ containing $t_1, \ldots, t_s$, respectively, and let G_i be the subgraph of G induced by the vertices of $\bigcup_{h \in V(T_i^{(1)})} X_h^{(1)}$ for $i \in \{1, \ldots, s\}$. For $i \in \{1, \ldots, s\}$, let $T_1^{(2)}, \ldots, T_s^{(2)}$ be the subtrees of $T^{(2)} - \alpha(t)$ that contain $\alpha(t_1), \ldots, \alpha(t_s)$, respectively, and let H_i be the subgraph of H induced by the vertices of $\bigcup_{h \in V(T_i^{(2)})} X_h^{(2)}$ for $i \in \{1, \ldots, s\}$.

Notice that G_i is φ-isomorphic to H_i for every $i \in \{1, \ldots, s\}$. If e_{j_i} is an edge of one of the paths $P_1, \ldots, P_\ell$, then this follows from conditions (iii) and (iv) of the definition of φ-good segments. Otherwise, because $\mathcal{S}_E$ does not contain a t'-switch for $X_{t'}^{(1)} = V(e_{j_i})$, G_i is not affected by switches in $\mathcal{S}$.

Denote by $e_1', \ldots, e_r'$ the edges of $H[X_{\alpha(t)}^{(2)}]$ taken in the cycle order and denote by $u_1, \ldots, u_r$ the vertices of this cycle such that $e_i' = u_{i-1}u_i$ (assuming that $u_0 = u_r$). For $i \in \{1, \ldots, s\}$, let $\varphi(e_{j_i}) = e_{j_i'}$ for $j_1', \ldots, j_s' \in \{1, \ldots, r\}$. Since each G_i is φ-isomorphic to H_i, we have that, by Lemma 7, for every $i \in \{1, \ldots, s\}$, either $\varphi(E_{G_i}(v_{j_i-1})) = E_{G_i}(u_{j_i'-1})$ and $\varphi(E_{G_i}(v_{j_i})) = E_{G_i}(u_{j_i'})$ or, symmetrically, $\varphi(E_{G_i}(v_{j_i-1})) = E_{G_i}(u_{j_i'})$ and $\varphi(E_{G_i}(v_{j_i})) = E_{G_i}(u_{j_i'-1})$. Let $I = \{i \mid 1 \le i \le s, \varphi(E_{G_i}(v_{j_i-1})) = E_{G_i}(u_{j_i'-1})$ and $\varphi(E_{G_i}(v_{j_i})) = E_{G_i}(u_{j_i'})\}$, and let $\bar{I} = \{1, \ldots, r\} \setminus I$.

We construct the following partially signed circular permutation $\overrightarrow{\pi}^c = (\langle \pi_1, s_1 \rangle, \ldots, \langle \pi_r, s_r \rangle)$ such that for every $i \in \{1, \ldots, r\}$, $e_{\pi_i} = \varphi^{-1}(e_i')$. For $i \in \{1, \ldots, s\}$, we define $s_{j_i} = +1$ if $i \in I$ and $s_{j_i} = -1$ if $i \in \bar{I}$. The other sign is zeros, that is, $s_j = 0$ if $j \notin \{i_1, \ldots, i_s\}$. Notice that by the definition of φ-good segments, $(\langle \pi_i, s_i \rangle, \ldots, \langle \pi_j, s_j \rangle)$ (with indices taken modulo r) is a signed block of $\overrightarrow{\pi}^c$ of length at least 5 if and only if the edges $e_{\pi_1}, \ldots, e_{\pi_j}$ form a φ-good segment.

Similarly to the first part of the proof, we have that obtaining the graph φ-isomorphic to H from G by Whitney switches is equivalent to sorting $\overrightarrow{\pi}^c$ by reversals. By Lemma 5, there is an optimal sorting sequence that does not cut strips of $\overrightarrow{\pi}^c$ of length at least 5. This implies that there is a sequence of Whitney switches $\mathcal{S}_E'$ of the same length as $\mathcal{S}_E$ such that applying $\mathcal{S}_E'$ to G creates a graph G' isomorphic to H and no switch of $\mathcal{S}_E'$ splits $P_1, \ldots, P_\ell$. This contradicts the choice of $\mathcal{S}_E$ and, therefore, proves the claim for the base case.

Now we consider the inductive step. If there is no bag $X_t^{(1)}$ with a φ-good segment, then the claim holds trivially. Assume that there is $t \in W_{\ge 3}^{(1)}$ such that $X_t^{(1)}$ has a φ-good segment. Let Z be the set of t'-switches of $\mathcal{S}_E$ for $t' \in \bar{N}_{T^{(1)}}(t)$ such that the unique edge of $G[X_t^{(1)}]$ is in a φ-good segment of $X_t^{(1)}$. By the iterative application of Lemma 11, we can assume that the t'-switches for $t' \in \{t\} \cup Z$ are the first switches in $\mathcal{S}_E$. Denote by $\mathcal{S}_E'$ the subsequence of these t'-switches in $\mathcal{S}_E$. Let $\mathcal{S}_E''$ be the subsequence of the switches that are after the switches in $\mathcal{S}_E'$ and let G' be the graph obtained from G by applying $\mathcal{S}_E'$. If $\mathcal{S}_E''$ is empty, then the claim holds as we have the base case. Assume that $\mathcal{S}_E''$ is nonempty. Since G' is obtained from G by a minimum sorted

G'-sequence of Whitney switches with a minimum number of switches splitting φ-good segments, we obtain that the switches in $\mathcal{S}'_E$ do not split φ-good segments by the proved base case. Then we apply the inductive assumption for $\mathcal{S}''_E$ and G', because $\mathcal{S}''_E$ is a minimum H-sequence for G'. Then the switches $\mathcal{S}$ do not split φ-good segments.

Finally, we show that every minimum H-sequence of Whitney switches $\mathcal{S}_E$ satisfying conditions (i) and (ii) of the definition of φ-good sequences satisfies (iii) as well, that is, $\mathcal{S}_E$ does not split mutually φ-good bags. The proof is by contradiction. Assume that $\mathcal{S}_E$ contains a switch that splits some mutually φ-good bags $X_{t_1}^{(1)}$ and $X_{t_2}^{(1)}$ bags. Because $\mathcal{S}_E$ does not split φ-good bags, a switch splitting $X_{t_1}^{(1)}$ and $X_{t_2}^{(1)}$ is a t-switch for the common neighbor $t \in W_2^{(1)}$ of t_1 and t_2 in T_1. By Lemma 11, we can assume that $\mathcal{S}_E$ is t-sorted. By Lemma 12, we have that $\mathcal{S}_E$ contains a unique t-switch. Denote by G' the graph obtained from G by applying this switch. Because $X_{t_1}^{(1)}$ and $X_{t_2}^{(1)}$ mutually φ-good bags, $G[X_{t_1}^{(1)} \cup X_{t_2}^{(1)}]$ is φ-isomorphic to $H[X_{\alpha(t_1)}^{(2)} \cup X_{\alpha(t_2)}^{(2)}]$. This implies that $G'[X_{t_1}^{(1)} \cup X_{t_2}^{(1)}]$ is not φ-isomorphic to $H[X_{\alpha(t_1)}^{(2)} \cup X_{\alpha(t_2)}^{(2)}]$. Because $\mathcal{S}_E$ does not have switches splitting $X_{t_1}^{(1)}$ and $X_{t_2}^{(1)}$ and other t-switches except the first, the subgraph induced by $V(E(G'[X_{t_1}^{(1)} \cup X_{t_2}^{(1)}]))$ in the graph obtained by applying $\mathcal{S}_E$ remains φ-isomorphic to $G'[X_{t_1}^{(1)} \cup X_{t_2}^{(1)}]$ and, therefore, not φ-isomorphic to $H[X_{\alpha(t_1)}^{(2)} \cup X_{\alpha(t_2)}^{(2)}]$. However, this contradicts that $\mathcal{S}_E$ is an H-sequence. This proves the claim and completes the proof of the lemma. □

Let $t \in W_{\geq 3}^{(1)}$ be such that $X_t^{(1)}$ is φ-bad. Denote by $t_1, \ldots, t_s \neq t$ the nodes of $N_{T^{(1)}}^2(t)$. Let $G_t = G[X_t^{(1)} \cup \bigcup_{i=1}^s X_{t_i}^{(1)}]$ and $H_{\alpha(t)} = G[X_{\alpha(t)}^{(2)} \cup \bigcup_{i=1}^s X_{\alpha(t_i)}^{(2)}]$. In other words, G_t is the subgraph of G induced by the vertices of $X_t^{(1)}$ and the vertices of the bags at distance two in $T^{(1)}$ from t, and $H_{\alpha(t)}$ the subgraph of H induced by the vertices of the bags that are images of the bags composing G_t according to α.

We say that a vertex $v \in X_t^{(1)}$ is a *crucial breakpoint* if $\varphi(E_{G_t}(v)) \neq E_{H_{\alpha(t)}}(u)$ for every $u \in V(H_{\alpha(t)})$. We denote by $b(G)$ the total number of crucial breakpoints in the φ-bad bags and say that $b(G)$ is the *breakpoint number* of G. Recall that by our convention, G and H are enhanced graphs, but we extend this definition for the general case needed in the next section. For (not necessarily enhanced) 2-isomorphic graphs G and H, and a 2-isomorphism φ, we construct their enhancements $\widehat{G}$ and $\widehat{H}$ and consider the enhanced mapping $\widehat{\varphi}$. Then $b(G)$ is defined as $b(\widehat{G})$.

Observe that if G and H are φ-isomorphic, then $b(G) = 0$ by Lemma 7. However, this is not true the other way around, because in the definition of $b(G)$ we count only breakpoints in φ-bad bags. In particular, if G has not φ-bad bags, then $b(G) = 0$ but this does not mean that G and H are φ-isomorphic.

We conclude the section by giving a lower bound for the length of an H-sequence.

LEMMA 14. *Let $\mathcal{S}$ be an H-sequence of Whitney switches. Then $b(G)/2 \leq |\mathcal{S}|$.*

Proof. The claim is trivial if $b(G) = 0$. Assume that $b(G) > 0$. Let $\mathcal{S}$ be an H-sequence of Whitney switches, that is, the graph G' obtained from G by applying $\mathcal{S}$ is φ-isomorphic to H. By Lemma 7, $b(G') = 0$. Hence, $\mathcal{S}$ should contain switches that decrease the breakpoint number. Note that the Whitney switch with respect to a Whitney partition (A, B) reduces $b(G)$ if and only if $A \cap B = \{u, v\}$, where at least one of u or v is a crucial breakpoint. Then the switch decreases $b(G)$ by at most 2 and the claim follows. □

5. Kernelization for Whitney Switches. In this section, we show that Whitney Switches parameterized by k admits a polynomial kernel. To do it, we obtain a more general result by proving that the problem has a polynomial kernel when parameterized by the breakpoint number of the first input graph.

Theorem 4. Whitney Switches *has a kernel such that each graph in the obtained instance has at most* $\min\{39 \cdot b - 27, 3\}$ *vertices, where b is the breakpoint number of the input graph.*

Proof. Let (G, H, φ, k) be an instance of Whitney Switches, where G and H are n-vertex 2-connected 2-isomorphic graphs, $\varphi \colon E(G) \to E(H)$ is a 2-isomorphism, and k is a nonnegative integer.

First, we use Proposition 1 to construct the Tutte decompositions of G and H. Denote by $\mathcal{T}^{(1)} = (T^{(1)}, \{X_t^{(1)}\}_{t \in V(T^{(1)})})$ and $\mathcal{T}^{(2)} = (T^{(2)}, \{X_t^{(2)}\}_{t \in V(T^{(2)})})$ the constructed Tutte decompositions of G and H, respectively, and let $(W_2^{(h)}, W_{\geq 3}^{(h)})$ be the partition of $V(T^{(h)})$ satisfying (T4)–(T8) for $h = 1, 2$.

In the next step, we construct the isomorphism $\alpha \colon V(T^{(1)}) \to V(T^{(2)})$ satisfying conditions (i)–(iii) of Lemma 8. Recall that Lemma 8 claims that such an isomorphism always exists. If $T^{(1)}$ and $T^{(2)}$ are single-vertex trees, then the construction is trivial. Assume that this is not the case. Let t be a leaf of $T^{(1)}$ and let t' be its unique neighbor. We have that $t \in W_{\geq 3}^{(1)}$ and $E(G[X_t^{(1)}]) \setminus E(G[X_{t'}^{(1)}]) \neq \emptyset$. By (iii), we have that there is a unique leaf t'' of $T^{(2)}$ such that $\varphi(E(G[X_t^{(1)}]) \setminus E(G[X_{t'}^{(1)}])) \subseteq E(H[X_{t''}^{(2)}])$ and $\alpha(t) = t''$. Observe that t'' can be found in polynomial time. This means that we can construct in polynomial time the restriction of α on the leaves of $T^{(1)}$ that maps them bijectively on the leaves of $T^{(2)}$. Since $T^{(1)}$ and $T^{(2)}$ are isomorphic, there is a unique way to extend α from leaves to $V(T^{(1)})$. This can be done by picking a *root* node r of $T^{(1)}$ and computing α bottom-up starting from the leaves. Given that α is already computed for the leaves, the construction of α can be completed in $\mathcal{O}(|V(T^{(1)})|)$ time.

Given α, we compute the enhancements $\widehat{G}$ and $\widehat{H}$ of G and H, respectively, and then define the enhanced mapping $\widehat{\varphi} \colon E(\widehat{G}) \to E(\widehat{H})$. This can be done in polynomial time. Note that α satisfies the conditions of Lemma 10. Observe also that we can verify in polynomial time whether a bag $X_t^{(1)}$ for $t \in W_{\geq 3}^{(1)}$ is φ-good or not. Then we can compute in polynomial time $b(G) = b(\widehat{G})$.

To simplify notation, let $G := \widehat{G}$, $H := \widehat{H}$, and $\varphi := \hat{\varphi}$.

Now we apply a series of reduction rules that are applied for G, H, φ, and the Tutte decompositions of G and H.

The aim of the first rule is to decrease the total size of bags that are φ-bad (see Figure 6 for an example).

Reduction Rule 1. *If for $t \in W_{\geq 3}^{(1)}$ such that $X_t^{(1)}$ is φ-bad there is an inclusion maximal φ-good segment $P = v_0 \cdots v_r$, then do the following:*
- *find the path $P' = u_0 \cdots u_r$ in $H[X_{\alpha(t)}^{(2)}]$ composed by the edges $u_{i-1}u_i = \varphi(v_{i-1}v_i)$ for $i \in \{1, \ldots, r\}$,*
- *add the edge $v_0 v_r$ to G and $u_0 u_r$ to H,*
- *extend φ by setting $\varphi(v_0 u_r) = u_0 u_r$,*
- *recompute the Tutte decompositions of the obtained graphs and the isomorphism α.*

Claim 1. *Reduction Rule 1 is safe, does not increase the breakpoint number, and can be executed in polynomial time.*
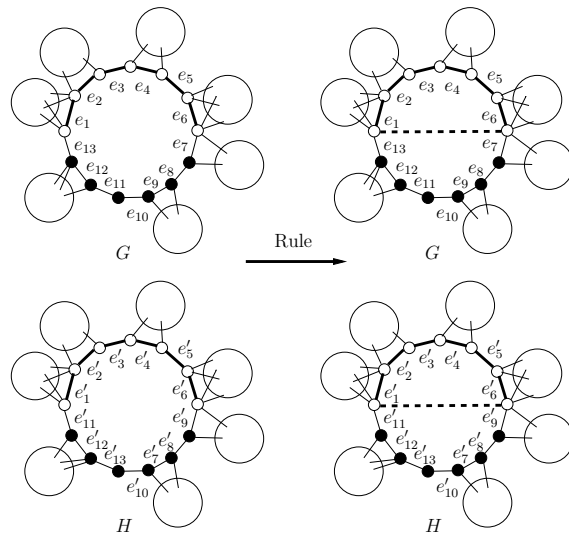
FIG. 6. *An example of an application of Reduction Rule 1; $\varphi(e_i) = e_i'$ for $i \in \{1, \ldots, 13\}$, the vertices of the φ-good segment in G and the corresponding segment in H are white, and the added edges are shown by dashed lines.*

Proof of Claim 1. Denote by $\tilde{G}$ the graph obtained from G by the application of Reduction Rule 1 for $P = v_0 \cdots v_r$. Let also $\tilde{H}$ be the graph obtained from H and denote by $\tilde{\varphi}$ in the extension of φ. Since φ maps the edges of P into the edges of P', we have that $\tilde{\varphi}$ is a 2-isomorphism of $\tilde{G}$ to $\tilde{H}$.

Suppose that (G, H, φ, k) is a yes-instance of WHITNEY SWITCHES. By Lemma 13, there is a φ-good H-sequence $\mathcal{S}_E$ of Whitney switches of length at most k that transforms G into the graph G' that is φ-isomorphic to H. By condition (ii) of the definition of a φ-good H-sequence, no Whitney switch in the sequence splits $E(P)$. This implies that $\mathcal{S}_E$ can be performed on $\tilde{G}$ and transforms $\tilde{G}$ into $\tilde{G}'$ that is $\tilde{\varphi}$-isomorphic to $\tilde{H}$. This means that $(\tilde{G}, \tilde{H}, \tilde{\varphi}, k)$ is a yes-instance.

It is straightforward to see that every sequence of Whitney switches transforming $\tilde{G}$ into a graph $\tilde{\varphi}$-isomorphic to $\tilde{H}$ can be applied to G and produces the graph φ-isomorphic to H. Therefore, if $(\tilde{G}, \tilde{H}, \tilde{\varphi}, k)$ is a yes-instance of WHITNEY SWITCHES, then (G, H, φ, k) is a yes-instance as well.

To show that $b(\tilde{G}) = b(G)$, we explain how to recompute the Tutte decompositions. For this, observe that we add a chord to a cycle of G that forms a bag of the Tutte decomposition. This operation splits the bag into two bags of size at least 3 and the bag of size 2 composed by the end-vertices of the chord. Formally, this is done as follows. We replace t in $T^{(1)}$ by three nodes t_1, t_2, and t' and define the corresponding bags $X_{t_1}^{(1)} = \{v_0, \ldots, v_r\}$, $X_{t_2}^{(1)} = X_t^{(1)} \setminus \{v_1, \ldots, v_{r-1}\}$, and $X_{t'}^{(1)} = \{v_0, v_r\}$. Notice that for every $t'' \in N_{V(T^{(1)})}(t)$, either $X_{t''}^{(1)} \subseteq X_{t_1}^{(1)}$ or $X_{t''}^{(1)} \subseteq X_{t_2}^{(1)}$. In the first case, we make t'' adjacent to t_1 and t'' is adjacent to t_2 in the second case. We modify $T^{(2)}$ and redefine α in similar way. The node $\alpha(t)$ is replaced by three nodes $\alpha(t_1)$, $\alpha(t_2)$, and $\alpha(t')$ with $X_{\alpha(t_1)}^{(2)} = \{u_0, \ldots, u_r\}$, $X_{\alpha(t_2)}^{(2)} = X_{\alpha(t)}^{(2)} \setminus \{u_1, \ldots, u_{r-1}\}$, and $X_{\alpha(t')}^{(2)} = \{u_0, u_r\}$. For every $t'' \in N_{V(T^{(2)})}(\alpha(t))$, either $X_{t''}^{(2)} \subseteq X_{\alpha(t_1)}^{(2)}$ or $X_{t''}^{(2)} \subseteq X_{\alpha(t_2)}^{(2)}$. We make t'' adjacent to $\alpha(t_1)$ in the first case

and t'' is adjacent to $\alpha(t_2)$ in the second case. It is straightforward to verify that we obtain the Tutte decompositions of $\tilde{G}$ and $\tilde{H}$, respectively, and the obtained α is an isomorphism of the modified tree $T^{(1)}$ to the modified tree $T^{(2)}$ satisfying the conditions of Lemma 10. Notice that the vertices of $X_{t'}^{(1)}$ are adjacent and the same holds for $X_{\alpha(t')}^{(2)}$, that is, $\tilde{G}$ and $\tilde{H}$ are enhanced.

Thus, we obtain two bags $X_{t_1}^{(1)}$ and $X_{t_2}^{(1)}$ of size at least 3 from $X_t^{(1)}$ and both of them induce cycles. Since P is a φ-good segment and $\tilde{\varphi}(v_0 v_r) = u_0 u_r$, we have that $X_{t_1}^{(1)}$ is $\tilde{\varphi}$-good. Moreover, for every $i \in \{0, \ldots, r\}$, $\tilde{\varphi}(E_{\tilde{G}[X_{t_1}^{(1)}]}(v_i)) = E_{\tilde{H}[X_{\alpha(t_1)}^{(2)}]}(u_i)$. This implies that the number of crucial breakpoints does not increase.

To argue that Reduction Rule 1 can be applied in polynomial time, observe first that inclusion maximal φ-good segments can be recognized in polynomial time. For each $t \in W_{\geq 3}^{(1)}$, we can verify whether $X_t^{(1)}$ is φ-good in polynomial time using Lemma 7. Then for each $t \in W_{\geq 2}^{(1)}$ such that $X_t^{(1)}$ is a φ-bad bag, we consider all at most n^2 paths P of the cycle $G[X_t^{(1)}]$ and for each P, we verify conditions (i)–(iv) of the definition of a φ-good segment. It is easy to see that each of these conditions can be verified in polynomial time. Further, given an inclusion maximal φ-good segment P, we can apply the rule in polynomial time. Note also that we can avoid recomputing the Tutte decompositions of $\tilde{G}$ of $\tilde{H}$ from scratch as the described-above computation procedure can be done in polynomial time. $\square$

Reduction Rule 1 is applied exhaustively while we are able to find φ-good segments. To simplify notation, we use G, H, and φ to denote the obtained graphs and the obtained 2-isomorphism. We also keep the notation used for the Tutte decompositions.

Our next reduction rule is used to simplify the structure of φ-good bags by turning them into cliques (see Figure 7 for an example).

REDUCTION RULE 2. *If for $t \in W_{\geq 3}^{(1)}$ such that $X_t^{(1)}$ is a φ-good there are non-adjacent vertices in $X_t^{(1)}$, then compute the φ-isomorphism ψ of $G[X_t^{(1)}]$ to $H[X_{\alpha(t)}^{(2)}]$ and for every nonadjacent $u, v \in X_t^{(1)}$, do the following:*
- *add the edge uv to G and $\psi(u)\psi(v)$ to H,*
- *extend φ by setting $\varphi(uv) = \psi(u)\psi(v)$.*

CLAIM 2. *Reduction Rule 2 is safe and changes neither the Tutte decomposition nor the breakpoint member. Furthermore, it can be executed in polynomial time.*

Proof of Claim 2. Let $t \in W_{\geq 3}^{(1)}$ be such that $X_t^{(1)}$ is φ-good and there are non-adjacent vertices in $X_t^{(1)}$. Recall that $G[X_t^{(1)}]$ and $H[X_{\alpha(t)}^{(2)}]$ are φ-isomorphic by the definition of φ-good bags. Therefore, there is a φ-isomorphism ψ of $G[X_t^{(1)}]$ to $H[X_{\alpha(t)}^{(2)}]$.

Denote by $\tilde{G}$ the graph obtained from G by the application of one step of Reduction Rule 2 for two nonadjacent $u, v \in X_t^{(1)}$, that is, $\tilde{(G)}$ is obtained by adding uv to G. Let $\tilde{H}$ be the graph obtained from H by adding $\psi(u)\psi(v)$, and let $\tilde{\varphi}$ be the extension of φ on uv. Since ψ is a φ-isomorphism, we conclude that $\tilde{\varphi}$ is a 2-isomorphism of $\tilde{G}$ to $\tilde{H}$.

Suppose that (G, H, φ, k) is a yes-instance of WHITNEY SWITCHES. By Lemma 13, there is a φ-good H-sequence of Whitney switches that transforms G into the graph G' that is φ-isomorphic to H. By condition (i) of the definition of a φ-good H-sequence,
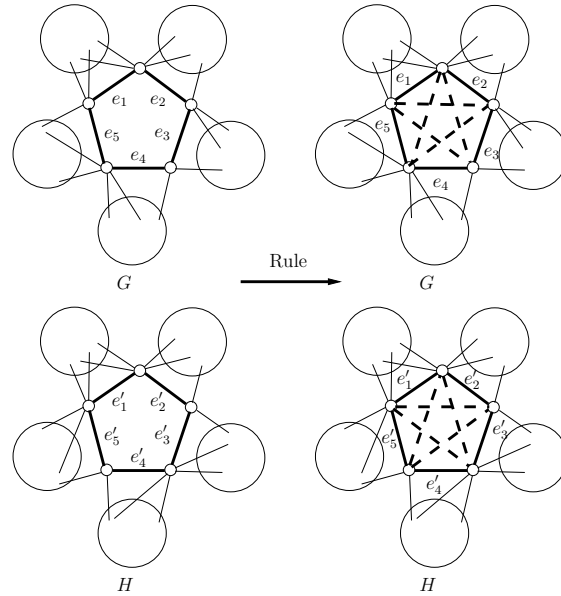
FIG. 7. *An example of an application of Reduction Rule 2; $\varphi(e_i) = e'_i$ for $i \in \{1, \ldots, 5\}$, the vertices of the φ-good bag in G and the corresponding bag of H are white, and the added edges are shown by dashed lines.*

no switch in the sequence splits $E(G[X_t^{(1)}])$. Therefore, $\mathcal{S}$ can be performed on $\tilde{G}$ and transforms $\tilde{G}$ into $\tilde{G}'$ that is $\tilde{S}$-isomorphic to $\tilde{H}$. This means that $(\tilde{G}, \tilde{H}, \tilde{\varphi}, k)$ is a yes-instance. The opposite claim, that if $(\tilde{G}, \tilde{H}, \tilde{\varphi}, k)$ is a yes-instance of WHIT-NEY SWITCHES, then (G, H, φ, k) is a yes-instance as well, is straightforward, because every sequence of Whitney switches transforming $\tilde{G}$ into a graph $\tilde{\varphi}$-isomorphic to $\tilde{H}$ can be applied to G and produces the graph φ-isomorphic to H.

This proves that the rule is safe. To show the remaining claims, observe that the rule transforms $X_t^{(1)}$ and $X_{\alpha(t)}^{(2)}$ into cliques and does not affect other bags. Moreover, for every $v \in X_t^{(1)}$, $\tilde{\varphi}(E_{\tilde{G}[X_t^{(1)}]}) = E_{\tilde{H}[X_{\alpha(t)}^{(2)}]}(\psi(v))$. Therefore, the rule changes neither the Tutte decomposition nor the breakpoint number. For every $t \in W_{\geq 3}^{(1)}$, we can verify whether $X_t^{(1)}$ is φ-good in polynomial time using Lemma 7. By the same lemma, we can compute ψ in polynomial time. Therefore, Reduction Rule 2 can be applied in polynomial time. $\square$

We apply Reduction Rule 2 for all bags of G that are not cliques. We use the same convention as for the first rule and keep the old notation for the obtained graphs, their Tutte decompositions, and the obtained 2-isomorphism.

The next aim is to reduce the number of mutually φ-good bags by "gluing" them into cliques (see Figure 8 for an example).

REDUCTION RULE 3. *For distinct $t_1, t_2 \in W_{\geq 3}^{(1)}$ such that $X_{t_1}^{(1)}$ and $X_{t_2}^{(1)}$ are mutually φ-good,*
  - *compute the φ-isomorphism ψ of $G[X_{t_1}^{(1)} \cup X_{t_2}^{(1)}]$ to $H[X_{\alpha(t_1)}^{(2)} \cup X_{\alpha(t_2)}^{(2)}]$,*
  - *for every $u \in X_{t_1}^{(1)} \setminus X_{t_2}^{(1)}$ and every $v \in X_{t_2}^{(1)} \setminus X_{t_1}^{(1)}$, do the following:*
    - *add the edge uv to G and $\psi(u)\psi(v)$ to H,*
    - *extend φ by setting $\varphi(uv) = \psi(u)\psi(v)$,*
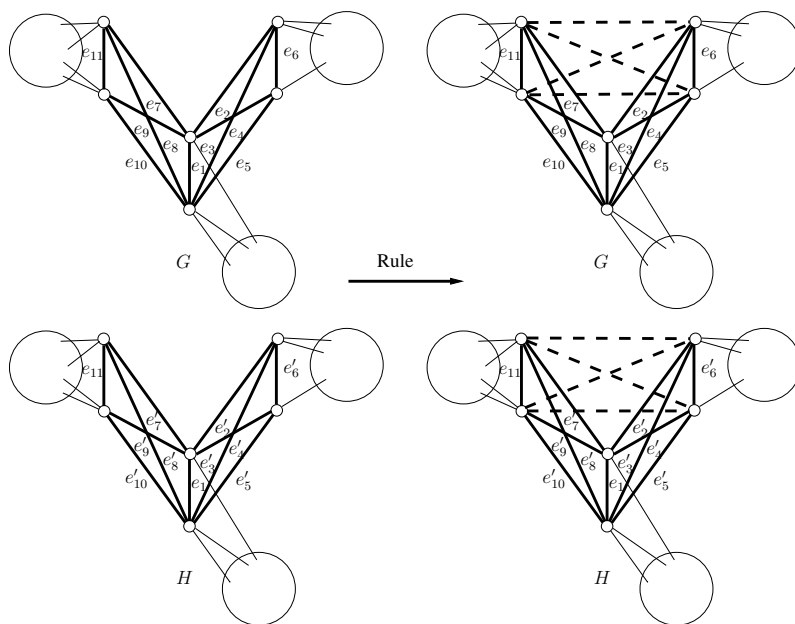
FIG. 8. *An example of an application of Reduction Rule* 3; $\varphi(e_i) = e_i'$ *for* $i \in \{1, \ldots, 11\}$, *the vertices of the mutually φ-good bags of G and the corresponding bags of H are white, and the added edges are shown by dashed lines.*

- *recompute the Tutte decompositions of the obtained graphs and the isomorphism α.*

CLAIM 3. *Reduction Rule* 3 *is safe, does not change the breakpoint number, and can be executed in polynomial time.*

Proof of Claim 3. The proof of safeness essentially repeats the proof for Reduction Rule 2.

Let $t_1, t_2 \in W_{\geq 3}^{(1)}$ be such that $X_{t_1}^{(1)}$ and $X_{t_2}^{(1)}$ are mutually φ-good. By the definition of mutually φ-good bags, $G[X_{t_1}^{(1)} \cup X_{t_2}^{(1)}]$ is φ-isomorphic to $H[X_{\alpha(t_1)}^{(2)} \cup X_{\alpha(t_2)}^{(2)}]$ and, therefore, ψ exists.

Denote by $\tilde{G}$ the graph obtained from G by the addition of one edge uv, and denote by $\tilde{H}$ the graph obtained from H by adding $\psi(u)\psi(v)$. Let also $\tilde{\varphi}$ be the extension of φ on uv. Because ψ is a φ-isomorphism of $G[X_{t_1}^{(1)} \cup X_{t_2}^{(1)}]$ to $H[X_{\alpha(t_1)}^{(2)} \cup X_{\alpha(t_2)}^{(2)}]$, $\tilde{\varphi}$ is a 2-isomorphism of $\tilde{G}$ to $\tilde{H}$.

Suppose that (G, H, φ, k) is a yes-instance of WHITNEY SWITCHES. By Lemma 13, there is a φ-good H-sequence of Whitney switches of length at most k that transforms G into the graph G' that is φ-isomorphic to H. By conditions (i) and (iii) of the definition of a φ-good H-sequence, no switch in the sequence splits $E(G[X_{t_1}^{(1)} \cup X_{t_2}^{(1)}])$. Therefore, $\mathcal{S}$ can be performed on $\tilde{G}$ and transforms $\tilde{G}$ into $\tilde{G}'$ that is $\tilde{S}$-isomorphic to $\tilde{H}$. Hence, $(\tilde{G}, \tilde{H}, \tilde{\varphi}, k)$ is a yes-instance. The opposite implication is straightforward. We conclude that the rule is safe.

To recompute the Tutte decompositions, observe that by Reduction Rule 2, $X_{t_1}^{(1)}$, $X_{t_2}^{(1)}$, $X_{\alpha(t_1)}^{(2)}$, and $X_{\alpha(t_2)}^{(2)}$ are cliques and, therefore, Reduction Rule 3 makes cliques from $X_{t_1}^{(1)} \cup X_{t_2}^{(1)}$ and $X_{\alpha(t_1)}^{(2)} \cup X_{\alpha(t_2)}^{(2)}$. Hence, to recompute the Tutte decompositions

of G and H, we have to identify the nodes t_1 and t_2 of $T^{(1)}$ and the nodes $\alpha(t_1)$ and $\alpha(t_2)$ of $T^{(2)}$, respectively. Every neighbor of t_1 or t_2 in $T^{(1)}$ (every neighbor of $\alpha(t_1)$ or $\alpha(t_2)$, respectively) distinct from these nodes becomes the neighbor of the obtained node t ($\alpha(t)$, respectively). Recall that there is $t' \in W_2^{(1)}$ such that $X_{t_1}^{(1)} \cap X_{t_2}^{(1)} = X_{t'}^{(1)}$. If $X_{t'}^{(1)}$ is not a separator of the graph constructed by the rule (i.e., if t' has exactly two neighbors t_1 and t_2 in the original tree $T^{(1)}$), then we delete t' and $\alpha(t')$ from the trees obtained from $T^{(1)}$ and $T^{(2)}$, respectively. It is straightforward to verify that this procedure, indeed, recomputes the Tutte decompositions and α.

Since Reduction Rule 2 does not affect the bags that are φ-bad and for every $v \in X_{t'}^{(1)}$, $\tilde{\varphi}(E_{\tilde{G}[X_t^{(1)}]}) = E_{\tilde{H}[X_{\alpha(t)}^{(2)}]}(\psi(v))$, we have that the breakpoint number remains the same.

To show that the rule can be executed in polynomial time, note that we can verify whether $X_{t_1}^{(1)}$ and $X_{t_2}^{(1)}$ are mutually φ-good and then compute ψ in polynomial time using Lemma 7. Clearly, recomputing the Tutte decomposition can be done in polynomial time. Then the total running time is polynomial. $\qquad \square$

Reduction Rule 3 is applied exhaustively whenever possible. As before, we do not change the notation for the obtained graphs, their Tutte decompositions, or the obtained 2-isomorphism.

Our next rule is used to perform the Whitney switches that are unavoidable. To state the rule, we define the following auxiliary instance of WHITNEY SWITCHES. Let $C^{(1)}$ and $C^{(2)}$ be copies of C_4 with the edges e_1, e_2, e_3, e_4 and e'_1, e'_2, e'_3, e'_4, respectively, taken in the cycle order. We define $\chi(e_1) = e_1$, $\chi(e_2) = e'_4$, $\chi(e_3) = e'_3$, and $\chi(e_4) = e'_2$. Then χ is a 2-isomorphism of $C^{(1)}$ to $C^{(2)}$, but $C^{(1)}$ and $C^{(2)}$ are not χ-isomorphic. This means that $I = (C^{(1)}, C^{(2)}, \chi, 0)$ is a no-instance of WHITNEY SWITCHES. We call this instance the *trivial* no-instance. Notice that for each no-instance, the input graphs should have at least four vertices each. Therefore, I is a no-instance of minimum size.

REDUCTION RULE 4. *If there is $t \in W_2^{(1)}$ such that $d_{T^{(1)}}(t) = 2$ and for the neighbors t_1 and t_2 of t it holds that $X_{t_1}^{(1)}$ and $X_{t_2}^{(1)}$ are φ-good but not mutually φ-good, then do the following:*
- *find the connected components T_1 and T_2 of $T^{(1)} - t$, and construct $A = \bigcup_{t' \in V(T_1)} X_{t'}^{(1)}$ and $B = \bigcup_{t' \in V(T_2)} X_{t'}^{(1)}$,*
- *perform the Whitney switch with respect to the separation (A, B),*
- *set $k := k - 1$, and if $k < 0$, then return the trivial no-instance and stop.*

An example is shown in Figure 9.

CLAIM 4. *Reduction Rule 4 is safe and changes neither the Tutte decomposition nor the breakpoint member. Furthermore, it can be executed in polynomial time.*

Proof of Claim 4. Suppose that (G, H, φ, k) is a yes-instance of WHITNEY SWITCHES By Lemma 13, there is a φ-good H-sequence $\mathcal{S}_E$ of Whitney switches of length at most k that transforms G into the graph G' that is φ-isomorphic to H.

Let $L = E(G[A])$ and $R = E(G[B])$. We claim that $\mathcal{S}_E$ contains the Whitney switch with respect to (L, R), that is, $\mathcal{S}_E$ has a t-switch. Suppose that this is not the case. Since $X_{t_1}^{(1)}$ and $X_{t_2}^{(1)}$ are φ-good, they are not split by the switches in $\mathcal{S}_E$ by condition (i) of the definition of φ-good switches. Because a t-switch is not in $\mathcal{S}_E$, $E(G[X_{t_1}^{(1)} \cup X_{t_2}^{(1)}])$ is not split. This implies that for the graph G' obtained from G by applying $\mathcal{S}_E$, $G'[X_{t_1}^{(1)} \cup X_{t_2}^{(1)}]$ is φ-isomorphic to $G[X_{t_1}^{(1)} \cup X_{t_2}^{(1)}]$. However,
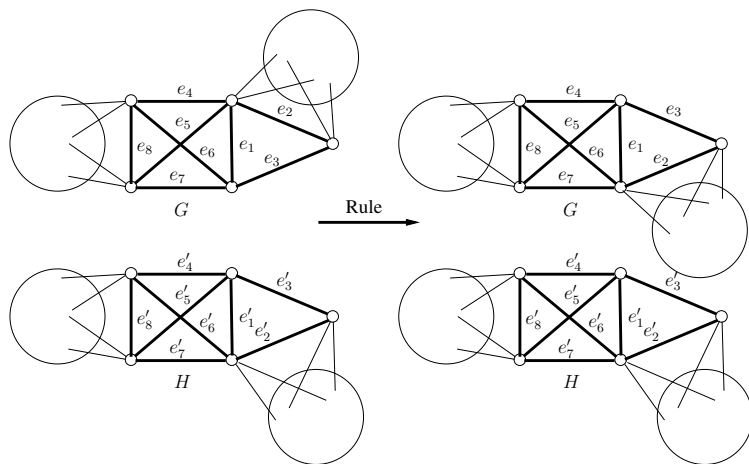
FIG. 9. *An example of an application of Reduction Rule 4; $\varphi(e_i) = e'_i$ for $i \in \{1, \ldots, 8\}$, and the vertices of the switched φ-good bags in G and the corresponding bags of H are white.*

$G[X^{(1)}_{t_1} \cup X^{(1)}_{t_2}]$ is not φ-isomorphic to $H[X^{(2)}_{t_1} \cup X^{(2)}_{t_2}]$, because $X^{(1)}_{t_1}$ and $X^{(1)}_{t_2}$ are not mutually φ-good. This contradicts that $\mathcal{H}$ is an H-sequence. Hence, the Whitney switch with respect to (L, R) is in $\mathcal{S}_E$.

By Lemma 12, the switch with respect to (L, R) is the unique t-switch in $\mathcal{S}_E$, and by Lemma 11, we can assume that this switch is the first switch in $\mathcal{S}_E$. Let $\tilde{G}$ be the graph obtained from G by performing this switch. Then $(\tilde{G}, H, \varphi, k-1)$ is a yes-instance of WHITNEY SWITCHES. Note that $k \geq 1$ in this case and we do not stop by Reduction Rule 4.

Let $\tilde{G}$ be the graph obtained from G by performing the Whitney switch with respect to (L, R). Trivially, if $(\tilde{G}, H, \varphi, k-1)$ is a yes-instance of WHITNEY SWITCHES, then (G, H, φ, k) is a yes-instance. This completes the safeness proof.

The Whitney switch with respect to (L, R) does not change the Tutte decomposition. Also, the switch does not affect the bags that are φ-bad and, moreover, if $X^{(1)}_t$ is φ-bad and t' is at distance two in $T^{(1)}$ from t, then $G[X^{(1)}_t \cup X^{(1)}_{t'}]$ is not modified. Therefore, the breakpoint number remains the same.

We can verify for every $t \in W^{(1)}_2$ such that $d_{T^{(1)}}(t) = 2$, whether for the neighbors t_1 and t_2 of t, it holds that $X^{(1)}_{t_1}$ and $X^{(2)}_{t_2}$ are φ-good but not mutually φ-good in polynomial time by Lemma 7. □

Reduction Rule 4 is applied exhaustively whenever it is possible. Note that after applying this rule, we are able to apply Reduction Rule 3 and we do it.

Suppose that the algorithm did not stop while executing Reduction Rule 4. In the same way as with previous rules, we maintain the initial notation for the obtained graphs, their Tutte decompositions, and the obtained 2-isomorphism.

Our final rule deletes simplicial vertices of degree at least 3.

REDUCTION RULE 5. *If there is a simplicial vertex $v \in V(G)$ with $d_G(v) \geq 3$, then do the following:*
- *find the vertex $u \in V(H)$ such that $E_H(u) = \varphi(E_G(v))$,*
- *set $G := G - v$ and $H := H - u$,*
- *set $\varphi := \varphi|_{E(G) \setminus E_G(v)}$.*

CLAIM 5. *Reduction Rule* 5 *is safe and can be executed in polynomial time. More-over, the rule does not increase the breakpoint number, and the Tutte decompositions obtained by the rule graphs are constructed by the deletions of v and u from the bags of the Tutte decompositions of G and H, respectively.*

Proof of Claim 5. Let v be a simplicial vertex of G of degree at least 3.

Because $N_G[v]$ is a clique of size at least 4, three is a unique $t \in W_{\geq 3}^{(1)}$ such that v is a simplicial vertex of $G[X_t^{(1)}]$ and $v \notin X_{t'}^{(1)}$ for every $t' \in V(T^{(1)})$ distinct from t. Recall that after the exhaustive application of Reduction Rules 2–4, the bags of the Tutte decompositions of G and H, respectively, are cliques. In particular, this means that $G[X_t^{(1)}]$ and $G[X_{\alpha(t)}^{(2)}]$ are 3-connected and, therefore, φ-isomorphic by Lemma 10. Then there is $u \in X_{\alpha(t)}^{(1)}$ such that $\varphi(E_{G[X_t^{(1)}]}(v)) = E_{G[X_{\alpha(t)}^{(2)}]}(u)$. Moreover, u does not belong to any other bag of the Tutte decomposition of H except $X_{\alpha(t)}^{(2)}$. This implies that u is a slimplicial vertex of H and $E_H(u) = \varphi(E_G(v))$. This means that, given v, there is unique $u \in V(H)$ such that $E_H(u) = \varphi(E_G(v))$.

Let $\tilde{G} = G - v$ and $\tilde{H} = H - u$. Since $X_t^{(1)}$ and $X_{\alpha(t)}^{(2)}$ are cliques and v and u do not belong to any separator of size 2 of G and H, respectively, $\tilde{G}$ and $\tilde{H}$ are 2-connected. Since $E_H(u) = \varphi(E_G(v))$, we have that $\tilde{\varphi} = \varphi|_{E(G) \setminus E_G(v)}$ is a 2-isomorphism of $\tilde{G}$ to $\tilde{H}$. Observe also that the Tutte decompositions of $\tilde{G}$ and $\tilde{H}$ are obtained by the deletion of v and u from $X_t^{(1)}$ and $X_{\alpha(t)}^{(2)}$, respectively, and this proves the last part of the claim. Since vertex deletion can only decrease the breakpoint number, $b(\tilde{G}) \leq b(G)$.

Now we show that (G, H, φ, k) is a yes-instance of WHITNEY SWITCHES if and only if $(\tilde{G}, \tilde{H}, \tilde{\varphi}, k)$ is a yes-instance. Here it is more convenient to consider vertex separations.

Since $X_t^{(1)}$ is a clique, for every Whitney separation (A, B) of G, $v \notin A \cap B$ and either $X_t^{(1)} \subseteq A$ or $X_t^{(1)} \subseteq B$, that is, $X_t^{(1)}$ cannot be split. Let $\mathcal{S}$ be an H-sequence. We modify this sequence as follows. For every Whitney separation (A, B) used in $\mathcal{S}$, we replace it by the separation $(A \setminus \{v\}, B \setminus \{v\})$. Denote by $\tilde{\mathcal{S}}$ the obtained sequence. Then $\tilde{\mathcal{S}}$ is an $\tilde{H}$-sequence. This means that if (G, H, φ, k) is yes-instance of WHITNEY SWITCHES, then $(\tilde{G}, \tilde{H}, \tilde{\varphi}, k)$ is a yes-instance.

For the opposite direction, notice that for every Whitney separation (A, B) of $\tilde{G}$, ether $X_t^{(1)} \setminus \{v\} \subseteq A$ or $X_t^{(1)} \setminus \{v\} \subseteq B$. Let $\tilde{\mathcal{S}}$ be an $\tilde{H}$-sequence. For every Whitney separation (A, B) used in $\tilde{\mathcal{S}}$, we replace it by the separation $(A \cup \{v\}, B)$ if $E(G(X_t^{(1)} \setminus \{v\})) \subseteq E(G'[A])$, where G' is the graph obtained from G by the switches prior (A, B), and by $(A, B \cup \{v\})$ otherwise. Then we have that the obtained sequence $\mathcal{S}$ is an H-sequence. Therefore, if $(\tilde{G}, \tilde{H}, \tilde{\varphi}, k)$ is a yes-instance, then (G, H, φ, k) is a yes-instance.

To complete the proof, it remains to observe that a simplicial vertex v can be recognized in polynomial time, and we can find the corresponding vertex u by checking whether $E_H(u) = \varphi(E_G(v))$ in polynomial time. Then the rule can be applied in polynomial time. □

Reduction Rule 5 is applied exhaustively. Let G, H, and φ be the resulting graphs. We also keep the same notation for the Tutte decompositions of G and H and the isomorphism α following the previous convention. This completes the description of our kernelization algorithm.
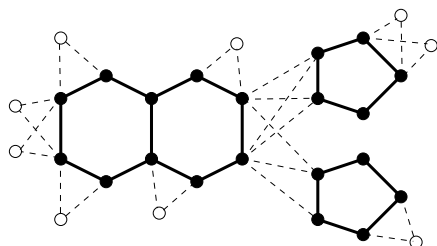
FIG. 10. *The structure of G after applying the reduction rules. The vertices of φ-bad bags are shown by black bullets, and the edges of the subgraph induced by φ-bad bags are shown by solid lines. The vertices of φ-good bags that are not included in φ-bad bags are shown by white bullets, and the edges of subgraphs induced by φ-good bags that are not in the subgraphs induced by φ-bad bags are shown by dashed lines.*

Our next aim is to show that the graphs G and H have bounded size. We prove that $|V(G)| = |V(H)| \leq \max\{39 \cdot b(G) - 27, 3\}$. Clearly, the graphs G and H have the same number of vertices and edges. Therefore, it is sufficient to upper bound $|V(G)|$.

Let $W' \subseteq W_{\geq 3}^{(1)}$ and $W'' \subseteq W_{\geq 3}^{(1)}$ be the sets of $t \in W_{\geq 3}^{(1)}$ such that $X_t^{(1)}$ are φ-good and φ-bad, respectively. Denote $U = \bigcup_{t \in W''} X_t^{(1)}$, that is, U is the set of vertces of the φ-bad bags. We prove the following claim about the structure of G (see Figure 10 for an example).

CLAIM 6. *If $W'' = \emptyset$, then $|V(G)| = 3$. Otherwise, the following holds:*
 (i) *For every $t \in W'$, either $X_t^{(1)} \subseteq U$ or (a) t is a leaf of $T^{(1)}$, (b) $|X_t^{(1)}| = 3$, and (c) $|X_t^{(1)} \setminus U| = 1$ and $X_t^{(1)}$ induces a triangle with two vertices in a φ-bad bag.*
 (ii) *For every $t \in W''$ and every two adjacent vertices $u, v \in X_t^{(1)}$, there are at most two $t' \in W'$ such that $X_{t'}^{(1)} \not\subseteq U$ and $X_t^{(1)} \cap X_{t'}^{(1)} = \{u, v\}$.*

Proof of Claim 6. First, we observe that if $W_2^{(1)} \neq \emptyset$, then for every $t \in W_2^{(1)}$, there is a neighbor t' in $T^{(1)}$ such that $X_{t'}^{(1)}$ is φ-bad. Suppose that this is not the case and there is $t \in W_2^{(1)}$ with the neighbors $t_1, \ldots, t_s$ such that $X_{t_i}^{(1)}$ is φ-good for every $i \in \{1, \ldots, s\}$. Note that $s \geq 2$ by the definition of the Tutte decomposition. Since Reduction Rule 3 is not applicable, for every distinct $i, j \in \{1, \ldots, s\}$, $X_{t_i}^{(1)}$ and $X_{t_j}^{(1)}$ are not mutually φ-good. This implies that $s = 2$, but then we are able to apply Reduction Rule 4, a contradiction.

Next, we show that if the set of φ-bad bags is empty, then $|V(G)| = 3$. For this, we observe that $W_2^{(1)} = \emptyset$. Otherwise, for arbitrary $t \in W_2^{(1)}$, we have that $X_{t'}^{(1)}$ is φ-good for every neighbor t' of t in $T^{(1)}$, contradicting the above observation. Since $W_2^{(1)} = \emptyset$, $|V(T^{(1)})| = 1$ and $G = G[X_t^{(1)}]$ for the unique $t \in W_{\geq 3}^{(1)}$. Recall that all φ-good bags are triangulated by Reduction Rule 2, that is, $X_t^{(1)}$ is a clique with at least three vertices. If $|X_t^{(1)}| \geq 4$, we would be able to apply Reduction Rule 5. We conclude that $|X_t^{(1)}| = 3$ and $|V(G)| = 3$.

Assume from now on that the set W'' of φ-bad bags is nonempty.

To show (i), let $t \in W'$. If $X_t(1) \subseteq U$, then (i) is fulfilled. Let $X_t^{(1)} \setminus U \neq \emptyset$.

We prove (a) by contradiction. Assume that t is not a leaf of $T^{(1)}$. Denote by $t_1, \ldots, t_s \in W_2^{(1)}$, $s \geq 2$, the neighbors of t in $T^{(1)}$. Let $Z = \bigcup_{i=1}^{s} X_{t_i}^{(1)}$. Assume that

$X_t^{(1)} \setminus Z \neq \emptyset$. Since $s \geq 2$, $|Z| \geq 3$. Therefore, $X_t^{(1)}$ is a clique of size at least 4. If there is $v \in X_t^{(1)} \setminus Z$, then v is a simplicial vertex of G. However, in this case, we would be able to apply Reduction Rule 5, a contradiction. Therefore $X_t^{(1)} = Z$. We observed that every node from $W_2^{(1)}$ has a neighbor in W''. Since $t_1, \ldots, t_s \in W_2^{(1)}$, for every $i \in \{1, \ldots, s\}$, t_i has a neighbor t_i' in $T^{(1)}$ such that $t_i' \in W''$. Because $X_{t_i}^{(1)} \subseteq X_{t_i'}^{(1)}$ for $i \in \{1, \ldots, s\}$, $X_t^{(1)} \subseteq Z \subseteq \bigcup_{i=1}^{s} X_{t_i'}^{(1)} \subseteq U$, contradicting the assumption that $X_t^{(1)} \setminus U \neq \emptyset$. This proves (a).

To show (b) and (c), we use the proved property that t is a leaf of $T^{(1)}$. Since t is a leaf, there is the unique neighbor t' of t in $T^{(1)}$. We proved that $t' \in W_2^{(1)}$ has a neighbor t'' such that $X_{t''}^{(1)}$ is φ-bad. Thus, $X_{t'}^{(1)} \subseteq X_{t''}^{(1)}$. Suppose that $|X_t^{(1)} \setminus X_t^{(1)}| \geq 2$. In this case, $X_t^{(1)}$ is a clique of size at least 4 and we would be able to apply Reduction Rule 5 for $v \in X_t^{(1)} \setminus X_t^{(1)}$, because this is a simplicial vertex of G. This cannot happen and we have that $|X_t^{(1)} \setminus X_{t'}^{(1)}| = 1$. Therefore, $|X_t^{(1)} \setminus U| = 1$. Thus, $G[X_t^{(t)}]$ is a triangle with exactly two vertices in a φ-bad bag, that is, (b) and (c) are fulfilled.

Finally, to show (ii), consider $t \in W''$ and let $u, v \in X_t^{(1)}$ be adjacent. Suppose that there are distinct $t_1, \ldots, t_s \in W'$ such that $X_{t_i}^{(1)} \not\subseteq U$ and $X_t^{(1)} \cap X_{t_i}^{(1)} = \{u, v\}$ for all $i \in \{1, \ldots, s\}$. Then there is $t' \in W_2^{(1)}$ such that $X_{t'}^{(1)} = \{u, v\}$ and $t_1 \ldots, t_s$ are neighbors of t' in $T^{(1)}$. By (i), $t_1, \ldots, t_s$ are leaves of $T^{(1)}$. Suppose $s \geq 3$. Then there are distinct $i, j \in \{1, \ldots, s\}$ such that $X_{t_i}^{(1)}$ and $X_{t_j}^{(1)}$ are mutually φ-good. Then we would be able to apply Reduction Rule 3, a contradiction. Hence $s \leq 2$ and (ii) holds. □

By Claim 6, $|V(G)| = 3$ if $W'' = \emptyset$, that is, we have the required upper bound for the number of vertices of G. From now on, we assume that this is not the case. In particular, $b(G) \geq 1$. Then, by Claim 6, we obtain that $|V(G) \setminus U| = |W'| \leq 2|E(G[U])|$. Since U is composed of cycles formed by φ-bad bags, $|E(G[U])| \geq |V(G[U])|$. Hence, to upper bound the number of vertices of G, it is sufficient to upper bound the number of edges of $G[U]$. We do it in the following claim.

CLAIM 7. $|E(G[U])| \leq 13 \cdot b(G) - 9$.

Proof of Claim 7. Denote by S the set of edges of $G[U]$ that are included in at least two φ-bad bags. We claim that $|W''| \leq b(G)$ and $|S| \leq b(G) - 1$. To see this, note that each φ-bad bag contains at least two critical breakpoints and, moreover, such a bag contains at least two nonadjacent critical breakpoints. We exploit this observation and the fact that φ-bad bags have a tree-like structure.

More formally, let T' be the forest obtained from $T^{(1)}$ by the deletion of the nodes $t \in W'$ and then the nodes $t' \in W_2^{(1)}$ that became leaves. Consider $t \in V(T')$ that is a leaf of T' or an isolated node. If t is an isolated node, then $X_t^{(1)}$ contains a crucial breakpoint that is not contained in other φ-bad bags $X_{t'}^{(1)}$ for $t' \in W''$. We *assign* this critical breakpoint to $X_t^{(1)}$. If t is a leaf, then there is the unique $t' \in W_2^{(1)}$ that is the neighbor of t in T'. Since $X_t^{(1)}$ is φ-bad, $X_t^{(1)}$ has a crucial breakpoint v such that $v \notin X_{t'}^{(1)}$. This means that v is not in any $X_{t''}^{(1)}$ for $t'' \in W''$. We *assign* v to $X_t^{(1)}$. Then we delete t from T' and then delete the nodes t' of T' such that $t' \in W_2^{(1)}$ that became leaves. Applying these arguments inductively, we obtain that each φ-bad bag receives an assigned critical breakpoint and each critical breakpoint is assigned to at

most one bag. Therefore, $|W''| \leq b(G)$. Then the number of edges S of $G[U]$ that are included in at least two φ-bad bags is at most $b(G) - 1$.

Further, we show that $|E(G[U]) \setminus S| \leq 12b(G) - 8$. For this, note that for each $t \in W''$, $G[X_t^{(1)}] - S$ is either a cycle (if no edge of S is an edge of $G[X_t^{(1)}]$) or a union of vertex disjoint paths. Denote by $\mathcal{P}$ the family of such cycles and paths taken over all $t \in W''$. We upper bound the total number of edges in the paths and cycles of $\mathcal{P}$.

Because $|W''| \leq b(G)$ and $|S| \leq b(G) - 1$, we have that $|\mathcal{P}| \leq |W''| - (|S| - 1) \leq 2b(G) - 2$. Let $\mathcal{P}'$ be the family off all inclusion maximal subpaths of the elements of $\mathcal{P}$ that do not have a crucial breakpoint as internal vertices. We obtain that $|\mathcal{P}'| \leq |\mathcal{P}| + b(G) \leq 3b(G) - 2$.

We claim that the total length of the paths of $\mathcal{P}'$ is at most $12b(G) - 8$. To obtain a contradiction, assume that the total length is at least $12b(G) - 7$. Then by the pigeonhole principle, there is a path $P \in \mathcal{P}'$ of length at least 5. Let $P = v_0 \cdots v_r$ and assume that P is a segment of the cycle $X_t^{(1)}$ for $t \in W''$. Let $\{t_1, \ldots, t_s\} = N_{T^{(1)}}^2(t)$ and denote $G_t = G[X_t^{(1)} \cup \bigcup_{i=1}^s X_{t_i}^{(1)}]$ and $H_{\alpha(t)} = H[X_{\alpha(t)}^{(2)} \cup \bigcup_{i=1}^s X_{\alpha(t_i)}^{(2)}]$. Since P does not contain edges of S, for every $i \in \{1, \ldots, r\}$ and for every $t' \in W_{\geq 3}^{(1)}$ such that $X_t^{(1)} \cap X_{t'}^{(1)}$, $X_{t'}^{(1)}$ is φ-good. Since $v_1, \ldots, v_{r-1}$ are not crucial breakpoints, there is a path $P' = u_0 \cdots u_r$ in $H[X_{\alpha 1}^{(1)}]$ such that $u_{i-1}u_i = \varphi(v_{i-1}v_i)$ for every $i \in \{1, \ldots, r\}$ and $\varphi(E_{G_t}(v_i)) = E_{H_{\alpha(t)}}(u_i)$ for $i \in \{1, \ldots, r - 1\}$. It follows that P is a φ-good segment of $X_t^{(1)}$. However, this means that we should be able to apply Reduction Rule 1, a contradiction.

Since the total length of paths of $\mathcal{P}'$ is at most $12 \cdot b(G) - 8$, we obtain that $G[U]$ has at most $12 \cdot b(G) - 8 + |S| \leq 13 \cdot b(G) - 9$ edges. □

Summarizing, we have that $|V(G)| = 3$ if $W'' = \emptyset$ and $|V(G)| = |V(G) \setminus U| + |U| \leq 3|E(G[U])| \leq 39 \cdot b(G) - 27$. Recall that Reduction Rules 1–5 do not increase the breakpoint number. Therefore, for the obtained instance (G, H, φ, k) of WHITNEY SWITCHES, $|V(G)| = |V(H)| \leq \min\{39 \cdot b - 27, 3\}$, where b is the breakpoint number of the initial input graph G.

Finally, we have to argue that the kernelization algorithm is polynomial. For this, recall that the intial construction of the Tutte decompositions of the input graphs and the isomorphism α is done in polynomial time. Further, we apply Reduction Rules 1–5, and we proved that each of them can be done in polynomial time in Claims 1–5, respectively. By each application of one of the Reduction Rules 1–4, we add at least one edge. Therefore, the rules are executed at most n^2 times. By Reduction Rule 5, we delete one vertex. Then the rule is called at most n times. This implies that the total running time is polynomial. □

Theorem 4 implies that WHITNEY SWITCHES has a polynomial kernel when parameterized by k and we can show Theorem 2, which we restate.

THEOREM 2. WHITNEY SWITCHES *admits a kernel with $\mathcal{O}(k)$ vertices and is solvable in $2^{\mathcal{O}(k \log k)} \cdot n^{\mathcal{O}(1)}$ time.*

Proof. Let (G, H, φ, k) be an instance of WHITNEY SWITCHES. We compute the breakpoint number $b(G)$. If $b(G) > 2k$, then by Lemma 14, (G, H, φ, k) is a no-instance. In this case, we return the trial no-instance of WHITNEY SWITCHES (defined in the proof of Theorem 4) and stop. Otherwise, we use the kernelization algorithm from Theorem 4 that returns an instance, where each of the graphs has at most $\max\{78 \cdot k - 54, 3\}$ vertices.

Combining the kernelization with the brute-force checking of at most k Whitney switches immediately leads to the algorithm running in $2^{\mathcal{O}(k \log k)} \cdot n^{\mathcal{O}(1)}$ time.   □

In Theorem 3, we proved that WHITNEY SWITCHES is NP-hard when the input graphs are constrained to be cycles. Theorem 4 indicates that it is the presence of bags in the Tutte decompositions that are cycles of length at least 4 that makes WHITNEY SWITCHES difficult, because only such cycles may contain a crucial breakpoint. In particular, we can derive the following straightforward corollary.

COROLLARY 5. *Let (G, H, φ, k) be an instance of* WHITNEY SWITCHES *such that $b(G) = 0$. Then* WHITNEY SWITCHES *for this instance can be solved in polynomial time.*

For example, the condition that $b(G) = 0$ holds when G and H have no induced cycles of length at least 4, that is, when G and H are chordal graphs.

COROLLARY 6. WHITNEY SWITCHES *can be solved in polynomial time on chordal graphs.*

6. Conclusion. We proved that WHITNEY SWITCHES admits a polynomial kernel when parameterized by the breakpoint number of the input graphs, and this implies that the problem has a polynomial kernel when parameterized by k. More precisely, we obtain a kernel, where the graphs have $\mathcal{O}(k)$ vertices. Using this kernel, we can solve WHITNEY SWITCHES in $2^{\mathcal{O}(k \log k)} \cdot n^{\mathcal{O}(1)}$ time. It is natural to ask whether the problem can be solved in a single-exponential in k time.

Another interesting direction of research is to investigate approximability for WHITNEY SWITCHES. In [3], Berman and Karpinski proved that for every $\varepsilon > 0$, it is NP-hard to approximate the reversal distance $d(\pi)$ for a linear permutation π within factor $\frac{1237}{1236} - \varepsilon$. This result can be translated for circular permutations and this allows one to obtain the inapproximability lower bound for WHITNEY SWITCHES on cycles similarly to Theorem 3. From the positive side, the currently best 1.375-approximation for $d(\pi)$ was given by Berman, Hannenhalli, and Karpinski [2]. Due to the close relations between WHITNEY SWITCHES and the sorting by reversal problem, it is interesting to check whether the same approximation ratio can be achieved for WHITNEY SWITCHES.

In WHITNEY SWITCHES, we are given two graphs G and H together with a 2-isomorphism and the task is to decide whether we can apply at most k Whitney switches to obtain a graph G' from G such that G' is φ-isomorphic to H. We can relax the task and ask whether we can obtain G' that is isomorphic to H, that is, we do not require an isomorphism of G to H to be a φ-isomorphism. Formally, we define the following problem.

UNLABELED WHITNEY SWITCHES

Input:	2-Isomorphic graphs G and H, and a nonnegative integer k.
Task:	Decide whether it is possible to obtain a graph G' from G by at most k Whitney switches such that G' is isomorphic to H.

Note that if φ is a 2-isomorphism of G to H, then the minimum number of Whitney switches needed to obtain G' that is φ-isomorphic to H gives an upper bound for the number of Whitney switches required to obtain from G a graph that is isomorphic to G. However, these values can be arbitrarily far apart. Consider two cycles G and H with the same number of vertices. Clearly, G and H are isomorphic but for a given 2-isomorphism φ of G to H, we may need many Whitney switches to
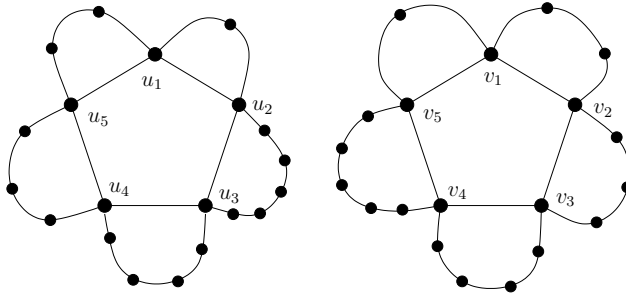
FIG. 11. *Construction of G and H for* $\pi^c = (2, 1, 5, 4, 5)$.

obtain G' that is φ-isomorphic to H and the number of switches is not bounded by any constant.

Using Proposition 2, we can show that UNLABELED WHITNEY SWITCHES is NP-hard for very restricted instances.

PROPOSITION 6. UNLABELED WHITNEY SWITCHES *is* NP-*complete when restricted to* 2-*connected series-parallel graphs even if the input graphs are given together with their* 2-*isomorphism.*

Proof. By Proposition 2, it is NP-complete to decide for a given circular permutation π^c and a nonnegative integer k whether $d^c(\pi^c) \leq k$. We reduce from this problem.

Let $\pi^c = (\pi_1, \ldots, \pi_n)$ be a circular permutation. We construct the graph G as follows:
- construct an n-vertex cycle $C = u_0 u_1 \cdots u_n$ assuming that $u_0 = u_n$,
- for every $i \in \{1, \ldots, n\}$, construct a (u_{i-1}, u_i)-path P_i of length $\pi_i + 1$.

The graph H is constructed in the same way for $\iota^c = (1, \ldots, n)$, that is, we do the following:
- construct an n-vertex cycle $C' = v_0 v_1 \cdots v_n$ assuming that $v_0 = v_n$,
- for every $i \in \{1, \ldots, n\}$, construct a (v_{i-1}, v_i)-path P'_i of length $i + 1$.

The construction of G and H is shown in Figure 11.

We define $\varphi \colon E(G) \to E(H)$ as follows:
- for $i \in \{1, \ldots, n\}$, set $\varphi(u_{i-1} u_i) = u_{\pi_i - 1} u_{\pi_i}$,
- for $i \in \{1, \ldots, n\}$, φ maps the edges of P_i to the edges of P'_i following the path order of the paths staring with the edges incident to v_{i-1} and $u_{\pi_i - 1}$, respectively.

It is straightforward to verify that φ is a 2-isomorphism of G to H.

We claim that $d^c(\pi^c) \leq k$ if and only if (G, H, k) is a yes-instance of UNLABELED WHITNEY SWITCHES.

Suppose that $d^c(\pi^c) \leq k$. Then there is a sorting sequence $\mathcal{S}$ of circular reversals of length at most k for π^c. We use the equivalence between reversals for circular permutations and Whitney switches on cycles described in section 3 and apply the equivalent to $\mathcal{S}$ sequence $\mathcal{S}'$ of Whitney switches for C. It is easy to see that $\mathcal{S}'$ produces the graph isomorphic to H.

For the opposite direction, assume that (G, H, k) is a yes-instance of UNLABELED WHITNEY SWITCHES. Then there is a sequence of Whitney switches $\mathcal{S}$ of length at most k such that the graph G' obtained from G by applying $\mathcal{S}$ is isomorphic to H. Note that the Whitney switch with respect to any Whitney separation (A, B) such
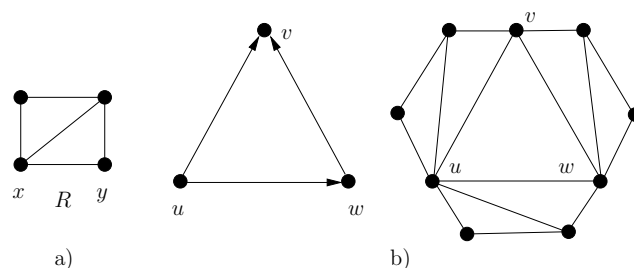
FIG. 12. *Construction of G from a tournament.*

that $A \cap B \subseteq V(P_i)$ for some $i \in \{1, \ldots, n\}$ results in a graph isomorphic to G. Therefore, we can assume that every Whitney switch in $\mathcal{S}$ is performed with respect to a Whitney separation (A, B) such that $A \cap B$ is a pair of nonadjacent vertices of C. We again use the equivalence between circular reversal and Whitney switches on cycles and consider the sequence $\mathcal{S}'$ of circular reversals for π^c that is equivalent to $\mathcal{S}$. Since each path P_i has length $\pi_i + 1$ and every switch from $\mathcal{S}$ does not affect P_i, we obtain that $\mathcal{S}'$ produces the identity circular permutation ι^c. Hence, $d^c(\pi^c) \leq k$. $\square$

Proposition 6 leads to the question about the parameterized complexity of UN-LABELED WHITNEY SWITCHES. In particular, does the problem admit a polynomial kernel when parameterized by k?

Notice that to deal with UNLABELED WHITNEY SWITCHES, we should be able to check whether the input graphs G and H are isomorphic. If we are given a 2-isomorphism φ of G to H, then checking whether G and H are φ-isomorphic can be done in polynomial time by Lemma 7. However, checking whether G and H are isomorphic, even if a 2-isomorphism φ is given, is a complicated task. For example, it can be observed that this is at least as difficult as solving GRAPH ISOMORPHISM on tournaments (recall that a *tournament* is a directed graph such that for every two distinct vertices u and v, either uv or vu is an arc). While GRAPH ISOMORPHISM on tournaments may be easier than the general problem (we refer to [27, 34] for the details), still it is unknown whether this special case can be solved in polynomial time and the best known algorithm is the quasi-polynomial algorithm of Babai [1].

Let T be a tournament. We construct the undirected graph $G(T)$:
  • construct a copy of $V(T)$,
  • for every arc uv of T, construct a copy of the graph R shown in Figure 12(a) and identify the vertex x with u and y with v in the copy of $V(T)$ (see Figure 12(b)).

If T_1 and T_2 are n-vertex tournaments with $n \geq 2$, then it is straightforward to verify that $G(T_1)$ and $G(T_2)$ are 2-isomorphic and it is easy to construct their 2-isomorphism. However, $G(T_1)$ and $G(T_2)$ are isomorphic if and only if T_1 and T_2 are isomorphic.

Given this observation, it is natural to consider UNLABELED WHITNEY SWITCHES on graph classes for which GRAPH ISOMORPHISM is polynomially solvable. For example, what can be said about UNLABELED WHITNEY SWITCHES on planar graphs?

The relation between Whitney switches and sorting by reversals together with the reduction in the proof of Proposition 6 indicates that as the first step, it could be reasonable to investigate the following problem for sequences that generalize SORTING BY REVERSALS for permutations. Let $\pi = (\pi_1, \ldots, \pi_n)$ be a sequence of positive integers; note that now some elements of π may be the same. For $1 \leq i < j \leq n$, we

define the *reversal* $\rho(i, j)$ in exactly the same way as for permutations. Then we can define the *reversal distance* between two n-element sequences such that the multisets of their elements are the same; we assume that the distance is $+\infty$ if the multisets of elements are distinct.

SEQUENCE REVERSAL DISTANCE

Input:	Two n-element sequences π and σ of positive integers and a nonnegative integer k.
Task:	Decide whether the reversal distance between π and σ is at most k.

By the result of Caprara in [7], this problem is NP-complete even if the input sequences are permutations. It is also known that the problem is NP–complete if the input sequences contain only two distinct elements [9]. The question whether SEQUENCE REVERSAL DISTANCE is FPT when parameterized by k was explicitly stated in the survey of Bulteau et al. [6] (in terms of strings) and is open, and only some partial results are known [5]. We also can define the version of SEQUENCE REVERSAL DISTANCE for circular sequences and ask the same question about parameterized complexity. Using the idea behind the reduction in the proof of Proposition 6, it is easy to observe that UNLABELED WHITNEY SWITCHES on 2-connected series-parallel graphs is at least as hard as the circular variant of SEQUENCE REVERSAL DISTANCE.

Acknowledgments. We are grateful to Erlend Raa Vågset for fruitful discussions that initiated the research resulting in the paper. We are also grateful to the anonymous reviewers, whose comments helped to improve the presentation of our results.

REFERENCES

[1] L. BABAI, *Graph isomorphism in quasipolynomial time*, in Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, ACM, 2016, pp. 684–697.

[2] P. BERMAN, S. HANNENHALLI, AND M. KARPINSKI, *1.375-approximation algorithm for sorting by reversals*, in Proceedings of the 10th Annual European Symposium, Rome, Italy, Lecture Notes in Comput. Sci. 2461, Springer, New York, 2002, pp. 200–210.

[3] P. BERMAN AND M. KARPINSKI, *On some tighter inapproximability results*, 26th International Colloquium on Automata, Languages and Programming, ICALP'99, Prague, Czech Republic, Lecture Notes in Comput. Sci. 1644, Springer, New York, 1999, pp. 200–209.

[4] P. BOSE AND F. HURTADO, *Flips in planar graphs*, Comput. Geom., 42 (2009), pp. 60–80.

[5] L. BULTEAU, G. FERTIN, AND C. KOMUSIEWICZ, *(Prefix) reversal distance for (signed) strings with few blocks or small alphabets*, J. Discrete Algorithms, 37 (2016), pp. 44–55.

[6] L. BULTEAU, F. HÜFFNER, C. KOMUSIEWICZ, AND R. NIEDERMEIER, *Multivariate algorithmics for NP-hard string problems*, Bull. Eur. Assoc. Theor. Comput. Sci. EATCS, 114 (2014).

[7] A. CAPRARA, *Sorting by reversals is difficult*, in Proceedings of the 1st Annual International Conference on Research in Computational Molecular Biology, RECOMB 1997, Santa Fe, NM, ACM, 1997, pp. 75–83.

[8] M. CHIMANI, P. HLINENÝ, AND P. MUTZEL, *Vertex insertion approximates the crossing number of apex graphs*, European J. Combin., 33 (2012), pp. 326–335.

[9] D. A. CHRISTIE AND R. W. IRVING, *Sorting strings by reversals and by transpositions*, SIAM J. Discrete Math., 14 (2001), pp. 193–206.

[10] S. CLEARY AND K. ST. JOHN, *Rotation distance is fixed-parameter tractable*, Inform. Process. Lett., 109 (2009), pp. 918–922.

[11] B. COURCELLE, *The monadic second-order logic of graphs* XI: *Hierarchical decompositions of connected graphs*, Theoret. Comput. Sci., 224 (1999), pp. 35–58.

[12] M. CYGAN, F. V. FOMIN, L. KOWALIK, D. LOKSHTANOV, D. MARX, M. PILIPCZUK, M. PILIPCZUK, AND S. SAURABH, *Parameterized Algorithms*, Springer, New York, 2015.

[13] R. DIESTEL, *Graph Theory*, 4th ed., Grad. Texts in Math. 173, Springer, New York, 2012.
[14] R. G. DOWNEY AND M. R. FELLOWS, *Fundamentals of Parameterized Complexity*, Texts Comput. Sci., Springer, New York, 2013.
[15] F. V. FOMIN, D. LOKSHTANOV, S. SAURABH, AND M. ZEHAVI, *Kernelization: Theory of parameterized preprocessing*, Cambridge University Press, Cambridge, 2019.
[16] S. HANNENHALLI AND P. A. PEVZNER, *To cut... or not to cut (applications of comparative physical maps in molecular evolution)*, in Proceedings of the 7th Annual ACM-SIAM Symposium on Discrete Algorithms, Atlanta, GA, ACM/SIAM, 1996, pp. 304–313.
[17] S. HANNENHALLI AND P. A. PEVZNER, *Transforming cabbage into turnip: Polynomial algorithm for sorting signed permutations by reversals*, J. ACM, 46 (1999), pp. 1–27.
[18] J. E. HOPCROFT AND R. E. TARJAN, *Dividing a graph into triconnected components*, SIAM J. Comput., 2 (1973), pp. 135–158.
[19] I. A. KANJ, E. SEDGWICK, AND G. XIA, *Computing the flip distance between triangulations*, Discrete Comput. Geom., 58 (2017), pp. 313–344.
[20] J. D. KECECIOGLU AND D. SANKOFF, *Efficient bounds for oriented chromosome inversion distance*, in Proceedings of the 5th Annual Symposium Combinatorial Pattern Matching, CPM 94, Asilomar, CA, Lecture Notes in Comput. Sci. 807, Springer, New York, 1994, pp. 307–325.
[21] J. D. KECECIOGLU AND D. SANKOFF, *Exact and approximation algorithms for sorting by reversals, with application to genome rearrangement*, Algorithmica, 13 (1995), pp. 180–210.
[22] D. LOKSHTANOV, A. E. MOUAWAD, F. PANOLAN, M. RAMANUJAN, AND S. SAURABH, *Reconfiguration on sparse graphs*, J. Comput. System Sci., 95 (2018), pp. 122–131.
[23] A. LUBIW AND V. PATHAK, *Flip distance between two triangulations of a point set is NP-complete*, Comput. Geom., 49 (2015), pp. 17–23.
[24] J. M. LUCAS, *An improved kernel size for rotation distance in binary trees*, Inform. Process. Lett., 110 (2010), pp. 481–484.
[25] A. E. MOUAWAD, N. NISHIMURA, V. RAMAN, N. SIMJOUR, AND A. SUZUKI, *On the parameterized complexity of reconfiguration problems*, Algorithmica, 78 (2017), pp. 274–297.
[26] P. A. PEVZNER, *Computational Molecular Biology—An Algorithmic Approach*, MIT Press, Cambridge, MA, 2000.
[27] P. SCHWEITZER, *A polynomial-time randomized reduction from tournament isomorphism to tournament asymmetry*, in Proceedings of the 44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, Warsaw, Poland, LIPIcs 80, Schloss Dagstuhl—Leibniz-Zentrum für Informatik, 2017, 66.
[28] A. SOLOMON, P. J. SUTCLIFFE, AND R. LISTER, *Sorting circular permutations by reversal*, in Proceedings of the 8th International Workshop on Algorithms and Data Structures, WADS 2003, Ottawa, Ontario, 2003, Lecture Notes in Comput. Sci. 2748, Springer, New York, 2003, pp. 319–328.
[29] C. THOMASSEN, *Embeddings and minors*, in Handbook of Combinatorics, Vol. 1, Elsevier Sci. B. V., Amsterdam, 1995, pp. 301–349.
[30] K. TRUEMPER, *On Whitney's 2-isomorphism theorem for graphs*, J. Graph Theory, 4 (1980), pp. 43–49.
[31] W. T. TUTTE, *Connectivity in Graphs*, Math. Expositions 15, University of Toronto Press, Toronto, 1966.
[32] W. T. TUTTE, *Graph Theory*, Encyclopedia Math. Appl., Cambridge University Press, Cambridge, 2001.
[33] D. L. VERTIGAN AND G. P. WHITTLE, *A 2-isomorphism theorem for hypergraphs*, J. Combin. Theory Ser. B, 71 (1997), pp. 215–230.
[34] F. WAGNER, *Hardness results for tournament isomorphism and automorphism*, in Proceedings of the 32nd International Symposium, MFCS 2007, Ceský Krumlov, Czech Republic, Lecture Notes in Comput. Sci. 4708, Springer, New York, 2007, pp. 572–583.
[35] H. WHITNEY, *Non-separable and planar graphs*, Trans. Amer. Math. Soc., 34 (1932), pp. 339–362.
[36] H. WHITNEY, *2-Isomorphic Graphs*, Amer. J. Math., 55 (1933), pp. 245–254.