

An Exponential Time Parameterized Algorithm for Planar Disjoint Paths

Daniel Lokshtanov* Pranabendu Misra† Michal Pilipczuk‡ Saket Saurabh§
Meirav Zehavi¶

November 3, 2019

Abstract

In the **Disjoint Paths** problem, the input is an undirected graph G on n vertices and a set of k vertex pairs, $\{s_i, t_i\}_{i=1}^k$, and the task is to find k pairwise vertex-disjoint paths such that the i 'th path connects s_i to t_i . In this paper, we give a parameterized algorithm with running time $2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$ for **Planar Disjoint Paths**, the variant of the problem where the input graph is required to be planar. Our algorithm is based on the unique linkage/treewidth reduction theorem for planar graphs by Adler et al. [JCTB 2017], the algebraic cohomology based technique of Schrijver [SICOMP 1994] and one of the key combinatorial insights developed by Cygan et al. [FOCS 2013] in their algorithm for **Disjoint Paths** on directed planar graphs. To the best of our knowledge our algorithm is the first parameterized algorithm to exploit that the treewidth of the input graph is small in a way completely different from the use of dynamic programming.

*University of California, Santa Barbara, USA. daniello@ucsb.edu

†Max Planck Institute for Informatics, Saarbrücken, Germany. pmisra@mpi-inf.mpg.de

‡Institute of Informatics, University of Warsaw, Poland. michal.pilipczuk@mimuw.edu.pl

§The Institute of Mathematical Sciences, HBNI, Chennai, India. saket@imsc.res.in

¶Ben-Gurion University, Beersheba, Israel. meiravze@bgu.ac.il

1 Introduction

In the Disjoint Paths problem, the input is an undirected graph G on n vertices and a set of k vertex pairs, $\{s_i, t_i\}_{i=1}^k$, and the task is to find k pairwise vertex-disjoint paths connecting s_i to t_i for each $i \in \{1, \dots, k\}$. The Disjoint Paths problem is a fundamental routing problem that finds applications in VLSI layout and virtual circuit routing, and has a central role in Robertson and Seymour’s Graph Minors series. We refer to surveys such as [21, 43] for a detailed overview. The Disjoint Paths problem was shown to be NP-complete by Karp (who attributed it to Knuth) in a followup paper [25] to his initial list of 21 NP-complete problems [24]. It remains NP-complete even if G is restricted to be a grid [33, 30]. On directed graphs, the problem remains NP-hard even for $k = 2$ [20]. For undirected graphs, Perl and Shiloach [35] designed a polynomial time algorithm for the case where $k = 2$. Then, the seminal work of Robertson and Seymour [37] showed that the problem is polynomial time solvable for every fixed k . In fact, they showed that it is *fixed parameter tractable (FPT)* by designing an algorithm with running time $f(k)n^3$. The currently fastest parameterized algorithm for Disjoint Paths has running time $h(k)n^2$ [26]. However, all we know about h and f is that they are computable functions. That is, we still have no idea about what the running time dependence on k really is. Similarly, the problem appears difficult in the realm of approximation, where one considers the optimization variant of the problem where the aim is to find disjoint paths connecting as many of the $\{s_i, t_i\}$ pairs as possible. Despite substantial efforts, the currently best known approximation algorithm remains a simple greedy algorithm that achieves approximation ratio $\mathcal{O}(\sqrt{n})$.

The Disjoint Paths problem has received particular attention when the input graph is restricted to be planar [2, 17, 42, 14]. Adler et al. [2] gave an algorithm for Disjoint Paths on planar graphs (Planar Disjoint Paths) with running time $2^{2^{\mathcal{O}(k)}} n^2$, giving at least a concrete form for the dependence of the running time on k for planar graphs. Schrijver [42] gave an algorithm for Disjoint Paths on *directed* planar graphs with running time $n^{\mathcal{O}(k)}$, in contrast to the NP-hardness for $k = 2$ on general directed graphs. Almost 20 years later, Cygan et al. [14] improved over the algorithm of Schrijver and showed that Disjoint Paths on directed planar graphs is FPT by giving an algorithm with running time $2^{2^{\mathcal{O}(k^2)}} n^{\mathcal{O}(1)}$. The Planar Disjoint Paths problem is well-studied also from the perspective of approximation algorithms, with a recent burst of activity [7, 8, 9, 10, 11]. Highlights of this work include an approximation algorithm with factor $\mathcal{O}(n^{9/19} \log^{\mathcal{O}(1)} n)$ [8] and, under reasonable complexity-theoretic assumptions, hardness of approximating the problem within a factor of $2^{\Omega(\frac{1}{(\log \log n)^2})}$ [10].

In this paper, we consider the parameterized complexity of Planar Disjoint Paths. Prior to our work, the fastest known algorithm was the $2^{2^{\mathcal{O}(k)}} n^2$ time algorithm of Adler et al. [2]. Double exponential dependence on k for a natural problem on planar graphs is something of an outlier—the majority of problems that are FPT on planar graphs enjoy running times of the form $2^{\mathcal{O}(\sqrt{k} \text{ polylog } k)} n^{\mathcal{O}(1)}$ (see, e.g., [15, 18, 19, 29, 36]). This, among other reasons (discussed below), led Adler [1] to pose as an open problem in GROW 2013¹ whether Planar Disjoint Paths admits an algorithm with running time $2^{k^{\mathcal{O}(1)}} n^{\mathcal{O}(1)}$. By integrating tools with origins in algebra and topology, we resolve this problem in the affirmative. In particular, we prove the following.

Theorem 1.1. *The Planar Disjoint Paths problem is solvable in time $2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$.*

In addition to its value as a stand-alone result, our algorithm should be viewed as a piece of an on-going effort of many researchers to make the Graph Minor Theory of Robertson and Seymour algorithmically efficient. The graph minors project is abound with powerful algorithmic and structural results, such as the algorithm for Disjoint Paths [37], Minor Testing [37] (given two undirected graphs, G and H on n and k vertices, respectively, the goal is to check whether G

¹The conference version of [2] appeared in 2011, before [1]. The document [1] erroneously states the open problem for Disjoint Paths instead of for Planar Disjoint Paths. That Planar Disjoint Paths is meant is evident from the statement that a $2^{2^{\mathcal{O}(k)}} n^{\mathcal{O}(1)}$ time algorithm is known.

contains H as a minor), the structural decomposition [38] and the Excluded Grid Theorem [40]. Unfortunately, all of these results suffer from such bad hidden constants and dependence on the parameter k that they have gotten their own term—“galactic algorithms” [31].

It is the hope of many researchers that, in time, algorithms and structural results from Graph Minors can be more algorithmically efficient, perhaps even practically applicable. Substantial progress has been made in this direction, examples include the simpler decomposition theorem of Kawarabayashi and Wollan [28], the faster algorithm for computing the structural decomposition of Grohe et al. [22], the improved unique linkage theorem of Kawarabayashi and Wollan [27], the linear excluded grid theorem on minor free classes of Demaine and Hajiaghayi [16], paving the way for the theory of Bidimensionality [15], and the polynomial grid minor theorem of Chekuri and Chuzhoy [6]. The algorithm for Disjoint Paths is a cornerstone of the entire Graph Minor Theory, and a vital ingredient in the $g(k)n^3$ -time algorithm for Minor Testing. Therefore, *efficient* algorithms for Disjoint Paths and Minor Testing are necessary and crucial ingredients in an algorithmically efficient Graph Minors theory. This makes obtaining $2^{\text{poly}(k)}n^{\mathcal{O}(1)}$ time algorithms for Disjoint Paths and Minor Testing a tantalizing and challenging goal.

Theorem 1.1 is a necessary basic step towards achieving this goal—a $2^{\text{poly}(k)}n^{\mathcal{O}(1)}$ time algorithms for Disjoint Paths on general graphs also has to handle planar inputs, and it is easy to give a reduction from Planar Disjoint Paths to Minor Testing in such a way that a $2^{\text{poly}(k)}n^{\mathcal{O}(1)}$ time algorithm for Minor Testing would imply a $2^{\text{poly}(k)}n^{\mathcal{O}(1)}$ time algorithms for Planar Disjoint Paths. In addition to being a necessary step in the formal sense, there is strong evidence that an efficient algorithm for the planar case will be useful for the general case as well—indeed the algorithm for Disjoint Paths of Robertson and Seymour [37] relies on topology and essentially reduces the problem to surface-embedded graphs. Thus, an efficient algorithm for Planar Disjoint Paths represents a speed-up of the base case of the algorithm for Disjoint Paths of Robertson and Seymour. Coupled with the other recent advances [6, 15, 16, 22, 27, 28], this gives some hope that $2^{\text{poly}(k)}n^{\mathcal{O}(1)}$ time algorithms for Disjoint Paths and Minor Testing may be within reach.

Known Techniques and Obstacles in Designing a $2^{\text{poly}(k)}$ Algorithm. All known algorithms for both Disjoint Paths and Planar Disjoint Paths have the same high level structure. In particular, given a graph G we distinguish between the cases of G having “small” or “large” treewidth. In case the treewidth is large, we distinguish between two further cases: either G contains a “large” clique minor or it does not. This results in the following case distinctions.

1. **Treewidth is small.** Let the treewidth of G be w . Then, we use the known dynamic programming algorithm with running time $2^{\mathcal{O}(w \log w)}n^{\mathcal{O}(1)}$ [41] to solve the problem. It is important to note that, assuming the Exponential Time Hypothesis (ETH), there is no algorithm for Disjoint Paths running in time $2^{o(w \log w)}n^{\mathcal{O}(1)}$ [32], nor an algorithm for Planar Disjoint Paths running in time $2^{o(w)}n^{\mathcal{O}(1)}$ [4].
2. **Treewidth is large and G has a large clique minor.** In this case, we use the good routing property of the clique to find an irrelevant vertex and delete it without changing the answer to the problem. Since this case will not arise for graphs embedded on a surface or for planar graphs, we do not discuss it in more detail.
3. **Treewidth is large and G has no large clique minor .** Using a fundamental structure theorem for minors called the “flat wall theorem”, we can conclude that G contains a large planar piece of the graph and a vertex v that is sufficiently insulated in the middle of it. Applying the unique linkage theorem [39] to this vertex, we conclude that it is irrelevant and remove it. For planar graphs, one can use the unique linkage theorem of Adler et al. [2]. In particular, we use the following result:

Any instance of Disjoint Paths consisting of a planar graph with treewidth at least $82k^{3/2}2^k$ and k terminal pairs contains a vertex v such that every solution to Disjoint Paths can be replaced by an equivalent one whose paths avoid v .




Figure 1: Flow at a vertex and its reduction.

This result says that if the treewidth of the input planar graph is (roughly) $\Omega(2^k)$, then we can find an irrelevant vertex and remove it. A natural question is whether we can guarantee an irrelevant vertex even if the treewidth is $\Omega(\text{poly}(k))$. Adler and Krause [3] exhibited a planar graph G with $k + 1$ terminal pairs such that G contains a $(2^k + 1) \times (2^k + 1)$ grid as a subgraph, Disjoint Paths on this input has a unique solution, and the solution uses all vertices of G ; in particular, *no vertex of G is irrelevant*. This implies that the irrelevant vertex technique can only guarantee a treewidth of $\Omega(2^k)$, even if the input graph is planar.

Combining items (1) and (3), we conclude that the known methodology for Disjoint Paths can only guarantee an algorithm with running time $2^{2^{\mathcal{O}(k)}} n^2$ for Planar Disjoint Paths. Thus, a $2^{\text{poly}(k)} n^{\mathcal{O}(1)}$ -time algorithm for Planar Disjoint Paths appears to require entirely new ideas. As this obstacle was known to Adler et al. [1], it is likely to be the main motivation for Adler to pose the existence of a $2^{\text{poly}(k)} n^{\mathcal{O}(1)}$ time algorithm for Planar Disjoint Paths as an open problem.

Our Methods. Our algorithm is based on a novel combination of two techniques that do not seem to give the desired outcome when used on their own. The first ingredient is the treewidth reduction theorem of Adler et al. [2] that proves that given an instance of Planar Disjoint Paths, the treewidth can be brought down to $2^{\mathcal{O}(k)}$ (explained in item (3) above). This by itself is sufficient for an FPT algorithm (this is what Adler et al. [2] do), but as explained above, it seems hopeless that it will bring a $2^{\text{poly}(k)} n^{\mathcal{O}(1)}$ -time algorithm.

We circumvent the obstacle by using an algorithm for a more difficult problem with a worse running time, namely, Schrijver’s $n^{\mathcal{O}(k)}$ -time algorithm for Disjoint Paths on directed planar graphs [42]. Schrijver’s algorithm has two steps: a “guessing” step where one (essentially) guesses the homology class of the solution paths, and then a surprising homology-based algorithm that, given a homology class, finds a solution in that class (if one exists) in polynomial time. Our key insight is that for Planar Disjoint Paths, if the instance that we are considering has been reduced according to the procedure of Adler et al. [2], then we only need to iterate over $2^{\mathcal{O}(k^2)}$ homology classes in order to find the homology class of a solution, if one exists. The proof of this key insight is highly non-trivial, and builds on a cornerstone ingredient of the recent FPT algorithm of Cygan et al. [14] for Disjoint Paths on directed planar graphs. To the best of our knowledge, this is the first algorithm that finds the exact solution to a problem that exploits that the treewidth of the input graph is small in a way that is different from doing dynamic programming. A technical overview of our methods will appear in the next section.

2 Overview

Homology. In this overview, we explain our main ideas in an *informal* manner. Our starting point is Schrijver’s view [42] of a collection of “non-crossing” (but possibly not vertex- or even edge-disjoint) sets of walks as flows. To work with flows (defined immediately), we deal with directed graphs. (In this context, undirected graphs are treated as directed graphs by replacing each edge by two parallel arcs of opposite directions.) Specifically, we denote an instance of Directed Planar Disjoint Paths as a tuple (D, S, T, g, k) where D is a directed plane graph, $S, T \subseteq V(D)$, $k = |S|$ and $g : S \rightarrow T$ is bijective. Then, a *solution* is a set \mathcal{P} of pairwise vertex-disjoint directed paths in D containing, for each vertex $s \in S$, a path directed from s to $g(s)$.

In the language of flows, each arc of D is assigned a word with letters in $T \cup T^{-1}$ (that is, we treat the set of vertices T also as an alphabet), where $T^{-1} = \{t^{-1} : t \in T\}$. This collection




Figure 2: Two different ways of extracting a walk from a flow.

of words is denoted by $(T \cup T^{-1})^*$ and let 1 denote the empty word. A word is *reduced* if, for all $t \in T$, the letters t and t^{-1} do not appear consecutively. Then, a *flow* is an assignment of reduced words to arcs that satisfies two constraints. First, when we concatenate the words assigned to the arcs incident to a vertex $v \notin S \cup T$ in clockwise order, where words assigned to ingoing arcs are reversed and their letters negated, the result (when reduced) is the empty word 1 (see Fig. 1). This is an algebraic interpretation of the standard flow-conservation constraint. Second, when we do the same operation with respect to a vertex $v \in S \cup T$, then when the vertex is in S , the result is $g(s)$ (rather than the empty word), and when it is in T , the result is t . There is a natural association of flows to solutions: for every $t \in T$, assign the letter t to all arcs used by the path from $g^{-1}(t)$ to t .

Roughly speaking, Schrijver proved that if a flow ϕ is given along with the instance (D, S, T, g, k) , then in *polynomial time* we can either find a solution or determine that there is no solution “similar to ϕ ”. Specifically, two flows are *homologous* (which is the notion of similarity) if one can be obtained from the other by a *set* of “face operations” defined as follows.

Definition 2.1. Let D be a directed plane graph with outer face f , and denote the set of faces of D by \mathcal{F} . Two flows ϕ and ψ are homologous if there exists a function $h : \mathcal{F} \rightarrow (T \cup T^{-1})^*$ such that (i) $h(f) = 1$, and (ii) for every arc $e \in A(D)$, $h(f_1)^{-1} \cdot \phi(e) \cdot h(f_2) = \psi(e)$ where f_1 and f_2 are the faces at the left-hand side and the right-hand side of e , respectively.

Then, a slight modification of Schrijver’s theorem [42] readily gives the following corollary.

Corollary 2.1. There is a polynomial-time algorithm that, given an instance (D, S, T, g, k) of Directed Planar Disjoint Paths, a flow ϕ and a subset $X \subseteq A(D)$, either finds a solution of $(D - X, S, T, g, k)$ or decides that there is no solution of it such that the “flow associated with it” and ϕ are homologous in D .

Discrete Homotopy and Our Objective. While the language of flows and homology can be used to phrase our arguments, it also makes them substantially longer and somewhat obscure because it brings rise to multiple technicalities. For example, different sets of non-crossing walks may correspond to the same flow (see Fig. 2). Instead, we define a notion of *discrete homotopy*, inspired by (standard) homotopy. Specifically, we deal only with collections of non-crossing *and edge-disjoint* walks, called *weak linkages*. Then, two weak linkages are *discretely homotopic* if one can be obtained from the other by using “face operations” that push/stretch its walks across faces and keep them non-crossing and edge-disjoint (see Fig. 3). More precisely, discrete homotopy is an equivalence relation that consists of three face operations, whose precise definition (not required to understand this overview) can be found in Section 5. We note that the order in which face operations are applied is important in discrete homotopy (unlike homology)—we cannot stretch a walk across a face if no walk passes its boundary, but we can execute operations that will move a walk to that face, and then stretch it. In Section 5, we translate Corollary 2.1 to discrete homotopy (and undirected graphs) to derive the following result.

Lemma 2.1. There is a polynomial-time algorithm that, given an instance (G, S, T, g, k) of Planar Disjoint Paths, a weak linkage \mathcal{W} in G and a subset $X \subseteq E(G)$, either finds a solution of $(G - X, S, T, g, k)$ or decides that no solution of it is discretely homotopic to \mathcal{W} in G .

In light of this result, our objective is reduced to the following task.




Figure 3: Moving a walk of a weak linkage (in blue) onto the Steiner tree (the walk in purple) with “face operations” (e.g. a sub-path of the blue path is pushed giving the green sub-path).

Compute a collection of weak linkages such that if there exists a solution, then there also exists a solution (possibly a different one!) that is discretely homotopic to one of the weak linkages in our collection. To prove Theorem 1.1, the size of the collection should be upper bounded by $2^{\mathcal{O}(k^2)}$.

Key Player: Steiner Tree. A key to the proof of our theorem is a very careful construction (done in three steps in Section 6) of a so-called *Backbone Steiner tree*. We use the term Steiner tree to refer to any tree in the *radial completion* of G (the graph obtained by placing a vertex on each face and making it adjacent to all vertices incident to the face) whose set of leaves is precisely $S \cup T$. In the first step, we consider an arbitrary Steiner tree as our Steiner tree R . Having R at hand, we have a more focused goal: we will zoom into weak linkages that are “pushed onto R ”, and we will only generate such weak linkages to construct our collection. Informally, a weak linkage is *pushed onto R* if all of the edges used by all of its walks are *parallel to* edges of R . We do not demand that the edges belong to R , because then the goal described immediately cannot be achieved—instead, we make $4n + 1$ parallel copies of each edge in the radial completion (the number $4n + 1$ arises from considerations in the “pushing process”), and then impose the above weaker demand. Now, our goal is to show that, if there exists a solution, then there also exists one that can be pushed onto R by applying face operations (in discrete homotopy) so that it becomes *identical* to one of the weak linkages in our collection (see Fig. 3).

At this point, one remark is in place. Our Steiner tree R is a subtree of the radial completion of G rather than G itself. Thus, if there exists a solution discretely homotopic to one of the weak linkages that we generate, it might not be a solution in G . We easily circumvent this issue by letting the set X in Lemma 2.1 contain all “fake” edges.

Partitioning a Weak Linkage Into Segments. For the sake of clarity, before we turn to present the next two steps taken to construct R , we begin with the (non-algorithmic) part of the proof where we analyze a (hypothetical) solution towards pushing it onto R . Our most basic notion in this analysis is that of a *segment*, defined as follows (see Fig. 4).

Definition 2.2. For a walk W in the radial completion of G that is edge-disjoint from R , a segment is a maximal subwalk of W that does not “cross” R .

Let $\text{Seg}(W)$ denote the set of segments of W . Clearly, $\text{Seg}(W)$ is a partition of W . Ideally, we would like to upper bound the number of segments of (all the paths of) a solution by $2^{\mathcal{O}(k^2)}$. However, this will not be possible because, while R is easily seen to have only $\mathcal{O}(k)$ vertices of degree 1 or at least 3, it can have “long” maximal degree-2 paths which can give rise to numerous segments (see Fig. 4). To be more concrete, we say that a maximal degree-2 path of R is *long* if it has more than 2^{ck} vertices (for some constant c), and it is *short* otherwise. Then, as the paths of a solution are vertex disjoint, the following observation is immediate.

Observation 2.1. Let \mathcal{P} be a solution. Then, its number of segments that have at least one endpoint on a short path, or a vertex of degree other than 2, of R , is upper bounded by $2^{\mathcal{O}(k)}$.




Figure 4: Segments arising from the crossings of a walk with the Steiner tree.




Figure 5: Separators and flows for a long maximal degree-2 path P in R .

To deal with segments crossing only long paths, several new ideas are required. In what follows, we first explain how to handle segments going across different long paths, whose number *can* be bounded (unlike some of the other types of segments we will encounter).

Segments Between Different Long Paths. To deal with such segments, we modify R (in the second step of its construction). For each long path P with endpoints u and v , we will compute two minimum-size vertex sets, S_u and S_v , such that S_u separates (i.e., intersects all paths with one endpoint in each of the two specified subgraphs) the following subgraphs in the radial completion of G : (i) the subtree of R that contains u after the removal of a vertex u_1 of P that is “very close” to u , and (ii) the subtree of R that contains v after the removal of a vertex u_2 that is “close” to u . The condition satisfied by S_v is symmetric (i.e. u and v switch their roles; see Fig. 5). Here, “very close” refers to distance $2^{c_1 k}$ and “close” refers to distance $2^{c_2 k}$ on the path, for some constants $c_1 < c_2$. Let u' and v' be the vertices of P in the intersection with the separators S_u and S_v respectively. (The selection of u' not to be u itself is of use in the third modification of R .)

To utilize these separators, we need their sizes to be upper bounded by $2^{\mathcal{O}(k)}$. For our initial R , such small separators may not exist. However, the modification we present now will guarantee their existence. Specifically, we will ensure that R does not have any *detour*, which roughly means that each of its maximal degree-2 paths is a shortest path connecting the two subtrees obtained once it is removed. More formally, we define a detour as follows (see Fig. 6).




Figure 6: Detours in the Steiner tree.

Definition 2.3. A detour in R is a pair of vertices $u, v \in V_{\geq 3}(R) \cup V_{=1}(R)$ (i.e. the non-degree 2 vertices in R) that are endpoints of a maximal degree-2 path L of R , and a path P in the radial completion of G , such that (i) P is shorter than L , (ii) one endpoint of P belongs to the component of $R - V(L) \setminus \{u, v\}$ containing u , and (iii) one endpoint of P belongs to the component of $R - V(L) \setminus \{u, v\}$ containing v .

By repeatedly “short-cutting” R , a process that terminates in a linear number of steps, we obtain a new Steiner tree R with no detour. Now, if the separator S_u is large, then there is a large number of vertex-disjoint paths that connect the two subtrees separated by S_u , and all of these paths are “long”, namely, of length at least $2^{c_2k} - 2^{c_1k}$. Based on a result by Bodlaender et al. [5] (whose application requires to work in the radial completion of G rather than G itself), we show that the existence of these paths implies that the treewidth of G is large. Thus, if the treewidth of G were small, all of our separators would have also been small. Fortunately, to guarantee this, we just need to invoke the following known result in a preprocessing step:

Proposition 2.1 ([2]). *There is a $2^{\mathcal{O}(k)}n^2$ -time algorithm that, given an instance (G, S, T, g, k) of Planar Disjoint Paths, outputs an equivalent instance (G', S, T, g, k) of Planar Disjoint Paths where G' is a subgraph of G whose treewidth is upper bounded by 2^{c_k} for some constant c .*

Having separators of size $2^{\mathcal{O}(k)}$, because segments going across different long paths must intersect these separators (or have an endpoint at distance $2^{\mathcal{O}(k)}$ in R from some endpoint of a maximal degree-2 path), we immediately deduce the following.

Observation 2.2. *Let \mathcal{P} be a solution. Then, its number of segments that have one endpoint on one long path, and a second endpoint on a different long path, is upper bounded by $2^{\mathcal{O}(k)}$.*

Segments with Both Endpoints on the Same Long Path. We are now left with segments whose both endpoints belong to the same long path, which have two different kinds of behavior: they may or may not *spiral* around R , where spiraling means that the two endpoints of the segment belong to different “sides” of the path (see Fig. 4). By making sure that at least one vertex in $S \cup T$ is on the outer face of the radial completion of G , we ensure that the cycle formed by any non-spiraling segment together with the subpath of R connecting its two endpoints does not enclose all of $S \cup T$; specifically, we avoid having to deal with segments as the one in Fig. 7.

While it is tempting to try to devise face operations that transform a spiraling segment into a non-spiraling one, this is not always possible. In particular, if the spiral “captures” a path P (of a solution), then when P and the spiral are pushed onto R , the spiral is not reduced to a simple path between its endpoints, but to a walk that “flanks” P . Due to such scenarios, dealing with spirals (whose number we are not able to upper bound) requires special attention. Before we turn to this task, let us consider the non-spiraling segments.

Non-Spiraling Segments. To achieve our main goal, we aim to push a (hypothetical) solution onto R so that the only few parallel copies of each edge will be used. Now, we argue that non-spiraling segments do not pose a real issue in this context. To see this, consider a less refined partition of a solution where some non-spiraling segments are “grouped” as follows (see Fig. 4).

Definition 2.4. A subwalk of a walk W is a preliminary group of W if either (i) it has endpoints on two different maximal degree-2 paths of R or an endpoint in $V_{=1}(R) \cup V_{\geq 3}(R)$ or it is spiraling, or (ii) it is the union of an inclusion-wise maximal collection of segments not of type (i).

The collection of preliminary groups of W is denoted by $\text{PreSegGro}(W)$. Clearly, it is a partition of W . For a weak linkage \mathcal{W} , $\text{PreSegGro}(\mathcal{W}) = \bigcup_{W \in \mathcal{W}} \text{PreSegGro}(W)$. Then,

Observation 2.3. *Let \mathcal{W} be a weak linkage. The number of type-(ii) preliminary groups in $\text{PreSegGro}(\mathcal{W})$ is at most 1 plus the number of type-(i) preliminary groups in $\text{PreSegGro}(\mathcal{W})$.*

Roughly speaking, a type-(i) preliminary group is easily pushed onto R so that it becomes merely a simple path (see Fig. 4). Thus, by Observation 2.3, all type-(ii) preliminary groups




Figure 7: A bad segment that contains all of $S \cup T$ in its cycle.

of a solution in total do not give rise to the occupation of more than $x + 1$ copies of an edge, where x is the number of type-(i) preliminary groups.

Rollback Spirals and Winding Number. Unfortunately, the number of spirals can be huge. Nevertheless, we can pair-up *some* of them so that they will “cancel” each other when pushed onto R (see Fig. 8), thereby behaving like a type-(ii) preliminary group. Intuitively, we pair-up two spirals of a walk if one of them goes from the left-side to the right-side of the path, the other goes from the right-side to the left-side of the same path, and “in between” them on the walk, there are only type-(ii) preliminary groups and spirals that have already been paired-up. We refer to paired-up spirals as *rollback spirals*. (Not all spirals can be paired-up in this manner.) This gives rise to the following strengthening of Definition 2.4.

Definition 2.5. *A subwalk of a walk W is called a group of W if either (i) it is a non-spiral type-(i) preliminary group, or (ii) it is the union of an inclusion-wise maximal collection of segments not of type (i) (i.e., all endpoints of the segments in the group are internal vertices of the same maximal degree-2 path of R). The potential of a group is (roughly) 1 plus its number of non-rollback spirals.*

Now, rather than upper bounding the total number of spirals, we only need to upper bound the number of non-rollback spirals. To this end, we use the notion of *winding number* (in Section 7), informally defined as follows. Consider a solution \mathcal{P} , a path $Q \in \mathcal{P}$, and a long path P of R with separators S_u and S_v . As S_u and S_v are minimal separators in a triangulated graph (the radial completion is triangulated), they are cycles, and as at least one vertex in T belongs to the outer face, they form a ring (see Fig. 9). Each maximal subpath of Q that lies inside this ring can either *visit* the ring, which means that both its endpoints belong to the same separator, or *cross* the ring, which means that its endpoints belong one to S_u and the other to S_v (see Fig. 9). Then, the (absolute value of the) *winding number* of a crossing subpath is the number of times it “winds around” P inside the ring (see Fig. 9). At least intuitively, it should be clear that winding numbers and non-rollback spirals are related. In particular, each ring can only have $2^{\mathcal{O}(k)}$ visitors and crossings subpaths (because the size of each separator is $2^{\mathcal{O}(k)}$), and we only have $\mathcal{O}(k)$ rings to deal with. Thus, it is possible to show that if the winding number of every crossing subpath is upper bounded by $2^{\mathcal{O}(k)}$, then the total number of non-rollback spirals is upper bounded by $2^{\mathcal{O}(k)}$ as well. The main tool we employ to bound the winding number of every crossing path is the following known result (rephrased to simplify the overview).

Proposition 2.2 ([14]). *Let G be a graph embedded in a ring with a crossing path P . Let \mathcal{P} and \mathcal{Q} be two collections of vertex-disjoint crossings paths of the same size. (A path in \mathcal{P} can intersect a path in \mathcal{Q} , but not another path in \mathcal{P} .) Then, G has a collection of crossing paths \mathcal{P}' such that (i) for every path in \mathcal{P} , there is a path in \mathcal{P}' with the same endpoints and vice versa, and (ii) the maximum difference between (the absolute value of) the winding numbers with respect to P of any path in \mathcal{P}' and any path in \mathcal{Q} is at most 6.*




Figure 8: Rollback spirals.

To see the utility of Proposition 2.2, suppose momentarily that none of our rings has visitors. Then, if we could ensure that for each of our rings, there is a collection \mathcal{Q} of vertex-disjoint paths of *maximum size* such that the winding number of each path in \mathcal{Q} is a constant, Proposition 2.2 would have the following implication: if there is a solution, then we can modify it within each ring to obtain another solution such that each crossing subpath of each of its paths will have a constant winding number (under the supposition that the rings are disjoint, which we will deal with later in the overview), see Fig. 9. Our situation is more complicated due to the existence of visitors—we need to ensure that the replacement \mathcal{P}' does not intersect them. On a high-level, this situation is dealt with by first showing how to ensure that visitors do not “go too deep” into the ring on either side of it. Then, we consider an “inner ring” where visitors do not exist, on which we can apply Proposition 2.2. Afterwards, we are able to bound the winding number of each crossing path by $2^{\mathcal{O}(k)}$ (but not by a constant) in the (normal) ring.

Modifying R within Rings. To ensure the existence of the aforementioned collection \mathcal{Q} for each ring, we need to modify R . To this end, consider a long path P with separators S_u and S_v , and let P' be the subpath of P inside the ring defined by the two separators. We compute a maximum-sized collection of vertex-disjoint paths $\text{Flow}(u, v)$ such that each of them has one endpoint in S_u and the other in S_v .² Then, we prove a result that roughly states the following.

Lemma 2.2. *There is a path P^* in the ring defined by S_u and S_v with the same endpoints as P' crossing each path in $\text{Flow}(u, v)$ at most once. Moreover, P^* is computable in linear time.*

Having P^* at hand, we replace P' by P^* . This is done for every maximal degree-2 path, and thus we complete the construction of R . However, at this point, it is not clear why after we perform these replacements, the separators considered earlier remain separators, or that we even still have a tree. Roughly speaking, a scenario as depicted in Fig. 10 can potentially happen. To show that this is not the case, it suffices to prove that there cannot exist a vertex that belongs to two different rings. Towards that, we apply another preprocessing operation: we ensure that the radial completion of G does not have 2^{ck} (for some constant c) *concentric cycles* that contain no vertex in $S \cup T$ by using another result by Adler et al. [2]. Informally, a sequence of concentric cycles is a sequence of vertex-disjoint cycles where each one of them is contained inside the next one in the sequence. Having no such sequences, we prove the following.

Lemma 2.3. *Let R' be any Steiner tree. For every vertex v , there exists a vertex in $V(R')$ whose distance to v (in the radial completion of G) is 2^{ck} for some constant c .*

To see why intuitively this lemma is correct, note that if v was “far” from R' in the *radial completion of G* , then in G itself v is surrounded by a large sequence of concentric cycles that

²This flow has an additional property: there is a tight collection of $\mathcal{C}(u, v)$ of concentric cycles separating S_u and S_v such that paths in $\text{Flow}(u, v)$ do not “oscillate” too much between any two cycles in the collection. Such a maximum flow is said to be *minimal* with respect to $\mathcal{C}(u, v)$.

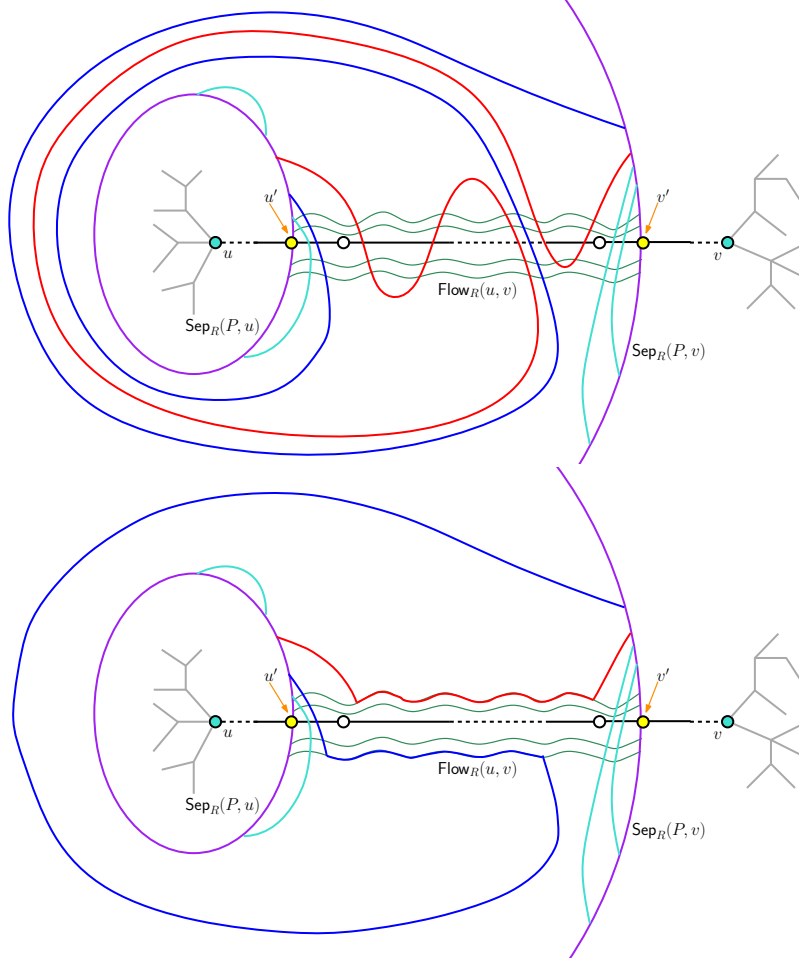


Figure 9: A solution winding in a ring (top), and the “unwinding” of it (bottom).

contain no vertex in $S \cup T$. Having Lemma 2.3 at hand, we show that if a vertex belongs to a certain ring, then it is “close” to at least one vertex of the restriction of R to that ring. In turn, that means that if a vertex belongs to two rings, it can be used to exhibit a “short” path between one vertex in the restriction of R to one ring and another vertex in the restriction of R to the second ring. By choosing constants properly, this path is shown to exhibit a detour in R , and hence we reach a contradiction. (In this argument, we use the fact that for every vertex u , towards the computation of the separator, we considered a vertex u' of distance $2^{c_1 k}$ from u —this subpath between u and u' is precisely that subpath that we will shortcut.)

Pushing a Solution Onto R . So far, we have argued that if there is a solution, then there is also one such that the sum of the potential of all of the groups of all of its paths is at most $2^{\mathcal{O}(k)}$. Additionally, we discussed the intuition why this, in turn, implies the following result.

Lemma 2.4. *If there is a solution \mathcal{P} , then there is a weak linkage pushed onto R that is discretely homotopic to \mathcal{P} and uses at most $2^{\mathcal{O}(k)}$ copies of every edge.*

The formal proof of Lemma 2.4 (in Section 8) is quite technical. On a high level, it consists of three phases. First, we push onto R all *sequences* of the solution—that is, maximal subpaths that touch (but not necessarily cross) R only at their endpoints. Second, we eliminate some U-turns of the resulting weak linkage (see Fig. 11), as well as “move through” R segments with both endpoints being internal vertices of the same maximal degree-2 path of R and crossing it in opposing directions (called *swollen segments*). At this point, we are able to bound by $2^{\mathcal{O}(k)}$ the number of segments of the pushed weak linkage. Third, we eliminate all of the remaining U-turns, and show that then, the number of copies of each edge used must be at most $2^{\mathcal{O}(k)}$. We also modify the pushed weak linkage to be of a certain “canonical form” (see Section 8).




Figure 10: The green and blue rings intersect, which can create cycles in R when replacing paths.




Figure 11: A walk going back and forth along a path of R , which gives rise to U-turns.

Generating a Collection of Pushed Weak Linkages. In light of Lemma 2.4 and Proposition 2.1, it only remains to generate a collection of $2^{\mathcal{O}(k^2)}$ pushed weak linkages that includes all pushed weak linkages (of some canonical form) using at most $2^{\mathcal{O}(k)}$ copies of each edge. (This part, along with the preprocessing and construction of R , are the algorithmic parts of our proof.)

This part of our proof is essentially a technical modification and adaptation of the work of Schrijver [42] (though we need to be more careful to obtain the bound $2^{\mathcal{O}(k^2)}$). Thus, we only give a brief description of it in the overview. Essentially, we generate pairs of a *pairing* and a *template*: a pairing assigns, to each vertex v of R of degree 1 or at least 3, a set of pairs of edges incident to v to indicate that copies of these edges are to be visited consecutively (by at least one walk of the weak linkage under construction); a template further specifies, for each of the aforementioned pairs of edges, how many times copies of these edges are to be visited consecutively (but not which copies are paired-up). Clearly, there is a natural association of a pairing and a template to a pushed weak linkage. Further, we show that to generate all pairs of pairings and templates associated with the weak linkages we are interested in, we only need to consider pairings that in total have $\mathcal{O}(k)$ pairs and templates that assign numbers bounded by $2^{\mathcal{O}(k)}$ (because we deal with weak linkages using $2^{\mathcal{O}(k)}$ copies of each edge):

Lemma 2.5. *There is a collection of $2^{\mathcal{O}(k^2)}$ pairs of pairings and templates that, for any canonical pushed weak linkage \mathcal{W} using only $2^{\mathcal{O}(k)}$ copies of each edge, contains a pair (of a pairing and a template) “compatible” with \mathcal{W} . Further, such a collection is computable in time $2^{\mathcal{O}(k^2)}$.*

Using somewhat more involved arguments (in Section 9), we also prove the following.

Lemma 2.6. *Any canonical pushed weak linkage is “compatible” with exactly one pair of a pairing and a template. Moreover, given a pair of a pairing and a template, if a canonical pushed weak linkage compatible with it exists, then it can be found in time polynomial in its size.*

These two lemmas complete the proof: we can indeed generate a collection of $2^{\mathcal{O}(k^2)}$ pushed weak linkages containing all canonical pushed weak linkages using only $2^{\mathcal{O}(k)}$ copies of any edge.

3 Preliminaries

Let A be a set of elements. A cyclic ordering $<$ on A is an ordering $(a_0, a_1, \dots, a_{|A|-1})$ of the elements in A such that, by enumerating A in *clockwise order* starting at $a_i \in A$, we refer to the ordering $a_i, a_{(i+1) \bmod |A|}, \dots, a_{(i+|A|-1) \bmod |A|}$, and by enumerating A in *counter-clockwise order* starting at $a_i \in A$, we refer to the ordering $a_i, a_{(i-1) \bmod |A|}, \dots, a_{(i-|A|+1) \bmod |A|}$. We

consider all cyclic orderings of A that satisfy the following condition to be the equivalent (up to cyclic shifts): the enumeration of A in cyclic clockwise order starting at a_i , for any $a_i \in A$, produces the same sequence. For a function $f : X \rightarrow Y$ and a subset $X' \subseteq X$, we denote the restriction of f to X' by $f|_{X'}$.

Graphs. Given an undirected graph G , we let $V(G)$ and $E(G)$ denote the vertex set and edge set of G , respectively. Similarly, given a directed graph (digraph) D , we let $V(D)$ and $A(D)$ denote the vertex set and arc set of D , respectively. Throughout the paper, we deal with graphs without self-loops but with parallel edges. Whenever it is not explicitly written otherwise, we deal with undirected graphs. Moreover, whenever G is clear from context, denote $n = |V(G)|$.

For a graph G and a subset of vertices $U \subseteq V(G)$, the subgraph of G induced by U , denoted by $G[U]$, is the graph on vertex set U and edge set $\{\{u, v\} \in E(G) : u, v \in U\}$. Additionally, $G - U$ denotes the graph $G[V(G) \setminus U]$. For a subset of edges $F \subseteq E(G)$, $G - F$ denotes the graph on vertex set $V(G)$ and edge set $E(G) \setminus F$. For a vertex $v \in V(G)$, the set of neighbors of v in G is denoted by $N_G(v)$, and for a subset of vertices $U \subseteq V(G)$, the open neighborhood of U in G is defined as $N_G(U) = \bigcup_{v \in U} N_G(v) \setminus U$. Given three subsets of vertices $A, B, S \subseteq V(G)$, we say that S *separates* A from B if $G - S$ has no path with an endpoint in A and an endpoint in B . For two vertices $u, v \in V(G)$, the *distance* between u and v in G is the length (number of edges) of the shortest path between u and v in G (if no such path exists, then the distance is ∞), and it is denoted by $\text{dist}_G(u, v)$; in case $u = v$, $\text{dist}_G(u, v) = 0$. For two subsets $A, B \subseteq V(G)$, define $\text{dist}_G(A, B) = \min_{u \in A, v \in B} \text{dist}_G(u, v)$. A *linkage* of order k in G is an ordered family \mathcal{P} of k vertex-disjoint paths in G . Two linkages $\mathcal{P} = (P_1, \dots, P_k)$ and $\mathcal{Q} = (Q_1, \dots, Q_k)$ are *aligned* if for all $i \in \{1, \dots, k\}$, P_i and Q_i have the same endpoints.

For a tree T and $d \in \mathbb{N}$, let $V_{\geq d}(T)$ (resp. $V_{=d}(T)$) denote the set of vertices of degree at least (resp. exactly d) in T . For two vertices $u, v \in V(T)$, the unique subpath of T between u and v is denoted by $\text{path}_T(u, v)$. We say that two vertices $u, v \in V(T)$ are *near each other* if $\text{path}_T(u, v)$ has no internal vertex from $V_{\geq 3}(T)$, and call $\text{path}_T(u, v)$ a *degree-2 path*. In case $u, v \in V_{\geq 3}(T) \cup V_{=1}(T)$, $\text{path}_T(u, v)$ is called a *maximal degree-2 path*.

Planarity. A planar graph is a graph that can be embedded in the Euclidean plane, that is, there exists a mapping from every vertex to a point on a plane, and from every edge to a plane curve on that plane, such that the extreme points of each curve are the points mapped to the endpoints of the corresponding edge, and all curves are disjoint except on their extreme points. A *plane graph* G is a planar graph with a fixed embedding. Its faces are the regions bounded by the edges, including the outer infinitely large region. For every vertex $v \in V(G)$, we let $E_G(v) = (e_0, e_1, \dots, e_{t-1})$ for $t \in \mathbb{N}$ where e_0, e_1, \dots, e_{t-1} are the edges incident to v in clockwise order (the decision which edge is e_0 is arbitrary). A planar graph G is *triangulated* if the addition of any edge (not parallel to an existing edge) to G results in a non-planar graph. A plane graph G that is triangulated is 2-connected, and each of its faces is a simple cycle that is a triangle or a cycle that consists of two parallel edges (when the graph is not simple). As we will deal with triangulated graphs, the following proposition will come in handy.

Proposition 3.1 (Proposition 8.2.3 in [34]). *Let G be a triangulated plane graph. Let $A, B \subseteq V(G)$ be disjoint subsets such that $G[A]$ and $G[B]$ are connected graphs. Then, for any minimal subset $S \subseteq V(G) \setminus (A \cup B)$ that separates A from B , it holds that $G[S]$ is a cycle.³*

The *radial graph* (also known as the *face-vertex incidence graph*) of a plane graph G is the planar graph G' whose vertex set consists of $V(G)$ and a vertex v_f for each face f of G , and whose edge set consists of an edge $\{u, v_f\}$ for every vertex $u \in V(G)$ and face f of G such that u is incident to (i.e. lies on the boundary of) f . The *radial completion* of G is the graph G'

³Here, the term cycle also refers to the degenerate case where $|S| = 1$.

obtained by adding the edges of G to the radial graph of G . The graph G' is planar, and we draw it on the plane so that its drawing coincides with that of G with respect to $V(G) \cup E(G)$. Moreover, G' is triangulated and, under the assumption that G had no self-loops, G' also has no self-loops (since all new edges in G' have one endpoint in $V(G)$ and the other endpoint in $V(G') \setminus V(G)$). For a plane graph G , the *radial distance* between two vertices u and v is one less than the minimum length of a sequence of vertices that starts at u and ends at v , such that every two consecutive vertices in the sequence lie on a common face.⁴ We denote the radial distance by $\text{rdist}_G(u, v)$. This definition extends to subsets of vertices: for $X, Y \subseteq V(G)$, $\text{rdist}_G(X, Y)$ is the minimum radial distance over all pairs of vertices in $X \times Y$.

For any $t \in \mathbb{N}$, a sequence $\mathcal{C} = (C_1, C_2, \dots, C_t)$ of t cycles in a plane graph G is said to be *concentric* if for all $i \in \{1, 2, \dots, t-1\}$, the cycle C_i is drawn in the strict interior of C_{i+1} (excluding the boundary, that is, $V(C_i) \cap V(C_{i+1}) = \emptyset$). The *length* of \mathcal{C} is t . For a subset of vertices $U \subseteq V(G)$, we say that \mathcal{C} is U -free if no vertex of U is drawn in the strict interior of C_t .

Treewidth. Treewidth is a measure of how “treelike” is a graph, formally defined as follows.

Definition 3.1 (Treewidth). A tree decomposition of a graph G is a pair (T, β) of a tree T and $\beta : V(T) \rightarrow 2^{V(G)}$, such that

1. for any edge $\{x, y\} \in E(G)$ there exists a node $v \in V(T)$ such that $x, y \in \beta(v)$, and
2. for any vertex $x \in V(G)$, the subgraph of T induced by the set $T_x = \{v \in V(T) : x \in \beta(v)\}$ is a non-empty tree.

The width of (T, β) is $\max_{v \in V(T)} \{|\beta(v)|\} - 1$. The treewidth of G , denoted by $\text{tw}(G)$, is the minimum width over all tree decompositions of G .

The following proposition, due to Cohen-Addad et al. [12], relates the treewidth of a plane graph to the treewidth of its radial completion.

Proposition 3.2 (Lemma 1.5 in [12]). Let G be a plane graph, and let H be its radial completion. Then, $\text{tw}(H) \leq \frac{7}{2} \cdot \text{tw}(G)$.⁵

3.1 Homology and Flows

For an alphabet Σ , denote $\Sigma^{-1} = \{\alpha^{-1} : \alpha \in \Sigma\}$. For a symbol $\alpha \in \Sigma$, define $(\alpha^{-1})^{-1} = \alpha$, and for a word $w = a_1 a_2 \dots a_t$ over $\Sigma \cup \Sigma^{-1}$, define $w^{-1} = a_t^{-1} a_{t-1}^{-1} \dots a_1^{-1}$ and $w^1 = w$. The empty word (the unique word of length 0) is denoted by 1. We say that a word $w = a_1 a_2 \dots a_t$ over $\Sigma \cup \Sigma^{-1}$ is *reduced* if there does not exist $i \in \{1, 2, \dots, t-1\}$ such that $a_i = a_{i+1}^{-1}$. We denote the (infinite) set of reduced words over $\Sigma \cup \Sigma^{-1}$ by $\text{RW}(\Sigma)$. The *concatenation* $w \circ \hat{w}$ of two words $w = a_1 a_2 \dots a_t$ and $\hat{w} = b_1 b_2 \dots b_\ell$ is the word $w^* = a_1 a_2 \dots a_t b_1 b_2 \dots b_\ell$. The *product* $w \cdot \hat{w}$ of two words $w = a_1 a_2 \dots a_t$ and $\hat{w} = b_1 b_2 \dots b_\ell$ in $\text{RW}(\Sigma)$ is a word w^* defined as follows:

$$w^* = a_1 a_2 \dots a_{t-r} b_{r+1} b_{r+2} \dots b_\ell$$

where r is the largest integer in $\{1, 2, \dots, \min(t, \ell)\}$ such that, for every $i \in \{1, 2, \dots, r\}$, $b_i = a_{t+1-i}^{-1}$. Note that w^* is a reduced word, and the product operation is associative. The *reduction* of a word $w = a_1 a_2 \dots a_t$ over $\Sigma \cup \Sigma^{-1}$ is the (reduced) word $w^* = a_1 \cdot a_2 \dots a_t$.

⁴We follow the definition of radial distance that is given in [23] in order to cite a result in that paper verbatim (see Proposition 6.1). We remark that in [5], the definition of a radial distance is slightly different.

⁵More precisely, Cohen-Addad et al. [12] prove that the *branchwidth* of G is at most twice the branchwidth of H . Since the treewidth (plus 1) of a graph is lower bounded by its branchwidth and upper bounded by $\frac{3}{2}$ its branchwidth (see [12]), the proposition follows.

Definition 3.2 (Homology). Let D be a directed plane graph with outer face f , and denote the set of faces of D by \mathcal{F} . Let Σ be an alphabet. Two functions $\phi, \psi : A(D) \rightarrow \text{RW}(\Sigma)$ are homologous if there exists a function $h : \mathcal{F} \rightarrow \text{RW}(\Sigma)$ such that $h(f) = 1$, and for every arc $e \in A(D)$, we have $h(f_1)^{-1} \cdot \phi(e) \cdot h(f_2) = \psi(e)$ where f_1 and f_2 are the faces at the left-hand side and the right-hand side of e , respectively.

The following observation will be useful in later results.

Observation 3.1. Let $\alpha, \beta, \gamma : A(D) \rightarrow \text{RW}(\Sigma)$ be three functions such that α, β and β, γ are pairs of homologous functions. Then, α, γ is also a pair of homologous functions.

Proof. Let f and g be the functions witnessing the homology of α, β and β, γ , respectively. Then, it is easy to check that the function $h = g \circ f$ (i.e. the composition of g and f) witnesses the homology of α, γ . \square

Towards the definition of flow, we denote an instance of Directed Planar Disjoint Paths by a tuple (D, S, T, g, k) where D is a directed plane graph, $S, T \subseteq V(D)$, $k = |S|$ and $g : S \rightarrow T$ is bijective. We remark that the implicit assumption that the pairs of terminals to connect are disjoint is made without loss of generality because, if a terminal occur $x > 1$ times, we can make x copies (twins) of that terminal and thereby obtain an equivalent instance. A *solution* of an instance (D, S, T, g, k) of Directed Planar Disjoint Paths is a set \mathcal{P} of pairwise vertex-disjoint directed paths in D that contains, for every vertex $s \in S$, a path directed from s to $g(s)$. When we write $\text{RW}(T)$, we treat T as an alphabet—that is, every vertex in T is treated as a symbol.

Definition 3.3 (Flow). Let (D, S, T, g, k) be an instance of Directed Planar Disjoint Paths. Let $\phi : A(D) \rightarrow \text{RW}(T)$ be a function. For any vertex $v \in V(D)$, denote the concatenation $\phi(e_1)^{\epsilon_1} \circ \phi(e_2)^{\epsilon_2} \cdots \phi(e_r)^{\epsilon_r}$ by $\text{conc}(v)$, where e_1, e_2, \dots, e_r are the arcs incident to v in clockwise order where the first arc e_1 is chosen arbitrarily, and for each $i \in \{1, 2, \dots, r\}$, $\epsilon_i = 1$ if v is the head of e_i and $\epsilon_i = -1$ if v is the tail of e_i . Then, the function ϕ is a flow if:⁶

1. For every vertex $v \in V(D) \setminus (S \cup T)$, the reduction of $\text{conc}(v)$ is 1.
2. For every vertex $v \in S$, (i) $\text{conc}(v) = a_0 a_1 \dots a_{\ell-1}$ is a word of length $\ell \geq 1$, and (ii) there exists $i \in \{0, 1, \dots, \ell-1\}$ such that the reduction of $a_i a_{(i+1) \bmod \ell} \cdots a_{(i+\ell-1) \bmod \ell}$ equals a_i , where $a_i = g(v)$ if the arc e associated with a_i has v as its tail, $a_i = g(v)^{-1}$ otherwise.
3. For every vertex $v \in T$, (i) $\text{conc}(v) = a_0 a_1 \dots a_{\ell-1}$ is a word of length $\ell \geq 1$, and (ii) there exists $i \in \{0, 1, \dots, \ell-1\}$ such that the reduction of $a_i a_{(i+1) \bmod \ell} \cdots a_{(i+\ell-1) \bmod \ell}$ equals a_i , where $a_i = v$ if the arc e associated with a_i has v as its head, $a_i = v^{-1}$ otherwise.

In the above definition, the conditions on the reduction of $\text{conc}(v)$ for each vertex $v \in V(D)$ are called *flow conservation* constraints. Informally speaking, these constraints resemble “usual” flow conservation constraints and ensure that every two walks that carry two different alphabets do not cross. The association between solutions to (D, S, T, g, k) and flows is defined as follows.

Definition 3.4. Let (D, S, T, g, k) be an instance of Directed Planar Disjoint Paths. Let \mathcal{P} be a solution of (D, S, T, g, k) . The flow $\phi : A(D) \rightarrow \text{RW}(T)$ associated with \mathcal{P} is defined as follows. For every arc $e \in A(D)$, define $\phi(e) = 1$ if there is no path in \mathcal{P} that traverses e , and $\phi(e) = t$ otherwise where $t \in T$ is the end-vertex of the (unique) path in \mathcal{P} that traverses e .

The following proposition, due to Schrijver [42], also holds for the above definition of flow.

⁶We note that there is slight technical difference between our definition and the definition in Section 3.1 in [42]. There, a flow must put only a single alphabet (v or v^{-1} in $T \cup T^{-1}$) on the arcs incident on vertices in $S \cup T$.

Proposition 3.3 (Proposition 5 in [42]). *There exists a polynomial-time algorithm that, given an instance (D, S, T, g, k) of Directed Planar Disjoint Paths and a flow ϕ , either finds a solution of (D, S, T, g, k) or determines that there is no solution of (D, S, T, g, k) such that the flow associated with it and ϕ are homologous.*

We need a slightly more general version of this proposition because we will work with an instance of Directed Planar Disjoint Paths where D contains some “fake” edges that emerge when we consider the radial completion of the input graph—the edges added to the graph in order to attain its radial completion are considered to be fake.

Corollary 3.1. *There exists a polynomial-time algorithm that, given an instance (D, S, T, g, k) of Directed Planar Disjoint Paths, a flow ϕ and a subset $X \subseteq A(D)$, either finds a solution of $(D - X, S, T, g, k)$ or determines that there is no solution of $(D - X, S, T, g, k)$ such that the flow associated with it and ϕ are homologous.⁷*

Proof. Given (D, S, T, g, k) , ϕ and X , the algorithm constructs an equivalent instance (D', S, T, g, k) of DIRECTED PLANAR DISJOINT PATHS and a flow ϕ' as follows. Each arc $(u, v) \in X$ is replaced by a new vertex w and two new arcs (whose drawing coincides with the former drawing of (u, v)), (u, w) and (v, w) , and we define $\phi'(u, w) = \phi(u, v)$ and $\phi'(v, w) = \phi(u, v)^{-1}$. For all other arcs $a \in A(D) \cap A(D')$, $\phi'(a) = \phi(a)$. It is immediate to verify that ϕ' is a flow in D' , and that (D', S, T, g, k) admits a solution that is homologous to ϕ' if and only if (D, S, T, g, k) admits a solution that is disjoint from X and homologous to ϕ . Indeed, any solution of one of these instances is also a solution of the other one. Now, we apply Proposition 3.3 to either obtain a solution to (D', S, T, g, k) homologous to ϕ' or conclude that no such solution exists. \square

4 Preprocessing to Obtain a Good Instance

We denote an instance of Planar Disjoint Paths similarly to an instance of Directed Planar Disjoint Paths (in Section 3) except that now the graph is denoted by G rather than D to stress the fact that it is undirected. Formally, an instance of Planar Disjoint Paths is a tuple (G, S, T, g, k) where G is a plane graph, $S, T \subseteq V(G)$, $k = |S|$ and $g : S \rightarrow T$ is bijective. We remind that the implicit assumption that the pairs of terminals to connect are disjoint was justified in Section 3. Moreover, we say that (G, S, T, g, k) is *nice* if every vertex in $S \cup T$ has degree 1 and $S \cap T = \emptyset$. The vertices in $S \cup T$ are called *terminals*. Let H_G be the radial completion of G . We choose a plane embedding of H_G so that one of the terminals, $t^* \in T$, will lie on the outer face.⁸ A *solution* of an instance (G, S, T, g, k) of Planar Disjoint Paths is a set \mathcal{P} of pairwise vertex-disjoint paths in G that contains, for every vertex $s \in S$, a path with endpoints s and $g(s)$.

The following proposition eliminates all long sequences of $S \cup T$ -free concentric cycles.

Proposition 4.1 ([2]). *There exists a $2^{\mathcal{O}(k)}n^2$ -time algorithm that, given an instance (G, S, T, g, k) of Planar Disjoint Paths, outputs an equivalent instance (G', S, T, g, k) of Planar Disjoint Paths where G' is a subgraph of G that has no sequence of $S \cup T$ -free concentric cycles whose length is larger than 2^{ck} for some fixed constant $c \geq 1$.*

Additionally, the following proposition reduces the treewidth of G . In fact, Proposition 4.1 was given by Adler et al. [2] as a step towards the proof of the following proposition. However, while the absence of a long sequence of concentric cycles implies that the treewidth is small, the reversed statement is not correct (i.e. small treewidth does not imply the absence of a long sequence of concentric cycles). Having small treewidth is required but not sufficient for our arguments, thus we cite both propositions.

⁷Note that ϕ and homology concern D rather than $D - X$.

⁸This can be ensured by starting with an embedding of H_G on a sphere, picking some face where t^* lies as the outer face, and then projecting the embedding onto the plane.

Proposition 4.2 (Lemma 10 in [2]). *There exists a $2^{\mathcal{O}(k)}n^2$ -time algorithm that, given an instance (G, S, T, g, k) of Planar Disjoint Paths, outputs an equivalent instance (G', S, T, g, k) of Planar Disjoint Paths where G' is a subgraph of G whose treewidth is upper bounded by 2^{ck} for some fixed constant $c \geq 1$.⁹*

The purpose of this section is to transform an arbitrary instance of Planar Disjoint Paths into a so called “good” instance, defined as follows.

Definition 4.1. *An instance (G, S, T, g, k) of Planar Disjoint Paths is good if it is nice, at least one terminal $t^* \in T$ belongs to the outer faces of both G and its radial completion H_G , the treewidth of G is upper bounded by 2^{ck} , and G has no $S \cup T$ -free sequence of concentric cycles whose length is larger than 2^{ck} . Here, $c \geq 1$ is the fixed constant equal to the maximum among the fixed constants in Propositions 4.1 and 4.2.*

Towards this transformation, note that given an instance (G, S, T, g, k) of Planar Disjoint Paths and a terminal $v \in S$, we can add to G a degree-1 vertex u adjacent to v , and replace v by u in S and in the domain of g . This operation results in an equivalent instance of Planar Disjoint Paths. Furthermore, it does not increase the treewidth of G (unless the treewidth of G is 0, in which it increases to be 1). The symmetric operation can be done for any terminal $v \in T$. By repeatedly applying these operations, we can easily transform (G, S, T, g, k) to an equivalent nice instance. Moreover, the requirement that at least one terminal $t^* \in T$ belongs to the outer faces of both G and its radial completion can be made without loss of generality by drawing G appropriately in the first place. Thus, we obtain the following corollary of Propositions 4.1 and 4.2. Note that k remains unchanged.

Corollary 4.1. *There exists a $2^{\mathcal{O}(k)}n^2$ -time algorithm that, given an instance (G, S, T, g, k) of Planar Disjoint Paths, outputs an equivalent good instance (G', S', T', g', k) of Planar Disjoint Paths where $|V(G')| = \mathcal{O}(|V(G)|)$.*

We remark that our algorithm, presented in Section 10, will begin by applying Corollary 4.1. To simplify arguments ahead, it will be convenient to suppose that every edge in H_G has $4|V(G)| + 1 = 4n + 1$ parallel copies. Thus, we slightly abuse the notation H_G and use it to refer to H_G enriched with such a number of parallel copies of each edge. For a pair of adjacent vertices $u, v \in V(H_G)$, we will denote the $4n + 1$ parallel copies of edges between them by $e_{-2n}, e_{-2n+1}, \dots, e_{-1}, e_0, e_1, e_2, \dots, e_{2n}$ where $e = \{u, v\}$, such that when the edges incident to u (or v) are enumerated in cyclic order, the occurrences of e_i and e_{i+1} are consecutive for every $i \in \{-2n, -2n + 1, \dots, 2n - 1\}$, and e_{-2n} and e_{2n} are the outermost copies of e . Thus, for every $i \in \{-2n + 1, -2n + 2, \dots, 2n - 1\}$, e_i lies on the boundary of exactly two faces: the face bounded by e_{i-1} and e_i , and the face bounded by e_i and e_{i+1} . When the specification of the precise copy under consideration is immaterial, we sometimes simply use the notation e .

5 Discrete Homotopy

The purpose of this section is to assert that rather than working with homology (Definition 3.2) or the standard notion of homotopy, to obtain our algorithm it will suffice to work with a notion called *discrete homotopy*. Working with discrete homotopy will *substantially shorten and simplify* our proof, particularly Section 8. Translation from discrete homotopy to homology is straightforward, thus readers are invited to skip the proofs in this section when reading the paper for the first time. We begin by defining the notion of a weak linkage. This notion is a

⁹While the running time in Lemma 10 in [2] is stated to be $2^{2^{\mathcal{O}(k)}}n^2$, the proof is easily seen to imply the bound $2^{\mathcal{O}(k)}n^2$ in our statement. The reason why Adler et al. [2] obtain a double-exponential dependence on k when they solve Planar Disjoint Paths is *not* due to the computation that attains a tree decomposition of width $\text{tw} = 2^{\mathcal{O}(k)}$, but it is because that upon having such a tree decomposition, they solve the problem in time $2^{\mathcal{O}(\text{tw})}n = 2^{2^{\mathcal{O}(k)}}n$.




Figure 12: Crossing (left) and non-crossing (right) walks.

generalization of a linkage (see Section 3) that concerns walks rather than paths, and which permits the walks to intersect one another in vertices. Here, we only deal with walks that may repeat vertices but which do not repeat edges. Moreover, weak linkages concern walks that are *non-crossing*, a property defined as follows (see Fig. 12).

Definition 5.1 (Non-Crossing Walks). *Let G be a plane graph, and let W and W' be two edge-disjoint walks in G . A crossing of W and W' is a tuple $(v, e, \hat{e}, e', \hat{e}')$ where e, \hat{e} are consecutive in W , e', \hat{e}' are consecutive in W' , $v \in V(G)$ is an endpoint of e, \hat{e}, e' and \hat{e}' , and when the edges incident to v are enumerated in clockwise order, then exactly one edge in $\{e', \hat{e}'\}$ occurs between e and \hat{e} . We say that W is self-crossing if, either it has a repeated edge, or it has two edge-disjoint subwalks that are crossing.*

We remark that when we say that a collection of edge-disjoint walks is non-crossing, we mean that none of its walks is self-crossing and no pair of its walks has a crossing.

Definition 5.2 (Weak Linkage). *Let G be a plane graph. A weak linkage in G of order k is an ordered family of k edge-disjoint non-crossing walks in G . Two weak linkages $\mathcal{W} = (W_1, \dots, W_k)$ and $\mathcal{Q} = (Q_1, \dots, Q_k)$ are aligned if for all $i \in \{1, \dots, k\}$, W_i and Q_i have the same endpoints.*

Given an instance (G, S, T, g, k) of Planar Disjoint Paths, a weak linkage \mathcal{W} in G (or H_G) is sensible if its order is k and for every terminal $s \in S$, \mathcal{W} has a walk with endpoints s and $g(s)$.

The following observation is clear from Definitions 5.1 and 5.2.

Observation 5.1. *Let G be a plane graph, and let \mathcal{W} be a weak linkage in G . Let e_1, e_2 and e_3, e_4 be two pairs of edges in $E(\mathcal{W})$ that are all distinct and incident on a vertex v , and there is some walk in \mathcal{W} where e_1, e_2 are consecutive, and likewise for e_3, e_4 . Then, in a clockwise enumeration of edges incident to v , the pairs e_1, e_2 and e_3, e_4 do not cross, that is, they do not occur as e_1, e_3, e_2, e_4 in clockwise order (including cyclic shifts).*

We now define the collection of operations applied to transform one weak linkage into another weak linkage aligned with it (see Fig. 13). We remark that the face push operation is not required for our arguments, but we present it here to ensure that discrete homotopy defines an equivalence relation (in case it will find other applications that need this property).

Definition 5.3 (Operations in Discrete Homotopy). *Let G be a triangulated plane graph with a weak linkage \mathcal{W} , and a face f that is not the outer face with boundary cycle C . Let $W \in \mathcal{W}$.*

- **Face Move.** *Applicable to (W, f) if there exists a subpath P of C such that (i) P is a subwalk of W , (ii) $1 \leq |E(P)| \leq |E(f)| - 1$, and (iii) no edge in $E(C) \setminus E(P)$ belongs to any walk in \mathcal{W} . Then, the face move operation replaces P in W by the unique subpath of C between the endpoints of P that is edge-disjoint from P .*
- **Face Pull.** *Applicable to (W, f) if C is a subwalk Q of W . Then, the face pull operation replaces Q in W by a single occurrence of the first vertex in Q .*
- **Face Push.** *Applicable to (W, f) if (i) no edge in $E(C)$ belongs to any walk in \mathcal{W} , and (ii) there exist two consecutive edges e, e' in W with common vertex $v \in V(C)$ (where W visits*




Figure 13: Face operations.

e first, and v is visited between e and e') and an order (clockwise or counter-clockwise) to enumerate the edges incident to v starting at e such that the only edges that belong to $E(\mathcal{W}) \cup E(C)$ enumerated between e and e' are the two edges of $E(C)$ incident to v . Let \tilde{e} be the first among the two edges of $E(C)$ that is enumerated. Then, the face push operation replaces the occurrence of v between e and e' in \mathcal{W} by the traversal of C starting at \tilde{e} .

We verify that the application of a single operation results in a weak linkage.

Observation 5.2. *Let G be a triangulated plane graph with a weak linkage \mathcal{W} , and a face f that is not the outer face. Let $W \in \mathcal{W}$ with a discrete homotopy operation applicable to (W, f) . Then, the result of the application is another weak linkage aligned to \mathcal{W} .*

Then, discrete homotopy is defined as follows.

Definition 5.4 (Discrete Homotopy). *Let G be a triangulated plane graph with weak linkages \mathcal{W} and \mathcal{W}' . Then, \mathcal{W} is discretely homotopic to \mathcal{W}' if there exists a finite sequence of discrete homotopy operations such that when we start with \mathcal{W} and apply the operations in the sequence one after another, every operation is applicable, and the final result is \mathcal{W}' .*

We verify that discrete homotopy gives rise to an equivalence relation.

Lemma 5.1. *Let G be a triangulated plane graph with weak linkages $\mathcal{W}, \mathcal{W}'$ and \mathcal{W}'' . Then, (i) \mathcal{W} is discretely homotopic to itself, (ii) if \mathcal{W} is discretely homotopic to \mathcal{W}' , then \mathcal{W}' is discretely homotopic to \mathcal{W} , and (iii) if \mathcal{W} is discretely homotopic to \mathcal{W}' and \mathcal{W}' is discretely homotopic to \mathcal{W}'' , then \mathcal{W} is discretely homotopic to \mathcal{W}'' .*

Proof. Statement (i) is trivially true. The proof of statement (ii) is immediate from the observation that each discrete homotopy operation has a distinct inverse. Indeed, every face move operation is invertible by a face move operation (applied to the same walk and cycle). Additionally, every face pull operation is invertible by a face push operation (applied to the same walk and cycle), and vice versa. Hence, given the sequence of operations to transform \mathcal{W} to \mathcal{W}' , say ϕ , the sequence of operations to transform \mathcal{W}' to \mathcal{W} is obtained by first writing the operations of ϕ in reverse order and then inverting each of them. Finally, Statement (iii) follows by considering the sequence of discrete homotopy operations obtained by concatenating the sequence of operations to transform \mathcal{W} to \mathcal{W}' and the sequence of operations to transform \mathcal{W}' to \mathcal{W}'' . \square

Towards the translation of discrete homotopy to homology, we need to associate a flow with every weak linkage and thereby extend Definition 3.4.

Definition 5.5. *Let (D, S, T, g, k) be an instance of Directed Planar Disjoint Paths. Let \mathcal{W} be a sensible weak linkage in D . The flow $\phi : A(D) \rightarrow \text{RW}(T)$ associated with \mathcal{W} is defined as follows. For every arc $e \in A(D)$, define $\phi(e) = 1$ if there is no walk in \mathcal{W} that traverses e , and $\phi(e) = t$ otherwise where $t \in T$ is the end-vertex of the (unique) walk in \mathcal{W} that traverses e .*

Additionally, because homology concerns directed graphs, we need the following notation. Given a graph G , we let \vec{G} denote the directed graph obtained by replacing every edge $e \in E(G)$ by two arcs of opposite orientations with the same endpoints as e . Notice that G and the underlying graph of \vec{G} are not equal (in particular, the latter graph contains twice as many edges as the first one). Given a weak linkage \mathcal{W} in G , the weak linkage in \vec{G} that corresponds to \mathcal{W} is the weak linkage obtained by replacing each edge e in each walk in \mathcal{W} , traversed from u to v , by the copy of e in \vec{G} oriented from u to v .

Now, we are ready to translate discrete homotopy to homology.

Lemma 5.2. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths where G is triangulated. Let \mathcal{W} be a sensible weak linkage in G . Let \mathcal{W}' be a weak linkage discretely homotopic to \mathcal{W} . Let $\widehat{\mathcal{W}}$ and $\widehat{\mathcal{W}'}$ be the weak linkages corresponding to \mathcal{W} and \mathcal{W}' in \vec{G} , respectively. Then, the flow associated with $\widehat{\mathcal{W}}$ is homologous to the flow associated with $\widehat{\mathcal{W}'}$.*

Proof. Let ϕ and ψ be the flows associated with $\widehat{\mathcal{W}}$ and $\widehat{\mathcal{W}'}$, respectively, in \vec{G} . Consider a sequence O_1, O_2, \dots, O_ℓ of discrete homotopy operations that, starting from \mathcal{W} , result in \mathcal{W}' . We prove the lemma by induction on ℓ . Consider the case when $\ell = 1$. Then, the sequence contains only one discrete homotopy operation, which a face move, face pull or face push operation. Let this operation be applied to a face f and a walk $W \in \mathcal{W}$, where the walk W goes from $s \in S$ to $g(s) \in T$. Let C be the boundary cycle of f in G , and let \vec{C} denote the collection of arcs in \vec{G} obtained from the edges of C . After this discrete homotopy operation, we obtain a walk $W' \in \mathcal{W}'$, which differs from W only in the subset of edges C . All other walks are identical between \mathcal{W} and \mathcal{W}' . Hence, $\widehat{\mathcal{W}}$ and $\widehat{\mathcal{W}'}$ differ in \vec{G} only in a subset of \vec{C} . Then observe that the flows ϕ and ψ are identical everywhere in $A(\vec{G})$ except for a subset of \vec{C} . More precisely, Let $P = \vec{C} \cap \widehat{\mathcal{W}}$ and $P' = \vec{C} \cap \widehat{\mathcal{W}'}$. Then $\phi(e) = g(s)$ if $e \in P$ and $\phi(e) = 1$ if $e \in \vec{C} - P$; a similar statement holds for ψ and P' . Furthermore, it is clear from the description of each of the discrete homotopy operations that P and P' have no common edges and $P \cup P'$ is the (undirected)¹⁰ cycle \vec{C} in \vec{G} .

It only remains to describe the homology between the flows ϕ and ψ , which is exhibited by a function h on the faces of \vec{G} . Then h assigns 1 to all faces of \vec{G} that lie in the exterior of \vec{C} , and $g(s)$ to all the faces that lie in the interior of \vec{C} . Note that h assigns 1 to the outer face of \vec{G} . It is easy to verify that h is indeed a homology between ψ and ϕ , that is, for any edge $e \in A(\vec{G})$ it holds that $\psi(e) = h(f_1)^{-1} \cdot \phi(e) \cdot h(f_2)$, where f_1 and f_2 are the faces on the left and the right of e with respect to its orientation. This proves the case where $\ell = 1$.

Now for $\ell > 1$, consider the weak linkage $\widehat{\mathcal{W}}^*$ obtained from $\widehat{\mathcal{W}}$ after applying the sequence $O_1, O_2, \dots, O_{\ell-1}$. Then by the induction hypothesis, we can assume that the flow associated with $\widehat{\mathcal{W}}^*$, say ψ^* is homologous to ϕ . Further, applying O_ℓ to $\widehat{\mathcal{W}}^*$ gives us $\widehat{\mathcal{W}'}$, and hence the flows ψ^* and ϕ are also homologous. Hence, by Observation 3.1 the flows ϕ and ψ are homologous. \square

Having Corollary 3.1 and Lemma 5.2 at hand, we prove the following theorem.

Lemma 5.3. *There exists a polynomial-time algorithm that, given an instance (G, S, T, g, k) of Planar Disjoint Paths where G is triangulated, a sensible weak linkage \mathcal{W} in G and a subset $X \subseteq E(G)$, either finds a solution of $(G - X, S, T, g, k)$ or determines that no solution of $(G - X, S, T, g, k)$ is discretely homotopic to \mathcal{W} in G .*

Proof. We first convert the given instance of Planar Disjoint Paths into an instance of Directed Planar Disjoint Paths as follows. We convert the graph G into the digraph \vec{G} , as described earlier. Then we construct \vec{X} from X by picking the two arcs of opposite orientation for each edge in X . Then we convert the sensible weak linkage \mathcal{W} into a weak linkage $\widehat{\mathcal{W}}$ in \vec{G} . Finally, we obtain the flow ϕ in \vec{G} associated with $\widehat{\mathcal{W}}$. Next, we apply Corollary 3.1 to the instance (\vec{G}, S, T, g, k) , \vec{X} and ϕ . Then either it returns a solution $\widehat{\mathcal{P}}$ that is disjoint from \vec{X} , or that there is no solution

¹⁰That is, the underlying undirected graph of \vec{C} is a cycle.

that is homologous to ϕ and disjoint from \vec{X} . In the first case, $\widehat{\mathcal{P}}$ can be easily turned into a solution \mathcal{P} for the undirected input instance that is disjoint from X . In the second case, we can conclude that the undirected input instance has no solution that is discretely homotopic to \mathcal{W} . Indeed, if this were not the case, then consider a solution \mathcal{P} to $(G - X, S, T, g, k)$ that is discretely homotopic to \mathcal{W} . Then we have a solution $\widehat{\mathcal{P}}$ to the directed instance that is disjoint to \vec{X} . Hence, by Lemma 5.2, the flow associated with $\widehat{\mathcal{P}}$ is homologous to ϕ , the flow associated with $\widehat{\mathcal{W}}$. Hence, $\widehat{\mathcal{P}}$ is a solution to the instance (\vec{G}, S, T, g, k) that is disjoint from \vec{X} and whose flow is homologous to ϕ . But this is a contradiction to Corollary 3.1. \square

As a corollary to this lemma, we derive the following result.

Corollary 5.1. *There exists a polynomial-time algorithm that, given an instance (G, S, T, g, k) of Planar Disjoint Paths and a sensible weak linkage \mathcal{W} in H_G , either finds a solution of (G, S, T, g, k) or decides that no solution of (G, S, T, g, k) is discretely homotopic to \mathcal{W} in H_G .*

Proof. Consider the instance (H_G, S, T, g, k) along with the set $X = E(H_G) \setminus E(G)$ of forbidden edges. We then apply Lemma 5.3 to (H_G, S, T, g, k) , X and \mathcal{W} (note that H_G is triangulated). If we obtain a solution to this instance, then it is also a solution in G since it traverses only edges in $E(H_G) \setminus X = E(G)$. Else, we correctly conclude that there is no solution of $(H_G - X, S, T, g, k)$ (and hence also of (G, S, T, g, k)) that is discretely homotopic to \mathcal{W} in H_G . \square

6 Construction of the Backbone Steiner Tree

In this section, we construct a tree that we call a backbone Steiner tree $R = R^3$ in H_G . Recall that H_G is the radial completion of G enriched with $4|V(G)| + 1$ parallel copies for each edge. These parallel copies will not be required during the construction of R , and therefore we will treat H_G as having just one copy of each edge. Hence, we can assume that H_G is a simple planar graph, and then $E(H_G) = O(n)$ where n is the number of vertices in G . We denote $H = H_G$ when G is clear from context. The tree R will be proven to admit the following property: if the input instance is a Yes-instance, then it admits a solution $\mathcal{P} = (P_1, \dots, P_k)$ that is discretely homotopic to a weak linkage $\mathcal{W} = (W_1, \dots, W_k)$ in H aligned with \mathcal{P} that uses at most $2^{O(k)}$ edges parallel to those in R , and none of the edges not parallel to those in R . We use the term Steiner tree to refer to any subtree of H whose set of leaves is precisely $S \cup T$. To construct the backbone Steiner tree $R = R^3$, we start with an arbitrary Steiner tree R^1 in H . Then over several steps, we modify the tree to satisfy several useful properties.

6.1 Step I: Initialization

We initialize R^1 to be an arbitrarily chosen Steiner tree. Thus, R^1 is a subtree of H such that $V_{=1}(R^1) = S \cup T$. The following observation is immediate from the definition of a Steiner tree.

Observation 6.1. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths. Let R' be a Steiner tree. Then, $|V_{=1}(R')| = 2k$ and $|V_{\geq 3}(R')| \leq 2k - 1$.*

Before we proceed to the next step, we claim that every vertex of H is, in fact, “close” to the vertex set of R_1 . For this purpose, we need the following proposition by Jansen et al. [23].

Proposition 6.1 (Proposition 2.1 in [23]). *Let G be a plane graph and with disjoint subsets $X, Y \subseteq V(G)$ such that $G[X]$ and $G[Y]$ are connected graphs and $\text{rdist}_G(X, Y) = d \geq 2$. For any $r \in \{0, 1, \dots, d - 1\}$, there is a cycle C in G such that all vertices $u \in V(C)$ satisfy $\text{rdist}_G(X, \{u\}) = r$, and such that $V(C)$ separates X and Y in G .*

Additionally, we need the following simple observation.

Observation 6.2. *Let G be a triangulated plane graph. Then, for any pair of vertices $u, v \in V(G)$, $\text{dist}_G(u, v) = \text{rdist}_G(u, v)$.*

Proof. Let $\text{rdist}_G(u, v) = t$, and consider a sequence of vertices $u = x_1, x_2, \dots, x_{t+1} = v$ that witnesses this fact—then, every two consecutive vertices in this sequence have a common face. Since G is triangulated, we have that $\{x_i, x_{i+1}\} \in E(G)$ for every two consecutive vertices x_i, x_{i+1} , $1 \leq i \leq t$. Hence, x_1, x_2, \dots, x_{t+1} is a walk from u to v in G with t edges, and therefore $\text{dist}_G(u, v) \leq \text{rdist}_G(u, v)$. Conversely, let $\text{dist}_G(u, v) = \ell$; then, there is a path with ℓ edges from u to v in G , which gives us a sequence of vertices $u = y_1, y_2, \dots, y_{\ell+1} = v$ where each pair of consecutive vertices forms an edge in G . Since G is planar, each such pair of consecutive vertices y_i, y_{i+1} , $1 \leq i \leq \ell$, must have a common face. Therefore, $\text{rdist}_G(u, v) \leq \text{dist}_G(u, v)$. \square

It is easy to see that Observation 6.2 is not true for general plane graph. However, this observation will be useful for us because the graph H , where we construct the backbone Steiner tree, is triangulated. We now present the promised claim, whose proof is based on Proposition 6.1, Observation 6.2 and the absence of long sequences of $S \cup T$ -free concentric cycles in good instances. Here, recall that c is the fixed constant in Corollary 4.1. We remark that, for the sake of clarity, throughout the paper we denote some natural numbers whose value depends on k by notations of the form $\alpha_{\text{subscript}}(k)$ where the subscript of α hints at the use of the value.

Lemma 6.1. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R' be a Steiner tree. For every vertex $v \in V(H)$, it holds that $\text{dist}_H(v, V(R')) \leq \alpha_{\text{dist}}(k) := 4 \cdot 2^{ck}$.*

Proof. Suppose, by way of contradiction, that $\text{dist}_H(v^*, V(R')) > \alpha_{\text{dist}}(k)$ for some vertex $v^* \in V(H)$. Since H is the (enriched) radial completion of G , it is triangulated. By Observation 6.2, $\text{rdist}_H(u, v) = \text{dist}_H(u, v)$ for any pair of vertices $u, v \in V(H)$. Thus, $\text{rdist}_H(v^*, V(R')) > \alpha_{\text{dist}}(k)$. By Proposition 6.1, for any $r \in \{0, 1, \dots, \alpha_{\text{dist}}(k)\}$, there is a cycle C_r in H such that all vertices $u \in V(C_r)$ satisfy $\text{rdist}_H(v^*, u) = \text{dist}_H(v^*, u) = r$, and such that $V(C_r)$ separates $\{v^*\}$ and $V(R')$ in H . In particular, these cycles must be pairwise vertex-disjoint, and each one of them contains either (i) v^* in its interior (including the boundary) and $V(R')$ in its exterior (including the boundary), or (ii) v^* in its exterior (including the boundary) and $V(R')$ in its interior (including the boundary). We claim that only case (i) is possible. Indeed, suppose by way of contradiction that C_i , for some $r \in \{0, 1, \dots, \alpha_{\text{dist}}(k)\}$, contains v^* in its exterior and $V(R')$ in its interior. Because the outer face of H contains a terminal $t^* \in T$ and $t^* \in V(R')$, we derive that $t^* \in V(C_i)$. Thus, $\text{rdist}_H(v^*, t^*) = i \leq \alpha_{\text{dist}}(k)$. However, because $t^* \in V(R')$, this is a contradiction to the supposition that $\text{dist}_H(v^*, V(R')) > \alpha_{\text{dist}}(k)$. Thus, our claim holds true. From this, we conclude that $\mathcal{C} = (C_0, C_1, \dots, C_{\alpha_{\text{dist}}(k)})$ is a $V(R')$ -free sequence of concentric cycles in H . Since $S \cup T \subseteq V(R')$, it is also $S \cup T$ -free.

Consider some odd integer $r \in \{1, 2, \dots, \alpha_{\text{dist}}(k)\}$. Note that every vertex $u \in V(C_r)$ that does not belong to $V(G)$ lies in some face f of G , and that the two neighbors of u in C_r must belong to the boundary of f (by the definition of radial completion). Moreover, each of the vertices on the boundary of f is at distance (in H) from u that is the same, larger by one or smaller by one, than the distance of f from u , and hence none of these vertices can belong to any C_i for $i > r + 1$ as well as $i < r + 1$. For every $r \in \{1, 2, \dots, \alpha_{\text{dist}}(k) - 1\}$ such that $r \bmod 3 = 1$, define C'_r as some cycle contained in the closed walk obtained from C_r by replacing every vertex $u \in V(C_r) \setminus V(G)$, with neighbors x, y on C_r , by a path from x to y on the boundary of the face of G that corresponds to u . In this manner, we obtain an $S \cup T$ -free sequence of concentric cycles in G whose length is at least 2^{ck} . However, this contradicts the supposition that (G, S, T, g, k) is good. \square

6.2 Step II: Removing Detours

In this step, we modify the Steiner tree to ensure that there exist no “shortcuts” via vertices outside the Steiner tree. This property will be required in subsequent steps to derive additional

properties of the Steiner tree. To formulate this, we need the following definition (see Fig. 6).

Definition 6.1 (Detours in Trees). *A subtree T of a graph G has a detour if there exist two vertices $u, v \in V_{\geq 3}(T) \cup V_{=1}(T)$ that are near each other, and a path P in G , such that*

1. P is shorter than $\text{path}_T(u, v)$, and
2. one endpoint of P belongs to the connected component of $T - V(\text{path}_T(u, v)) \setminus \{u, v\}$ that contains u , and the other endpoint of P belongs to the connected component of $T - V(\text{path}_T(u, v)) \setminus \{u, v\}$ that contains v .

Such vertices u, v and path P are said to witness the detour. Moreover, if P has no internal vertex from $(V(T) \setminus V(\text{path}_T(u, v))) \cup \{u, v\}$ and its endpoints do not belong to $V_{=1}(T) \setminus \{u, v\}$, then u, v and P are said to witness the detour compactly.

We compute a witness for a detour as follows. Note that this lemma also implies that, if there exists a detour, then there exists a compact witness rather than an arbitrary one.

Lemma 6.2. *There exists an algorithm that, given a good instance (G, S, T, g, k) of Planar Disjoint Paths and a Steiner tree R' , determines in time $\mathcal{O}(k^2 \cdot n)$ whether R' has a detour. In case the answer is positive, it returns u, v and P that witness the detour compactly.*

Proof. Let $Q = \text{path}_{R'}(u, v) - \{u, v\}$ for some two vertices $u, v \in V_{\geq 3}(T) \cup V_{=1}(T)$ that are near each other. Then, $R' - V(Q)$ contains precisely two connected components: R'_u and R'_v that contain u and v , respectively. Consider a path P of minimum length between vertices $x \in V(R'_u)$ and $y \in V(R'_v)$ in H , over all choices of x and y . Further, we choose P so that contains as few vertices of $(S \cup T) \setminus \{u, v\}$ as possible. Suppose that $|E(P)| \leq |E(\text{path}_{R'}(u, v))| - 1$. Then, we claim that P is a compact detour witness. To prove this claim, we must show that (i) no internal vertex of P lies in $(V(R') \setminus V(\text{path}_{R'}(u, v))) \cup \{u, v\} = V(R'_u) \cup V(R'_v)$, and (ii) the endpoints of P do not lie in $V_{=1}(R') \setminus \{u, v\} = (S \cup T) \setminus \{u, v\}$. The first property follows directly from the choice of P . Indeed, if P were a path from $x \in V(R'_u)$ to $y \in V(R'_v)$, which contained an internal vertex $z \in V(R'_u)$, then the subpath P' of P with endpoints z and y is a strictly shorter path from $V(R'_u)$ to $V(R'_v)$ (the symmetric argument holds when $z \in V(R'_v)$).

For the second property, we give a proof by contradiction. To this end, suppose that some terminal $w \in (S \cup T) \setminus \{u, v\}$ belongs to P . Necessarily, $w \in V(R'_u) \cup V(R'_v)$ (by the definition of a Steiner tree). Without loss of generality, suppose that $w \in V(R'_u)$. By the first property, w must be an endpoint of P . Let $z \in V(R'_v)$ be the other endpoint of P . Because the given instance is good, w has degree 1 in G , thus we can let $n(w)$ denote its unique neighbor in G . Observe that w lies on only one face of G , which contains both w and $n(w)$. Hence, w is adjacent to exactly two vertices in H : $n(w)$ and a vertex $f(w) \in V(H) \setminus V(G)$. Furthermore, $\{n(w), f(w)\} \in E(H)$, i.e. $w, n(w)$ and $f(w)$ form a triangle in H . Thus, P contains exactly one of $n(w)$ or $f(w)$ (otherwise, we can obtain a strictly shorter path connecting w and the other endpoint of P that contradicts the choice of P). Let $a(w) \in \{n(w), f(w)\}$ denote the neighbor of w in P , and note that, by the first property, $a(w) \notin V(R'_u)$. Note that it may be the case that $a(w) = z$. Since w is a leaf of R' , exactly one of $n(w)$ and $f(w)$ is adjacent to w in R' , and we let $b(w)$ denote this vertex. Because $w \neq u$, we have that $V(R'_u)$ contains but is not equal to $\{w\}$, and therefore $b(w) \in V(R'_u)$. In turn, by the first property, this means that $a(w) \neq b(w)$ (because otherwise $a(w) \neq z$ and hence it is an internal vertex of P , which cannot belong to $V(R'_u)$). Because $w, a(w)$ and $b(w)$ form a triangle in H , we obtain a path $P' \neq P$ in H by replacing w with $b(w)$ in P . Observe that P' connects the vertex $b(w) \in V(R'_u)$ to the vertex $z \in V(R'_v)$. Furthermore, because $|E(P')| = |E(P)|$, and P' contains strictly fewer vertices of $(S \cup T) \setminus \{u, v\}$ compared to P , we contradict the choice of P . Therefore, P also satisfies the second property, and we conclude that u, v, P compactly witness a detour in R' .

We now show that a compact detour in R' can be computed in $\mathcal{O}(k^2 \cdot n)$ time. First, observe that if there is a detour witnessed by some u, v and P , then $u, v \in V_{\geq 3}(R') \cup V_{=1}(R')$. By

Observation 6.1, $|V_{\geq 3}(R') \cup V_{=1}(R')| \leq 4k$. Therefore, there are at most $16k^2$ choices for the vertices u and v . We consider each choice, and test if there is detour for it in linear time as follows. Fix a choice of distinct vertices $u, v \in V_{\geq 3}(R') \cup V_{=1}(R')$, and check if they are near each other in R' in $O(|V(R')|)$ time by validating that each internal vertex of $\text{path}_{R'}(u, v)$ has degree 2. If they are not near each other, move on to the next choice. Otherwise, consider the path $Q = \text{path}_{R'}(u, v) - \{u, v\}$, and the trees R'_u and R'_v of $R' - V(Q)$ that contain u and v , respectively. Now, consider the graph \tilde{H} derived from H by first deleting $(V(Q) \cup S \cup T) \setminus \{u, v\}$ and then introducing a new vertex r adjacent to all vertices in $V(R'_u)$. We now run a breadth first search (BFS) from r in \tilde{H} . This step takes $\mathcal{O}(n)$ time since $|E(\tilde{H})| = \mathcal{O}(n)$ (because H is planar). From the BFS-tree, we can easily compute a shortest path P between a vertex $x \in V(R'_u)$ and a vertex $y \in V(R'_v)$. Observe that $V(P) \cap (S \cup T) \subseteq \{u, v\}$ by the construction of \tilde{H} . If $|E(P)| < |E(\text{path}_{R'}(u, v))|$, then we output u, v, P as a compact witness of a detour in R' . Else, we move on to the next choice of u and v . If we fail to find a witness for all choices of u and v , then we output that R' has no detour. Observe that the total running time of this process is bounded by $\mathcal{O}(k^2 \cdot n)$. This concludes the proof. \square

Accordingly, as long as R^1 has a detour, compactly witnessed by some vertices u, v and a path P , we modify it as follows: we remove the edges and the internal vertices of $\text{path}_{R^1}(u, v)$, and add the edges and the internal vertices of P . We refer to a single application of this operation as *undetouring* R^1 . For a single application, because we consider compact witnesses rather than arbitrary ones, we have the following observation.

Observation 6.3. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths with a Steiner tree R' . The result of undetouring R' is another Steiner tree with fewer edges than R' .*

Proof. Consider a compact detour witness u, v, P of R' . Then, $|E(P)| < |E(\text{path}_{R'}(u, v))|$. Let $Q = \text{path}_{R'}(u, v) - \{u, v\}$, and let R'_u and R'_v be the two trees of $R' - V(Q)$ that contain u and v , respectively. Consider the graph \tilde{R} obtained from $(R' - V(Q)) \cup P$ by iteratively removing any leaf vertex that does not lie in $S \cup T$. We claim that the graph \tilde{R} that result from undetouring R' (with respect to u, v, P) is a Steiner tree with strictly fewer edges than R' . Clearly, \tilde{R} is connected because P reconnects the two trees R'_u and R'_v of $R' - V(Q)$. Further, as P contains no internal vertex from $(V(R') \setminus V(\text{path}_{R'}(u, v))) \cup \{u, v\} = V(R'_u) \cup V(R'_v)$, and R'_u and R'_v are trees, \tilde{R} is cycle-free. Additionally, all the vertices in $S \cup T$ are present in \tilde{R} by construction and they remain leaves due to the compactness of the witness. Hence, \tilde{R} is a Steiner tree in G . Because $|E(P)| < |E(\text{path}_{R'}(u, v))|$, it follows that \tilde{R} contains fewer edges than R' . \square

Initially, R^1 has at most $n - 1$ edges. Since every iteration decreases the number of edges (by Observation 6.3) and can be performed in time $\mathcal{O}(k^2 \cdot n)$ (by Lemma 6.2), we obtain the following result.

Lemma 6.3. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths with a Steiner tree R' . An exhaustive application of the operation undetouring R' can be performed in time $\mathcal{O}(k^2 \cdot n^2)$, and results in a Steiner tree that has no detour.*

We denote the Steiner tree obtained at the end of Step II by R^2 .

6.3 Step III: Small Separators for Long Paths

We now show that any two parts of R^2 that are “far” from each other can be separated by small separators in H . This is an important property used in the following sections to show the existence of a “nice” solution for the input instance. Specifically, we consider a “long” maximal degree-2 path in R^2 (which has no short detours in H), and show that there are two separators of small cardinality, each “close” to one end-point of the path. The main idea behind the proof of this result is that, if it were false, then the graph H would have had large treewidth (see




Figure 14: Separators and flows for long degree-2 paths.

Proposition 6.2), which contradicts that H has bounded treewidth (by Corollary 4.1). We first define the threshold that determines whether a path is long or short.

Definition 6.2 (Long Paths in Trees). *Let G be a graph with a subtree T . A subpath of T is k -long if its length is at least $\alpha_{\text{long}}(k) := 10^4 \cdot 2^{ck}$, and k -short otherwise.*

As k will be clear from context, we simply use the terms long and short. Towards the computation of two separators for each long path, we also need to define which subsets of $V(R^2)$ we would like to separate.

Definition 6.3 ($P'_u, P''_u, A_{R^2, P, u}$ and $B_{R^2, P, u}$). *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 be a Steiner tree that has no detour. For any long maximal degree-2 path P of R^2 and for each endpoint u of P , define P'_u, P''_u and $A_{R^2, P, u}, B_{R^2, P, u} \subseteq V(R^2)$ as follows.*

- P'_u (resp. P''_u) is the subpath of P consisting of the $\alpha_{\text{pat}}(k) := 100 \cdot 2^{ck}$ (resp. $\alpha_{\text{pat}}(k)/2 = 50 \cdot 2^{ck}$) vertices of P closest to u .
- $A_{R^2, P, u}$ is the union of $V(P''_u)$ and the vertex set of the connected component of $R^2 - (V(P'_u) \setminus \{u\})$ containing u .
- $B_{R^2, P, u} = V(R^2) \setminus (A_{R^2, P, u} \cup V(P'_u))$.

For each long maximal degree-2 path P of R^2 and for each endpoint u of P , we compute a “small” separator $\text{Sep}_{R^2}(P, u)$ as follows. Let $A = A_{R^2, P, u}$ and $B = B_{R^2, P, u}$. Then, compute a subset of $V(H) \setminus (A \cup B)$ of minimum size that separates A and B in H , and denote it by $\text{Sep}_{R^2}(P, u)$ (see Fig. 14). Since $A \cap B = \emptyset$ and there is no edge between a vertex in A and a vertex in B (because R^2 has no detours), such a separator exists. Moreover, it can be computed in time $\mathcal{O}(n|\text{Sep}_{R^2}(P, u)|)$: contract each set among A and B into a single vertex and then obtain a minimum vertex $s - t$ cut by using Ford-Fulkerson algorithm.

To argue that the size of $\text{Sep}_{R^2}(P, u)$ is upper bounded by $2^{\mathcal{O}(k)}$, we make use of the following proposition due to Bodlaender et al. [5].

Proposition 6.2 (Lemma 6.11 in [5]). *Let G be a plane graph, and let H be its radial completion. Let $t \in \mathbb{N}$. Let C, Z, C_1, Z_1 be disjoint subsets of $V(H)$ such that*

1. $H[C]$ and $H[C_1]$ are connected graphs,
2. Z separates C from $Z_1 \cup C_1$ and Z_1 separates $C \cup Z$ from C_1 in H ,
3. $\text{dist}_H(Z, Z_1) \geq 3t + 4$, and
4. G contains $t + 2$ pairwise internally vertex-disjoint paths with one endpoint in $C \cap V(G)$ and the other endpoint in $C_1 \cap V(G)$.

Then, the treewidth of $G[V(M) \cap V(G)]$ is larger than t where M is the union of all connected components of $H \setminus (Z \cup Z_1)$ having at least one neighbor in Z and at least one neighbor in Z_1 .

Additionally, the following immediate observation will come in handy.

Observation 6.4. *Let G be a plane graph. Let H be the radial completion of G , and let H' be the radial completion of H . Then, for all $u, v \in V(H)$, $\text{dist}_H(u, v) \leq \text{dist}_{H'}(u, v)$.*

Proof. Note that H is triangulated. Thus, for all $u, v \in V(H)$ and path P in H' between u and v , we can obtain a path between u and v whose length is not longer than the length of P by replacing each vertex $w \in V(H') \setminus V(H)$ by at most one vertex of the boundary of the face in H that w represents. \square

We now argue that $\text{Sep}_{R^2}(P, u)$ is small.

Lemma 6.4. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 be a Steiner tree that has no detour, P be a long maximal degree-2 path of R^2 , and u be an endpoint of P . Then, $|\text{Sep}_{R^2}(P, u)| \leq \alpha_{\text{sep}}(k) := \frac{7}{2} \cdot 2^{ck} + 2$.*

Proof. Denote $P' = P'_u, P'' = P''_u, A = A_{R^2, P, u}$ and $B = B_{R^2, P, u}$. Recall that H is the radial completion of G (enriched with parallel edges), and let H' denote the radial completion of H . Towards an application of Proposition 6.2, define $C = A, C_1 = B, Z = N_{H'}(C), Z_1 = N_{H'}(C_1)$ and $t = \frac{7}{2} \cdot 2^{ck}$. Since R^2 is a subtree of H , it holds that $H[C]$ and $H[C_1]$ are connected, and therefore $H'[C]$ and $H'[C_1]$ are connected as well. From the definition of Z and Z_1 , it is immediate that Z separates C from $Z_1 \cup C_1$ and Z_1 separates $C \cup Z$ from C_1 in H . Clearly, $C \cap C_1 = \emptyset, C \cap Z = \emptyset$, and $C_1 \cap Z_1 = \emptyset$. We claim that, in addition, $Z \cap C_1 = \emptyset, Z_1 \cap C = \emptyset$ and $Z \cap Z_1 = \emptyset$. To this end, it suffices to show that $\text{dist}_{H'}(Z, Z_1) \geq 3t + 4$. Indeed, because $Z = N_{H'}(C)$ and $Z_1 = N_{H'}(C_1)$, we have that each inequality among $Z \cap C_1 \neq \emptyset, Z_1 \cap C = \emptyset$ and $Z \cap Z_1 = \emptyset$, implies that $\text{dist}_{H'}(Z, Z_1) \leq 2$.

Lastly, we show that $\text{dist}_{H'}(Z, Z_1) \geq 3t + 4$. As $\text{dist}_{H'}(C, C_1) \leq \text{dist}_H(Z, Z_1) + 2$, it suffices to show that $\text{dist}_{H'}(C, C_1) \geq 3t + 6$. Because $C \cup C_1 \subseteq V(H)$, Observation 6.4 implies that $\text{dist}_{H'}(C, C_1) \geq \text{dist}_H(C, C_1)$. Hence, it suffices to show that $\text{dist}_H(C, C_1) \geq 3t + 6$. However, $\text{dist}_H(C, C_1) \geq |E(P')| - |E(P'')|$ since otherwise we obtain a contradiction to the supposition that R^2 has no detour. This means that $\text{dist}_H(C, C_1) \geq \alpha_{\text{pat}}(k)/2 - 1 \geq 3t + 6$ as required.

Recall that $\text{Sep}_{R^2}(P, u)$ is a subset of $V(H) \setminus (C \cup C_1)$ of minimum size that separates C and C_1 in H . We claim that $|\text{Sep}_{R^2}(P, u)| \leq \alpha_{\text{sep}}(k)$. Suppose, by way of contradiction, that $|\text{Sep}_{R^2}(P, u)| > \alpha_{\text{sep}}(k) = t + 2$. By Menger's theorem, the inequality $|\text{Sep}_{R^2}(P, u)| > \alpha_{\text{sep}}(k)$ implies that H contains $t + 2$ pairwise internally vertex-disjoint paths with one endpoint in $C \subseteq V(H)$ and the other endpoint in $C_1 \subseteq V(H)$. From this, we conclude that all of the conditions in the premise of Proposition 6.2 are satisfied. Thus, the treewidth of $H[V(M) \cap V(H)]$ is larger than t where M is the union of all connected components of $H' \setminus (Z \cup Z_1)$ having at least one neighbor in Z and at least one neighbor in Z_1 . However, $H[V(M) \cap V(H)]$ is a subgraph of H , which means that the treewidth of H is also larger than t . By Proposition 3.2, this implies that the treewidth of G is larger than 2^{ck} . This contradicts the supposition that (G, S, T, g, k) is good. From this, we conclude that $|\text{Sep}_{R^2}(P, u)| \leq \alpha_{\text{sep}}(k)$. \square

Recall that $\text{Sep}_{R^2}(P, u)$ is computable in time $\mathcal{O}(n|\text{Sep}_{R^2}(P, u)|)$. Thus, by Lemma 6.4, we obtain the observation below. We remark that the reason we had to argue that the separator is small is not due to this observation, but because the size bound will be crucial in later sections.

Observation 6.5. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 be a Steiner tree that has no detour, P be a long maximal degree-2 path of R^2 , and u be an endpoint of P . Then, $\text{Sep}_{R^2}(P, u)$ can be computed in time $2^{\mathcal{O}(k)}n$.*

Moreover, we have the following immediate consequence of Proposition 3.1.

Observation 6.6. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 be a Steiner tree that has no detour, P be a long maximal degree-2 path of R^2 , and u be an endpoint of P . Then, $H[\text{Sep}_{R^2}(P, u)]$ is a cycle.*

6.4 Step IV: Internal Modification of Long Paths

In this step, we replace the “middle” of each long maximal degree-2 path $P = \text{path}_{R^2}(u, v)$ of R^2 by a different path P^* . This “middle” is defined by the two separators obtained in the previous step. Let us informally explain the reason behind this modification. In Section 7 we will show that, if the given instance (G, S, T, g, k) admits a solution (which is a collection of disjoint paths connecting S and T), then it also admits a “nice” solution that “spirals” only a few times around parts of the constructed Steiner tree. This requirement is crucial, since it is only such solutions \mathcal{P} that are discretely homotopic to weak linkages \mathcal{W} in H aligned with \mathcal{P} that use at most $2^{O(k)}$ edges parallel to those in R , and none of the edges not parallel to those in R . To ensure the existence of nice solutions, we show how an arbitrary solution can be rerouted to avoid too many spirals. This rerouting requires a collection of vertex-disjoint paths between $\text{Sep}_{R^2}(P, u)$ and $\text{Sep}_{R^2}(P, v)$ which itself does not spiral around the Steiner tree. The replacement of P by P^* in the Steiner tree, described below, will ensure this property.

To describe this modification, we first need to assert the statement in the following simple lemma, which partitions every long maximal degree-2 path P of R^2 into three parts (see Fig. 14).

Lemma 6.5. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 be a Steiner tree with no detour, and P be a long maximal degree-2 path of R^2 with endpoints u and v . Then, there exist vertices $u' = u'_P \in \text{Sep}_{R^2}(P, u) \cap V(P)$ and $v' = v'_P \in \text{Sep}_{R^2}(P, v) \cap V(P)$ such that:*

1. *The subpath $P_{u,u'}$ of P with endpoints u and u' has no internal vertex from $\text{Sep}_{R^2}(P, u) \cup \text{Sep}_{R^2}(P, v)$, and $\alpha_{\text{pat}}(k)/2 \leq |V(P_{u,u'})| \leq \alpha_{\text{pat}}(k)$. Additionally, the subpath $P_{v,v'}$ of P with endpoints v and v' has no internal vertex from $\text{Sep}_{R^2}(P, u) \cup \text{Sep}_{R^2}(P, v)$, and $\alpha_{\text{pat}}(k)/2 \leq |V(P_{v,v'})| \leq \alpha_{\text{pat}}(k)$.*
2. *Let $P_{u',v'}$ be the subpath of P with endpoints u' and v' . Then, $P = P_{u,u'} - P_{u',v'} - P_{v',v}$.*

Proof. We first prove that there exists a vertex $u' \in \text{Sep}_{R^2}(P, u) \cap V(P)$ such that the subpath $P_{u,u'}$ of P between u and u' has no internal vertex from $\text{Sep}_{R^2}(P, u) \cup \text{Sep}_{R^2}(P, v)$. To this end, let $P' = P'_u$, and let \tilde{P} denote the subpath of P that consists of the $\alpha_{\text{pat}}(k) + 1$ vertices of P that are closest to u . Let $A = A_{R^2, P, u}$ and $B = B_{R^2, P, u}$. Recall that $\text{Sep}_{R^2}(P, u) \subseteq V(H) \setminus (A \cup B)$ separates A and B in H . Since \tilde{P} is a path with the endpoint u in A and the other endpoint in B , it follows that $\text{Sep}_{R^2}(P, u) \cap V(\tilde{P}) \neq \emptyset$. Accordingly, let u' denote the vertex of P' closest to u that belongs to $\text{Sep}_{R^2}(P, u)$. Then, $u' \in \text{Sep}_{R^2}(P, u) \cap V(P)$ and the subpath $P_{u,u'}$ of P between u and u' has no internal vertex from $\text{Sep}_{R^2}(P, u)$. As the number of vertices of $P_{u,u'}$ is between those of P''_u and P' , the inequalities $\alpha_{\text{pat}}(k)/2 \leq |V(P_{u,u'})| \leq \alpha_{\text{pat}}(k)$ follow. It remains to argue that $P_{u,u'}$ has no internal vertex from $\text{Sep}_{R^2}(P, v)$. Because $\text{Sep}_{R^2}(P, v) \subseteq V(H) \setminus (A_{R^2, P, v} \cup B_{R^2, P, v})$, the only vertices of P that $\text{Sep}_{R^2}(P, v)$ can possibly contain are the $\alpha_{\text{pat}}(k)$ vertices of P that are closest to v . Since P is long, none of these vertices belongs to P' , and hence $P_{u,u'}$ (which is a subpath of P') has no internal vertex from $\text{Sep}_{R^2}(P, v)$.

Symmetrically, we derive the existence of a vertex $v' \in \text{Sep}_{R^2}(P, v) \cap V(P)$ such that the subpath $P_{v,v'}$ of P between v and v' has no internal vertex from $\text{Sep}_{R^2}(P, u) \cup \text{Sep}_{R^2}(P, v)$.

Lastly, we prove that $P = P_{u,u'} - P_{u',v'} - P_{v',v}$. Since $P_{u,u'}$, $P_{u',v'}$ and $P_{v',v}$ are subpaths of P such that $V(P) = V(P_{u,u'}) \cup V(P_{u',v'}) \cup V(P_{v',v})$, it suffices to show that (i) $V(P_{u,u'}) \cap V(P_{u',v'}) = \{u'\}$, (ii) $V(P_{v,v'}) \cap V(P_{u',v'}) = \{v'\}$, and (iii) $V(P_{u,u'}) \cap V(P_{v,v'}) = \emptyset$. Because P is long and $|V(P_{u,u'})|, |V(P_{v,v'})| \leq \alpha_{\text{pat}}(k)$, it is immediate that item (iii) holds. For item (i), note that $V(P_{u,u'}) \cap V(P_{u',v'})$ can be a strict superset of $\{u'\}$ only if $P_{u',v'}$ is a subpath of $P_{u,u'}$; then, $v' \in V(P_{u,u'})$, which means that $P_{u,u'}$ has an internal vertex from $\text{Sep}_{R^2}(P, v)$ and results in a contradiction. Thus, item (i) holds. Symmetrically, item (ii) holds as well. \square

In what follows, when we use the notation u'_P , we refer to the vertex in Lemma 6.5. Before we describe the modification, we need to introduce another notation and make an immediate observation based on this notation.

Definition 6.4 ($\tilde{A}_{R^2, P, u}$). *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 be a Steiner tree that has no detour, P be a long maximal degree-2 path of R^2 , and u be an endpoint of P . Then, $\tilde{A}_{R^2, P, u} = (V(P_{u, u'_P}) \setminus \{u'_P\}) \cup A_{R^2, p, u}$.*

Observation 6.7. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 be a Steiner tree that has no detour, and P be a long maximal degree-2 path of R^2 with endpoints u and v . Then, there exists a single connected component $C_{R^2, P, u}$ in $H - (\text{Sep}_{R^2}(u, P) \cup \text{Sep}_{R^2}(v, P))$ that contains $\tilde{A}_{R^2, P, u}$ and a different single connected component $C_{R^2, P, v}$ in $H - (\text{Sep}_{R^2}(u, P) \cup \text{Sep}_{R^2}(v, P))$ that contains $\tilde{A}_{R^2, P, v}$.*

We proceed to describe the modification. For brevity, let $S_u = \text{Sep}_{R^2}(P, u)$ and $S_v = \text{Sep}_{R^2}(P, v)$. Recall that there is a terminal $t^* \in T$ that lies on the outer face of H (and G). By Observation 6.6, S_u and S_v induce two cycles in H , and t^* lies in the exterior of both these cycles. Assume w.l.o.g. that u lies in the interior of both S_u and S_v , while v lies in the exterior of both S_u and S_v . Then, S_u belongs to the strict interior of S_v . We construct a sequence of concentric cycles between S_u and S_v as follows.

Lemma 6.6. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 be a Steiner tree with no detour, and P be a long maximal degree-2 path of R^2 with endpoints u and v . Let $S_u = \text{Sep}_{R^2}(P, u)$ and $S_v = \text{Sep}_{R^2}(P, v)$, where S_u lies in the strict interior of S_v . Then, there is a sequence of concentric cycles $\mathcal{C}(u, v) = (C_1, C_2, \dots, C_p)$ in G of length $p \geq 100\alpha_{\text{sep}}(k)$ such that S_u is in the strict interior of C_1 in H , S_v is in the strict exterior of C_p in H , and there is a path η in H with one endpoint $v_0 \in S_u$ and the other endpoint $v_{p+1} \in S_v$, such that the intersection of $V(\eta)$ with $V(G) \cup S_u \cup S_v$ is $\{v_0, v_1, \dots, v_{p+1}\}$ for some $v_i \in V(C_i)$ for every $i \in \{1, \dots, p\}$. Furthermore, $\mathcal{C}(u, v)$ can be computed in linear time.*

Proof. Towards the computation of $\mathcal{C}(u, v)$, delete all vertices that lie in the strict interior of S_u or in the strict exterior of S_v , as well as all vertices of $V(H) \setminus (V(G) \cup S_u \cup S_v)$. Denote the resulting graph by $G_{u, v}^+$, and note that it has a plane embedding in the “ring” defined by $H[S_u]$ and $H[S_v]$. Observe that $S_u, S_v \subseteq V(G_{u, v}^+)$, where S_v defines the outer face of the embedding of $G_{u, v}^+$. Thus, any cycle in this graph that separates S_u and S_v must contain S_u in its interior and S_v in its exterior. Furthermore, $G_{u, v} = G_{u, v}^+ - V(H)$ is an induced subgraph of G , which consists of all vertices of G that, in H , lie in the strict exterior of S_u and in the strict interior of S_v simultaneously, or lie in $S_u \cup S_v$. In particular, any cycle of $G_{u, v}$ is also a cycle in G .

Now, $\mathcal{C}(u, v)$ is computed as follows. Start with an empty sequence, and the graph $G_{u, v} - (S_u \cup S_v)$. As long as there is a cycle in the current graph such that all vertices of S_u are in the strict interior of C with respect to H , remove vertices of degree at most 1 in the current graph until no such vertices remain, and append the outer face of the current graph as a cycle to the constructed sequence. It is clear that this process terminates in linear time, and that by the above discussion, it constructs a sequence of concentric cycles $\mathcal{C}(u, v) = (C_1, C_2, \dots, C_p)$ in G such that S_u is in the strict interior of C_1 in H , S_v is in the strict exterior of C_p in H .

To assert the existence of a path η in H with one endpoint $v_0 \in S_u$ and the other endpoint $v_{p+1} \in S_v$, such that the intersection of $V(\eta)$ with $V(G) \cup S_u \cup S_v$ is $\{v_0, v_1, \dots, v_{p+1}\}$ for some $v_i \in V(C_i)$ for every $i \in \{1, \dots, p\}$, we require the following claim.

Claim 6.1. *Let $C_{p+1} = H[S_v]$. For every $i \in \{1, 2, \dots, p\}$ and every vertex $w \in V(C_i)$, there exists a vertex $w' \in V(C_{i+1})$ such that w and w' lie on a common face in $G_{u, v}^+$. Moreover, there exist vertices $w \in S_u$ and $w' \in V(C_1)$ that lie on a common face in $G_{u, v}^+$.*

Proof of Claim 6.1. Consider $i \in \{1, 2, \dots, p\}$ and a vertex $w \in V(C_i)$. We claim that $\text{rdist}(w, V(C_{i+1})) \leq 1$, i.e. there must be $w' \in V(C_{i+1})$ such that w, w' have a common face in $G_{u,v}^+$. By way of contradiction, suppose that $\text{rdist}(w, V(C_{i+1})) \geq 2$. Then, by Proposition 6.1, there is a cycle C that separates w and $V(C_{i+1})$ in $G_{u,v}^+$ such that $\text{rdist}(w, w'') = 1$ for every vertex $w'' \in V(C)$. Here, w lies in the strict interior of C , and C lies in the strict interior of C_{i+1} . Further, C is vertex disjoint from C_{i+1} , since $\text{rdist}(w, w') \geq 2$ for every $w' \in V(C_{i+1})$. Now, consider the outer face of $G[V(C_i) \cup V(C)]$. By the construction of C_i , this outer face must be C_i . However, $w \in V(C_i)$ cannot belong to it, hence we reach a contradiction.

For the second part, we claim that $\text{rdist}(S_u, V(C_1)) \leq 1$, i.e. there must be $w \in S_u$ and $w' \in V(C_1)$ such that w, w' have a common face in $G_{u,v}^+$. By way of contradiction, suppose that $\text{rdist}(S_u, V(C_1)) \geq 2$. Then, by Proposition 6.1, there is a cycle C that separates S_u and $V(C_1)$ in $G_{u,v}^+$ such that $\text{rdist}(S_u, w'') = 1$ for every vertex $w'' \in V(C)$. Further, C is vertex disjoint from C_1 , since $\text{rdist}(S_u, w') \geq 2$ for every $w' \in V(C_1)$. However, this is a contradiction to the termination condition of the construction of $\mathcal{C}(u, v)$. \diamond

Having this claim, we construct η as follows. Pick vertices $v_0 \in S_u$ and $v_1 \in V(C_1)$ that lie on a common face in $G_{u,v}^+$. Then, for every $i \in \{2, \dots, p+1\}$, pick a vertex $v_i \in V(C_i)$ such that v_{i-1} and v_i lie on a common face in $G_{u,v}^+$. Thus, for every $i \in \{0, 1, \dots, p\}$, we have that v_i and v_{i+1} are either adjacent in H or there exists a vertex $u_i \in V(H) \setminus V(G)$ such that u_i is adjacent to both v_i and v_{i+1} . Because $\mathcal{C}(u, v) = (C_1, C_2, \dots, C_p)$ is a sequence of concentric cycles in G such that S_u is in the strict interior of C_1 and S_v is in the strict exterior of C_p , the u_i 's are distinct. Thus, $\eta = v_0 - u_0 - v_1 - u_1 - v_2 - u_2 - \dots - v_p - u_p - v_{p+1}$, where undefined u_i 's are dropped, is a path as required.

Finally, we argue that $p \geq 100 \cdot \alpha_{\text{sep}}(k)$. Note that $100\alpha_{\text{sep}}(k) = 100(\frac{7}{2} \cdot 2^{ck} + 2) \leq 400 \cdot 2^{ck}$, thus it suffices to show that $p \geq 400 \cdot 2^{ck}$. To this end, we obtain a lower bound on the radial distance between S_u and S_v in $G_{u,v}^+$. Recall that $|S_u|, |S_v| \leq \alpha_{\text{sep}}(k) = \frac{7}{2} \cdot 2^{ck} + 2 \leq 4 \cdot 2^{ck}$. Let $P = \text{path}_{R^2}(u, v)$, and recall that its length is at least $\alpha_{\text{long}}(k) = 10^4 2^{ck}$. Since R^2 has not detour, P is a shortest path in H between u and v , thus for any two vertices in $V(P)$, the subpath of P between them is a shortest path between them. Now, recall the vertices $u' = u'_p, v' = v'_p$ (defined in Lemma 6.5), and denote the subpath between them by P' . By construction, $|E(P')| \geq \alpha_{\text{long}}(k) - 2 \cdot \alpha_{\text{pat}}(k) = (10^4 - 200) \cdot 2^{ck}$.

We claim that the radial distance between S_u and S_v in $G_{u,v}^+$ is at least $|E(P')|/2 - |S_u| - |S_v|$. Suppose not, and consider a sequence of vertices in $G_{u,v}^+$ that witnesses this fact: $x_1, x_2, x_3, \dots, x_{p-1}, x_p \in V(G_{u,v}^+)$ where $x_1 \in S_u, x_p \in S_v, p < |E(P')|/2 - |S_v| - |S_u|$, and every two consecutive vertices lie on a common face. Consider a shortest such sequence, which visits each face of $G_{u,v}^+$ at most once. In particular, x_1 and x_p are the only vertices of $S_u \cup S_v$ in this sequence. Then, we can extend this sequence on both sides to derive another sequence of vertices of $G_{u,v}^+$ starting at u' and ending at v' such that the prefix of the new sequence is a path in $G_{u,v}^+[S_u]$ from u' to x_1 , the midfix is $x_1, x_2, x_3, \dots, x_{p-1}, x_p$, and the suffix is a path in $G_{u,v}^+[S_v]$ from x_p to v' . Further, the length of the new sequence of vertices is smaller than $|E(P')|/2$. Hence, the radial distance between u' and v' in $H' = H[V(G) \cup S_u \cup S_v]$ (the graph derived from $G_{u,v}^+$ by reintroducing the vertices of G that lie inside S_u or outside S_v) is smaller than $|E(P')|/2$, and let it be witnessed by a sequence $Q[u', v']$. As H is the radial completion of G , observe that $Q[u', v']$ gives rise to a path Q in H between u' and v' of length smaller than $|E(P')|$. However, then u', v' and Q witness a detour in R^2 , which is a contradiction. Hence, the radial distance between S_u and S_v in $G_{u,v}^+$ is at least

$$\begin{aligned} & \frac{(10^4 - 200)}{2} \cdot 2^{ck} - (|S_u| + |S_v|) \\ & \geq 4900 \cdot 2^{ck} - 2\alpha_{\text{sep}}(k) \\ & = 4900 \cdot 2^{ck} - (7 \cdot 2^{ck} + 4) \geq 400 \cdot 2^{ck}. \end{aligned}$$

Now, observe that S_u and S_v are connected sets in $G_{u,v}^+$ and S_v forms the outer-face of $G_{u,v}^+$.

Then, by Proposition 6.1, we obtain a collection of at least $400 \cdot 2^{ck}$ disjoint cycles in $G_{u,v}^+$, where each cycle separates S_u and S_v . Note that these cycles are disjoint from $S_u \cup S_v$, and hence they lie in G . Moreover, each of them contains S_u in its strict interior, and S_v in its strict exterior. Thus, it is clear that the sequence $\mathcal{C}(u, v)$ computed above must contain at least $400 \cdot 2^{ck} > 100\alpha_{\text{sep}}(k)$ cycles. \square

Recall the graph $G_{u,v}$, which is an induced subgraph of G , which consists of all vertices of G that, in H , lie in the strict exterior of S_u and in the strict interior of S_v simultaneously, or lie in $S_v \cup S_v$. With $\mathcal{C}(u, v)$ at hand, we compute a maximum size collection of disjoint paths from S_u to S_v in $G_{u,v}$ that minimizes the number of edges it traverses outside $E(\mathcal{C}(u, v))$. In this observation, the implicit assumption that $\ell \leq \alpha_{\text{sep}}(k)$ is justified by Lemma 6.4.

Observation 6.8. *Let the maximum flow between $S_u \cap V(G)$ and $S_v \cap V(G)$ in $G_{u,v}$ be $\ell \leq \alpha_{\text{sep}}(k)$. Given the sequence $\mathcal{C}(u, v)$ of Lemma 6.6, a collection $\text{Flow}_{R^2}(u, v)$ of ℓ vertex-disjoint paths in $G_{u,v}$ from $S_u \cap V(G)$ to $S_v \cap V(G)$ that minimizes $|E(\text{Flow}_{R^2}(u, v)) \setminus E(\mathcal{C}(u, v))|$ is computable in time $\mathcal{O}(n^{3/2} \log^3 n)$.*

Proof. We determine $\ell \leq \alpha_{\text{sep}}(k)$ in time $2^{\mathcal{O}(k)}n$ by using Ford-Fulkerson algorithm. Next, we define a weight function w on $E(G_{u,v})$ as follows:

$$w(e) = \begin{cases} 0 & \text{if } e \in E(\mathcal{C}(u, v)) \\ 1 & \text{otherwise} \end{cases}$$

We now compute a minimum cost flow between $S_u \cap V(G)$ and $S_v \cap V(G)$ of value ℓ in $G_{u,v}$ under the weight function w . This can be done in time $\mathcal{O}(n^{3/2} \log^3 n)$ by [13, Theorem 1], as the cost of such a flow is bounded by $\mathcal{O}(n)$. Clearly, the result is a collection $\text{Flow}_{R^2}(u, v)$ of ℓ vertex-disjoint paths from $S_u \cap V(G)$ to $S_v \in V(G)$ minimizing $|E(\text{Flow}_{R^2}(u, v)) \setminus E(\mathcal{C}(u, v))|$. \square

Having $\text{Flow}_{R^2}(u, v)$ at hand, we proceed to find a certain path between u'_P and v'_P that will be used to replace $P_{u'_P, v'_P}$. We remark that all vertices and edges of $\text{Flow}_{R^2}(u, v)$ lie between S_u and S_v in the plane embedding of H . The definition of this path is given by the following lemma, and the construction will make it intuitively clear that the paths in $\text{Flow}_{R^2}(u, v)$ do not “spiral around” the Steiner tree once we replace $P_{u'_P, v'_P}$ with $P_{u'_P, v'_P}^*$. Note that in the lemma, we consider a path P in H , while the paths in $\text{Flow}_{R^2}(u, v)$ are in G .

Lemma 6.7. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 be a Steiner tree with no detour, and P be a long maximal degree-2 path of R^2 , with endpoints u and v . Let $u' = u'_P$ and $v' = v'_P$. Then, there exists a path $P_{u', v'}^*$ in $H - (V(C_{R^2, P, u}) \cup V(C_{R^2, P, u}))$ between u' and v' with the following property: there do not exist three vertices $x, y, z \in V(P_{u', v'}^*)$ such that (i) $\text{dist}_{P_{u', v'}^*}(u', x) < \text{dist}_{P_{u', v'}^*}(u', y) < \text{dist}_{P_{u', v'}^*}(u', z)$, and (ii) there exist a path in $\text{Flow}_{R^2}(u, v)$ that contains x and z and a different path in $\text{Flow}_{R^2}(u, v)$ that contains y . Moreover, such a path $P_{u', v'}^*$ can be computed in time $\mathcal{O}(n)$.*

Proof. Let $P_{u', v'}^*$ be a path in $H - (V(C_{R^2, P, u}) \cup V(C_{R^2, P, u}))$ between u' and v' that minimizes the number of paths $Q \in \text{Flow}_{R^2}(u, v)$ for which there exist at least one triple $x, y, z \in V(P_{u', v'}^*)$ that has the two properties in the lemma and $x, z \in V(Q)$. Due to the existence of $P_{u', v'}$, such a path $P_{u', v'}^*$ exists. We claim that this path $P_{u', v'}^*$ has no triple $x, y, z \in V(P_{u', v'}^*)$ that has the two properties in the lemma. Suppose, by way of contradiction, that our claim is false, and let $x, y, z \in V(P_{u', v'}^*)$ be a triple that has the two properties in the lemma. Note that, when traversed from u' to v' , $P_{u', v'}^*$ first visits x , then visits y and afterwards visits z . Let Q' be the path in $\text{Flow}_{R^2}(u, v)$ that contains x and z . Let x' and z' be the first and last vertices of Q' that are visited by $P_{u', v'}^*$. Then, replace the subpath of $P_{u', v'}^*$ between x' and z' by the subpath of Q' between x' and z' . This way we obtain a path P' in $H - (V(C_{R^2, P, u}) \cup V(C_{R^2, P, u}))$ between u'

and v' for which there exist fewer paths $Q \in \text{Flow}_{R^2}(u, v)$, when compared to $P_{u', v'}^*$, for which there exists at least one triple $x, y, z \in V(P')$ that has the two properties in the lemma and such that $x, z \in V(Q)$. As we have reached a contradiction, we conclude that our initial claim is correct. While the proof is existential, it can clearly be turned into a linear-time algorithm. \square

The following is a direct corollary of the above lemma.

Corollary 6.1. *For each path $Q \in \text{Flow}_{R^2}(u, v)$, there are at most two edges in $P_{u', v'}^*$ such that one endpoint of the edge lies in $V(Q)$ and the other lies in $V(G) \setminus V(Q)$.*

Having Lemma 6.7 at hand, we modify R^2 as follows: for every long maximal degree-2 path P of R^2 with endpoints u and v , replace $P_{u', v'}$ by $P_{u', v'}^*$. Denote the result of this modification by $R = R^3$. We refer to R^3 as a *backbone Steiner tree*. Let us remark that the backbone Steiner tree R^3 is always accompanied by the separators $\text{Sep}_{R^2}(P, u)$ and $\text{Sep}_{R^2}(P, v)$, and the collection $\text{Flow}_{R^2}(u, v)$ for every long maximal degree-2 path $P = \text{path}_{R^2}(u, v)$ of R^2 . These separators and flows will play crucial role in our algorithm. In the following subsection, we will prove that R^3 is indeed a Steiner tree, and in particular it is a tree. Additionally and crucially, we will prove that the separators computed previously remain separators. Let us first conclude the computational part by stating the running time spent so far. From Lemma 6.3, Observations 6.5 and 6.8, and by Lemma 6.7, we have the following result.

Lemma 6.8. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Then, a backbone Steiner tree R can be computed in time $2^{\mathcal{O}(k)} n^{3/2} \log^3 n$.*

6.5 Analysis of R^3 and the Separators $\text{Sep}_{R^2}(P, u)$

Having constructed the backbone Steiner tree R^3 , we turn to analyse its properties. Among other properties, we show that useful properties of R^2 also transfer to R^3 . We begin by proving that the two separators of each long maximal degree-2 path P of R^2 partition $V(H)$ into five “regions”, and that the vertices in each region are all close to the subtree of R that (roughly) belongs to that region. Specifically, the regions are $V(C_{R^2, P, u})$, $\text{Sep}_{R^2}(P, u)$, $V(C_{R^2, P, v})$, $\text{Sep}_{R^2}(P, v)$, and $V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v}) \cup \text{Sep}_{R^2}(P, u) \cup \text{Sep}_{R^2}(P, v))$, and our claim is as follows.

Lemma 6.9. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 be a Steiner tree that has no detour, P be a long maximal degree-2 path of R^2 , and u and v be its endpoints.*

1. For all $w \in \text{Sep}_{R^2}(P, u)$, it holds that $\text{dist}_H(w, u'_P) \leq \alpha_{\text{sep}}(k)$.
2. For all $w \in \text{Sep}_{R^2}(P, v)$, it holds that $\text{dist}_H(w, v'_P) \leq \alpha_{\text{sep}}(k)$.
3. For all $w \in V(C_{R^2, P, u})$, it holds that $\text{dist}_H(w, \tilde{A}_{R^2, P, u} \cup \{u'_P\}) \leq \alpha_{\text{dist}}(k) + \alpha_{\text{sep}}(k)$.
4. For all $w \in V(C_{R^2, P, v})$, it holds that $\text{dist}_H(w, \tilde{A}_{R^2, P, v} \cup \{v'_P\}) \leq \alpha_{\text{dist}}(k) + \alpha_{\text{sep}}(k)$.
5. For all $w \in V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v}) \cup \text{Sep}_{R^2}(P, u) \cup \text{Sep}_{R^2}(P, v))$, it holds that $\text{dist}_H(w, V(R^2) \setminus (\tilde{A}_{R^2, P, u} \cup \tilde{A}_{R^2, P, v})) \leq \alpha_{\text{dist}}(k) + \alpha_{\text{sep}}(k)$.

Proof. First, note that Conditions 1 and 2 follow directly from Lemma 6.4 and Observation 6.6.

For Condition 3, consider some vertex $w \in V(C_{R^2, P, u})$. By Lemma 6.1, $\text{dist}_H(w, V(R^2)) \leq \alpha_{\text{dist}}(k)$. Thus, there exists a path Q in H with w as one endpoint and the other endpoint x in $V(R^2)$ such that the length of Q is at most $\alpha_{\text{dist}}(k)$. In case $x \in \tilde{A}_{R^2, P, u} \cup \{u'_P\}$, we have that $\text{dist}_H(w, \tilde{A}_{R^2, P, u}) \leq \alpha_{\text{dist}}(k)$, and hence the condition holds. Otherwise, by the definition of $\text{Sep}_{R^2}(P, u)$, the path Q must traverse at least one vertex from $\text{Sep}_{R^2}(P, u)$. Thus, $\text{dist}_H(w, \text{Sep}_{R^2}(P, u)) \leq \alpha_{\text{dist}}(k)$. Combined with Condition 1, we derive that $\text{dist}_H(w, \tilde{A}_{R^2, P, u} \cup \{u'_P\}) \leq \alpha_{\text{dist}}(k) + \alpha_{\text{sep}}(k)$. The proof of Condition 4 is symmetric.

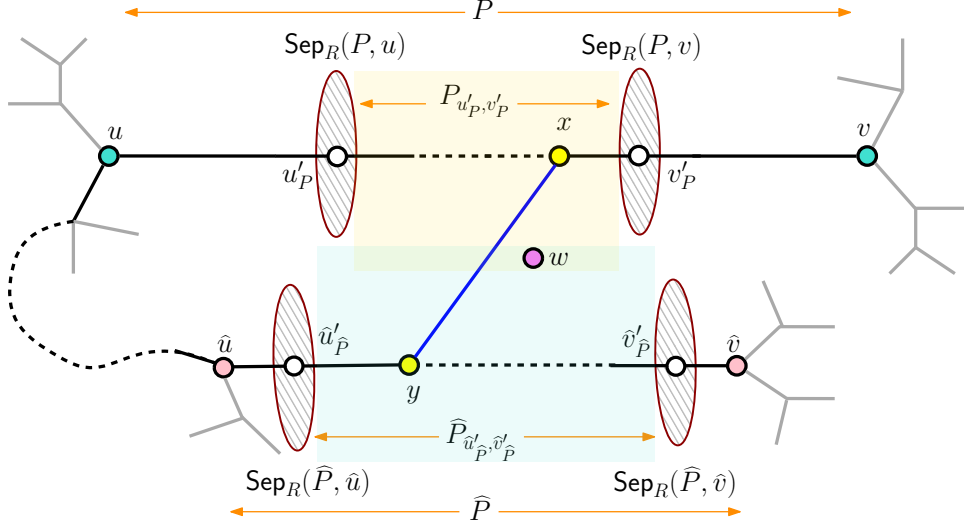


Figure 15: Illustration of Lemma 6.10.

The proof of Condition 5 is similar. Consider some vertex $w \in V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v}) \cup \text{Sep}_{R^2}(P, u) \cup \text{Sep}_{R^2}(P, v))$. As before, there exists a path Q in H with w as one endpoint and the other endpoint x in $V(R^2)$ such that the length of Q is at most $\alpha_{\text{dist}}(k)$. In case $x \in \tilde{A}_{R^2, P, u} \cup \{u'_P\}$, we are done. Otherwise, the path Q must traverse at least one vertex from $\text{Sep}_{R^2}(P, u) \cup \text{Sep}_{R^2}(P, v)$. Specifically, if $x \in V(C_{R^2, P, u})$, then it must traverse at least one vertex from $\text{Sep}_{R^2}(P, u)$, and otherwise $x \in V(C_{R^2, P, v})$ and it must traverse at least one vertex from $\text{Sep}_{R^2}(P, v)$. Combined with Conditions 1 and 2, we derive that $\text{dist}_H(w, V(R^2) \setminus (\tilde{A}_{R^2, P, u} \cup \tilde{A}_{R^2, P, v})) \leq \alpha_{\text{dist}}(k) + \alpha_{\text{sep}}(k)$. \square

An immediate corollary of Lemma 6.9 concerns the connectivity of the “middle region” as follows. (This corollary can also be easily proved directly.)

Corollary 6.2. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 be a Steiner tree that has no detour, P be a long maximal degree-2 path of R^2 , and u and v be its endpoints. Then, $H[V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v}))]$ is a connected graph.*

Proof. By Lemma 6.9 and the definition of $\text{Sep}_{R^2}(P, u)$ and $\text{Sep}_{R^2}(P, v)$, for every vertex in $V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v}))$, the graph H has a path from that vertex to some vertex in $\text{Sep}_{R^2}(P, u) \cup \text{Sep}_{R^2}(P, v)$ that lies entirely in $H[V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v}))]$. Thus, the corollary follows from Observation 6.6. \square

Next, we utilize Lemma 6.9 and Corollary 6.2 to argue that the “middle regions” of different long maximal degree-2 paths of R^2 are distinct. Recall that we wish to reroute a given solution to be a solution that “spirals” only a few times around the Steiner tree. This lemma allows us to independently reroute the solution in each of these “middle regions”. In fact, we prove the following stronger statement concerning these regions. The idea behind the proof of this lemma is that if it were false, then R^2 admits a detour, which is contradiction.

Lemma 6.10. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 be a Steiner tree that has no detour. Additionally, let P and \hat{P} be two distinct long maximal degree-2 paths of R^2 . Let u and v be the endpoints of P , and \hat{u} and \hat{v} be the endpoints of \hat{P} . Then, one of the two following conditions holds:*

- $V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v})) \subseteq V(C_{R^2, \hat{P}, \hat{u}})$.
- $V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v})) \subseteq V(C_{R^2, \hat{P}, \hat{v}})$.

Proof. We first prove that the intersection of $V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v}))$ with $V(H) \setminus (V(C_{R^2, \hat{P}, \hat{u}}) \cup V(C_{R^2, \hat{P}, \hat{v}}))$ is empty. To this end, suppose by way of contradiction that there exists a vertex w in this intersection. By Lemma 6.9, the inclusion of w in both $V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v}))$ and $V(H) \setminus (V(C_{R^2, \hat{P}, \hat{u}}) \cup V(C_{R^2, \hat{P}, \hat{v}}))$ implies that the two following inequalities are satisfied:

- $\text{dist}_H(w, V(R^2) \setminus (\tilde{A}_{R^2, P, u} \cup \tilde{A}_{R^2, P, v})) \leq \alpha_{\text{dist}}(k) + \alpha_{\text{sep}}(k)$.
- $\text{dist}_H(w, V(R^2) \setminus (\tilde{A}_{R^2, \hat{P}, \hat{u}} \cup \tilde{A}_{R^2, \hat{P}, \hat{v}})) \leq \alpha_{\text{dist}}(k) + \alpha_{\text{sep}}(k)$.

From this, we derive the following inequality:

$$\text{dist}_H(V(R^2) \setminus (\tilde{A}_{R^2, P, u} \cup \tilde{A}_{R^2, P, v}), V(R^2) \setminus (\tilde{A}_{R^2, \hat{P}, \hat{u}} \cup \tilde{A}_{R^2, \hat{P}, \hat{v}})) \leq 2(\alpha_{\text{dist}}(k) + \alpha_{\text{sep}}(k)).$$

In particular, this means that there exist vertices $x \in V(P_{u'_P, v'_P})$ and $y \in V(\hat{P}_{\hat{u}'_{\hat{P}}, \hat{v}'_{\hat{P}}})$ and a path Q in H between them whose length is at most $2(\alpha_{\text{dist}}(k) + \alpha_{\text{sep}}(k))$. Note that the unique path in R^2 between x and y traverses exactly one vertex in $\{u, v\}$. Suppose w.l.o.g. that this vertex is u . Then, consider the walk W (which might be a path) obtained by traversing P from v to x and then traversing Q from x to y . Now, notice that

$$\begin{aligned} |E(P)| - |E(W)| &\geq |E(P'_{u, u'_P})| - 2(\alpha_{\text{dist}}(k) + \alpha_{\text{sep}}(k)) \\ &\geq \alpha_{\text{pat}}(k)/2 - 2(\alpha_{\text{dist}}(k) + \alpha_{\text{sep}}(k)) \\ &= 50 \cdot 2^{ck} - 2(4 \cdot 2^{ck} + \frac{7}{2} \cdot 2^{ck} + 2) > 0. \end{aligned}$$

Here, the inequality $|E(P'_{u, u'_P})| \geq \alpha_{\text{pat}}(k)/2$ followed from Lemma 6.5. As $|E(P)| - |E(W)| > 0$, we have that u, v and any subpath of the walk W between u and v witness that R^2 has a detour (see Fig. 15). This is a contradiction, and hence we conclude that the intersection of $V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v}))$ with $V(H) \setminus (V(C_{R^2, \hat{P}, \hat{u}}) \cup V(C_{R^2, \hat{P}, \hat{v}}))$ is empty.

Having proved that the intersection is empty, we know that

$$V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v})) \subseteq V(C_{R^2, \hat{P}, \hat{u}}) \cup V(C_{R^2, \hat{P}, \hat{v}}).$$

Thus, it remains to show that $V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v}))$ cannot contain vertices from both $V(C_{R^2, \hat{P}, \hat{u}})$ and $V(C_{R^2, \hat{P}, \hat{v}})$. Since $H[V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v}))]$ is a connected graph (by Corollary 6.2), if $V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v}))$ contains vertices from both $V(C_{R^2, \hat{P}, \hat{u}})$ and $V(C_{R^2, \hat{P}, \hat{v}})$, then it must also contain at least one vertex from $\text{Sep}_{R^2}(\hat{P}, \hat{u}) \cup \text{Sep}_{R^2}(\hat{P}, \hat{v}) \subseteq V(H) \setminus (V(C_{R^2, \hat{P}, \hat{u}}) \cup V(C_{R^2, \hat{P}, \hat{v}}))$, which we have already shown to be impossible. Thus, the proof is complete. \square

We are now ready to prove that R^3 is a Steiner tree.

Lemma 6.11. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 be a Steiner tree that has no detour, and R^3 be the subgraph constructed from R^2 in Step IV. Then, R^3 is a Steiner tree with the following properties.*

- R^3 has the same set of vertices of degree at least 3 as R^2 .
- Every short maximal degree-2 path P of R^2 is a short maximal degree-2 path of R^3 .
- For every long maximal degree-2 path P of R^2 with endpoints u and v , the paths P_{u, u'_P} and P_{v, v'_P} are subpaths of the maximal degree-2 path of R^3 with endpoints u and v .

Proof. To prove that R^3 is a Steiner tree, we only need to show that R^3 is acyclic. Indeed, the construction of R^3 immediately implies that it is connected and has the same set of degree-1 vertices as R^2 , which together with an assertion that R^3 is acyclic, will imply that it is a Steiner tree. The other properties in the lemma are immediate consequences of the construction of R^3 .

By its construction, to show that R^3 is acyclic, it suffices to prove two conditions:

- For every long maximal degree-2 path P of R^2 with endpoints u and v , it holds that $V(P_{u'_P, v'_P}^*) \cap (V(R^2) \setminus V(P_{u'_P, v'_P})) = \emptyset$.
- For every two distinct long maximal degree-2 paths P and \widehat{P} of R^2 , it holds that $V(P_{u'_P, v'_P}^*) \cap V(\widehat{P}_{\widehat{u}', \widehat{v}'}) = \emptyset$, where u and v are the endpoints of P , \widehat{u} and \widehat{v} are the endpoints of \widehat{P} , $u' = u'_P, v' = v'_P, \widehat{u}' = \widehat{u}'_{\widehat{P}}$ and $\widehat{v}' = \widehat{v}'_{\widehat{P}}$.

The first condition follows directly from the fact that $P_{u'_P, v'_P}^*$ is a path in $H - (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v}))$ while $V(R^2) \setminus V(P_{u'_P, v'_P}) \subseteq V(C_{R^2, P, u}) \cup V(C_{R^2, P, v})$.

For the second condition, note that Lemma 6.10 implies that $V(H) \setminus (V(C_{R^2, P, u}) \cup V(C_{R^2, P, v})) \subseteq V(C_{R^2, \widehat{P}, \widehat{u}}) \cup V(C_{R^2, \widehat{P}, \widehat{v}})$. Thus, we have that $V(P_{u'_P, v'_P}^*) \subseteq V(C_{R^2, \widehat{P}, \widehat{u}}) \cup V(C_{R^2, \widehat{P}, \widehat{v}})$. However, $V(\widehat{P}_{\widehat{u}', \widehat{v}'}) \cap (V(C_{R^2, \widehat{P}, \widehat{u}}) \cup V(C_{R^2, \widehat{P}, \widehat{v}})) = \emptyset$, and hence $V(P_{u'_P, v'_P}^*) \cap V(\widehat{P}_{\widehat{u}', \widehat{v}'}) = \emptyset$. \square

We remark that R^3 might have detours. (These detours are restricted to $P_{u'_P, v'_P}$ for some long path $P = \text{path}_{R^3}(u, v)$ in R^3 .) However, what is important for us is that we can still use the same small separators as before. To this end, we first define the appropriate notations, in particular since later we will like to address objects corresponding to R^3 directly (without referring to R^2). The validity of these notations follows from Lemma 6.11. Recall that u'_P and P_{u, u'_P} refer to the vertex and path in Lemma 6.5.

Definition 6.5 (Translating Notations of R^2 to R^3). *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^2 and R^3 be the Steiner trees constructed in Steps II and IV. For any long maximal degree-2 path \widehat{P} of R^3 and for each endpoint u of \widehat{P} :*

- Define \widehat{P}_{R^2} as the unique (long maximal degree-2) path in R^2 with the same endpoints as \widehat{P} .
- Let $P = \widehat{P}_{R^2}$. Then, denote $u'_P = u'_P, \widehat{P}_{u, u'_P} = P_{u, u'_P}$ and $\text{Sep}_{R^3}(\widehat{P}, u) = \text{Sep}_{R^2}(P, u)$.
- Define $A_{R^3, \widehat{P}, u}^*$ as the union of $V(\widehat{P}_{u, u'_P})$ and the vertex set of the connected component of $R^3 - (V(\widehat{P}_{u, u'_P}) \setminus \{u\})$ containing u .
- Define $B_{R^3, \widehat{P}, u}^* = V(R^3) \setminus A_{R^3, \widehat{P}, u}^*$.

In the context of Definition 6.5, note that by Lemma 6.5, $u' \in \text{Sep}_{R^3}(\widehat{P}, u) \cap V(\widehat{P})$ where $u' = u'_P, \widehat{P}_{u, u'}$ is the subpath of \widehat{P} between u and u' , $\widehat{P}_{u, u'}$ has no internal vertex from $\text{Sep}_{R^3}(\widehat{P}, u) \cup \text{Sep}_{R^3}(\widehat{P}, v)$, and $\alpha_{\text{pat}}(k)/2 \leq |V(\widehat{P}_{u, u'})| \leq \alpha_{\text{pat}}(k)$. Additionally, note that $A_{R^3, \widehat{P}, u}^*$ might not be equal to $A_{R^2, P, u}$ where $P = \widehat{P}_{R^2}$. When R^3 is clear from context, we omit it from the subscripts.

Now, let us argue why, in a sense, we can still use the same small separators as before. Recall that a backbone Steiner tree is a Steiner tree constructed in Step IV.

Lemma 6.12. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R^3 be a backbone Steiner tree. Additionally, let \widehat{P} be a long maximal degree-2 path of R^3 , and u be an endpoint of \widehat{P} . Then, $\text{Sep}(\widehat{P}, u)$ separates $A_{\widehat{P}, u}^*$ and $B_{\widehat{P}, u}^*$ in H .*

Proof. Denote $P = \widehat{P}_{R^2}$ where R^2 is the Steiner tree computed in Step II to construct R^3 . Then, $\text{Sep}(\widehat{P}, u)$ separates $V(C_{R^2, P, v})$ and $V(C_{R^2, P, u})$ in H . Thus, to prove that $\text{Sep}(\widehat{P}, u)$ separates $A_{\widehat{P}, u}^*$ and $B_{\widehat{P}, u}^*$ in H , it suffices to show that $A_{\widehat{P}, u}^* \subseteq V(C_{R^2, P, u})$ and $A_{\widehat{P}, v}^* \subseteq V(C_{R^2, P, v})$ (because $B_{\widehat{P}, u}^* \setminus A_{\widehat{P}, v}^* \subseteq V(P_{u'_P, v'_P}^*)$ and $P_{u'_P, v'_P}^* \cap V(C_{R^2, P, u}) = \emptyset$ by the construction of $P_{u'_P, v'_P}^*$). We only prove that $A_{\widehat{P}, u}^* \subseteq V(C_{R^2, P, u})$. The proof of the other containment is symmetric.

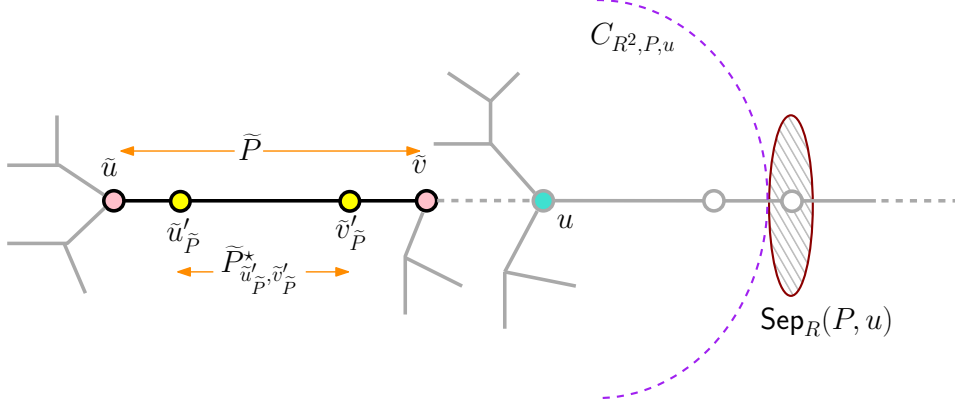


Figure 16: Illustration of Lemma 6.12.

Clearly, $A_{R^2, P, u} \cap A_{\tilde{P}, v}^* \subseteq V(C_{R^2, P, u})$. Thus, due to Lemma 6.11, to show that $A_{\tilde{P}, u}^* \subseteq V(C_{R^2, P, u})$, it suffices to show the following claim: For every long maximal degree-2 path \tilde{P} of R^2 whose vertex set is contained in $V(C_{R^2, P, u})$, it holds that the vertex set of $\tilde{P}^*_{\tilde{u}'_{\tilde{P}}, \tilde{v}'_{\tilde{P}}}$ (computed by Lemma 6.7) is contained in $V(C_{R^2, P, u})$ as well, where \tilde{u} and \tilde{v} are the endpoints of \tilde{P} . We refer the reader to Fig. 16 for an illustration of this statement. For the purpose of proving it, consider some long maximal degree-2 path \tilde{P} of R^2 whose vertex set is contained in $V(C_{R^2, P, u})$.

By Lemma 6.10, we know that either $V(H) \setminus (V(C_{R^2, \tilde{P}, \tilde{u}}) \cup V(C_{R^2, \tilde{P}, \tilde{v}})) \subseteq V(C_{R^2, P, u})$ or $V(H) \setminus (V(C_{R^2, \tilde{P}, \tilde{u}}) \cup V(C_{R^2, \tilde{P}, \tilde{v}})) \subseteq V(C_{R^2, P, v})$. Moreover, by the definition of $\tilde{P}^*_{\tilde{u}'_{\tilde{P}}, \tilde{v}'_{\tilde{P}}}$, its vertex set is contained in $V(H) \setminus (V(C_{R^2, \tilde{P}, \tilde{u}}) \cup V(C_{R^2, \tilde{P}, \tilde{v}}))$. Thus, to conclude the proof, it remains to rule out the possibility that $V(H) \setminus (V(C_{R^2, \tilde{P}, \tilde{u}}) \cup V(C_{R^2, \tilde{P}, \tilde{v}})) \subseteq V(C_{R^2, P, v})$. For this purpose, recall that we chose \tilde{P} such that $V(\tilde{P}) \subseteq V(C_{R^2, P, u})$, and that $V(\tilde{P}) \cap V(C_{R^2, P, u}) \neq \emptyset$. Because $V(C_{R^2, P, u}) \cap V(C_{R^2, P, v}) = \emptyset$, we derive that the containment $V(H) \setminus (V(C_{R^2, \tilde{P}, \tilde{u}}) \cup V(C_{R^2, \tilde{P}, \tilde{v}})) \subseteq V(C_{R^2, P, v})$ is indeed impossible. \square

6.6 Enumerating Parallel Edges with Respect to R^3

Recall that H is enriched with $4n + 1$ parallel copies of each edge of the (standard) radial completion of G . While the copies did not play a role in the construction of R , they will be important in how we relate a solution of the given instance of Planar Disjoint Paths to a weak linkage in H . We remind that for a pair of adjacent vertices $u, v \in V(H)$, we denoted the $4n + 1$ parallel copies of edges between them by $e_{-2n}, e_{-2n+1}, \dots, e_{-1}, e_0, e_1, e_2, \dots, e_{2n}$ where $e = \{u, v\}$, such that when the edges incident to u (or v) are enumerated in cyclic order, the occurrences of e_i and e_{i+1} are consecutive (that is, e_i appears immediately before e_{i+1} or vice versa) for every $i \in \{-2n, -2n+1, \dots, 2n-1\}$, and e_{-2n} and e_{2n} are the outermost copies of e . We say that such an embedding is *valid*. Now, we further refine the embedding of H_G (so that it remains valid)—notice that for each edge, there are two possible ways to order its copies in the embedding so that it satisfies the condition above. Here, we will specify for the edges of R a particular choice among the two to embed their copies. Towards this, for a vertex $v \in V(R)$, let $\widehat{E}_R(v) = \{e \in E_H(v) \mid e \text{ is parallel to an edge } e' \in E(R)\}$. (The set $E_H(v)$ contains all edges in $E(H)$ incident to v .)

For the definition of the desired embedding, we remind that any tree can be properly colored in two colors (that is, every vertex is assigned a color different than the colors of its neighbors), and that in such a coloring, for every vertex, all the neighbors of the vertex get the same color. We let $\text{color} : V(R) \rightarrow \{\text{red}, \text{green}\}$ be some such coloring of R . Then, we embed parallel copies such that for every $v \in V(R)$, the following conditions hold (see Fig. 17).

- If $\text{color}(v) = \text{red}$, then when we enumerate $\widehat{E}_R(v)$ in *clockwise* order, for every $e \in E_R(v)$,




Figure 17: The clockwise / anti-clockwise enumeration of parallel edges with respect to R .

the $4n + 1$ copies of e are enumerated in this order: $e_{-2n}, e_{-2n+1}, \dots, e_0, \dots, e_{2n}$. We let order_v denote such an enumeration starting with an edge indexed $-2n$.

- If $\text{color}(v) = \text{green}$, then when we enumerate $\widehat{E}_R(v)$ in *counter-clockwise* order, for every $e \in E_R(v)$, the $4n+1$ copies of e are enumerated in this order: $e_{-2n}, e_{-2n+1}, \dots, e_0, \dots, e_{2n}$. We let order_v denote such an enumeration starting with an edge indexed $-2n$.

Let us observe that the above scheme is well defined.

Observation 6.9. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths with a backbone Steiner tree R . Then, there is a valid embedding of H such that, for every $v \in V(R)$, the enumeration order_v is well defined with respect to some proper coloring $\text{color} : V(R) \rightarrow \{\text{red}, \text{green}\}$. Furthermore, such an embedding can be computed in time $\mathcal{O}(n^2)$.*

Proof. Since $e = \{u, v\} \in E(R)$, u and v get different colors under color . Let us assume that $\text{color}(u) = \text{red}$ and $\text{color}(v) = \text{green}$. Then the parallel copies of e are enumerated in clockwise order in order_v and anti-clockwise order in order_u . Hence, they agree and by construction, it is a good enumeration. Finally, to bound the time required to obtain such an embedding, observe that it can be obtained by starting with any arbitrary embedding of H and then renaming the edges. Since the total number of edges in $E(H)$ (including parallel copies) is at most $\mathcal{O}(n^2)$, this can be done in $\mathcal{O}(n^2)$ time. \square

From now on, we assume that H is embedded in a way so that the enumerations order_v are well defined. We also remind that R only contains the 0-th copies of edges in H . Finally, we have the following observation.

Observation 6.10. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths with a backbone Steiner tree R . For every $v \in V(H)$, order_v is an enumeration of $\widehat{E}_R(v)$ in either clockwise or counter-clockwise order around v (with a fixed start). Further, for any pair $e, e' \in E_R(v)$ such that e occurs before e' in order_v , the edges e_0, e_1, \dots, e_{2n} occur before $e'_0, e'_1, \dots, e'_{2n}$.*

7 Existence of a Solution with Small Winding Number

In this section we show that if the given instance admits a solution, then it admits a “nice solution”. The precise definition of nice will be in terms of “winding number” of the solution, which counts the number of times the solution spirals around the backbone steiner tree. Our goal is to show that there is a solution of small winding number.

7.1 Rings and Winding Numbers

Towards the definition of a ring, let us remind that H is the triangulated plane multigraph obtained by introducing $4n + 1$ parallel copies of each edge to the a radial completion of the input graph G . Hence, each face of H is either a triangle or a 2-cycle.

Definition 7.1 (Ring). Let $I_{\text{in}}, I_{\text{out}}$ be two disjoint cycles in H such that the cycle I_{in} is drawn in the strict interior of the cycle I_{out} . Then, $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ is the plane subgraph of H induced by the set of vertices that are either in $V(I_{\text{in}}) \cup V(I_{\text{out}})$ or drawn between I_{in} and I_{out} (i.e. belong to the exterior of I_{in} and the interior of I_{out}).

We call I_{in} and I_{out} are the *inner* and *outer interfaces* of $\text{Ring}(I_{\text{in}}, I_{\text{out}})$. We also say that this ring is induced by I_{in} and I_{out} . Recall the notion of self-crossing walks defined in Section 5. Unless stated otherwise, all walks considered here are *not self-crossing*. A walk α in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ is *traversing* the ring if one of its endpoints lies in I_{in} and the other lies in I_{out} . A walk α is *visiting* the ring if both its endpoints together lie in either I_{in} or in I_{out} ; moreover α is an *inner visitor* if both its endpoints lie in I_{in} , and otherwise it is an *outer visitor*.

Definition 7.2 (Orienting Walks). Fix an arbitrary ordering of all vertices in I_{in} and another one for all vertices in I_{out} . Then for a walk α in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ with endpoints in $V(I_{\text{in}}) \cup V(I_{\text{out}})$, orient α from one endpoint to another as follows. If α is a traversing walk, then orient it from its endpoint in I_{in} to its endpoint in I_{out} . If α is a visiting walk, then both its endpoints lie either in I_{in} or in I_{out} ; then, orient α from its smaller endpoint to its greater endpoint.

Observe that if α is a traversing path in the ring, then the orientation of α also defines its *left-side* and *right-side*. These are required for the following definition.

Definition 7.3 (Winding Number of a Walk w.r.t. a Traversing Path). Let α be an a walk in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ with endpoints in $V(I_{\text{in}}) \cup V(I_{\text{out}})$, and let β be a traversing path in this ring, such that α and β are edge disjoint. The winding number, $\text{WindNum}(\alpha, \beta)$, of α with respect to β is the signed number of crossings of α with respect to β . That is, while walking along α (according to the orientation in Definition 7.2, for each intersection of α and β record $+1$ if α crosses β from left to right, -1 if α crosses β from right to left, and 0 if it α does not cross β . Then, the winding number $\overline{\text{WindNum}}(\alpha, \beta)$ is the sum of the recorded numbers.

Observe that if α and β are edge-disjoint traversing paths, then both $\overline{\text{WindNum}}(\alpha, \beta)$ and $\overline{\text{WindNum}}(\beta, \alpha)$ are well defined. We now state some well-known properties of the winding number. We sketch a proof of these properties in Appendix A, using homotopy.

Proposition 7.1. Let α, β and γ be three edge-disjoint paths traversing $\text{Ring}(I_{\text{in}}, I_{\text{out}})$. Then,

$$(i) \quad \overline{\text{WindNum}}(\beta, \gamma) = -\overline{\text{WindNum}}(\gamma, \beta).$$

$$(ii) \quad \left| \overline{\text{WindNum}}(\alpha, \beta) - \overline{\text{WindNum}}(\alpha, \gamma) \right| - \left| \overline{\text{WindNum}}(\beta, \gamma) \right| \leq 1.$$

We say that $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ is *rooted* if it is equipped with some fixed path η that is traversing it, called the *reference path* of this ring. In a rooted ring $\text{Ring}(I_{\text{in}}, I_{\text{out}})$, we measure all winding numbers with respect to η , hence we shall use the shorthand $\overline{\text{WindNum}}(\alpha) = \overline{\text{WindNum}}(\alpha, \eta)$ when η is implicit or clear from context. Here, we implicitly assume that the walk α is edge disjoint from η . This requirement will always be met by the following assumptions: (i) H is a plane multigraph where we have $4n + 1$ parallel copies of every edge, and we assume that the reference path η consists of only the 0-th copy e_0 ; and (ii) whenever we consider the winding number of a walk α , it will edge-disjoint from the reference curve η as it will not contain the 0-th copy of any edge. (In particular, the walks of the (weak) linkages that we consider will always satisfy this property.)

Note that any visitor walk in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ with both endpoints in I_{in} is discretely homotopic to a segment of I_{in} , and similarly for I_{out} . Thus, we derive the following observation.

Observation 7.1. Let α be a visitor in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$. Then, $|\overline{\text{WindNum}}(\alpha)| \leq 1$.

Recall the notion of a weak linkage defined in Section 5, which is a collection of edge-disjoint non-crossing walks. When we use the term *weak linkage of order k in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$* , we refer to a weak linkage such that each walk has both endpoints in $V(I_{\text{in}}) \cup V(I_{\text{out}})$. For brevity, we abuse the term ‘weak linkage’ to mean a weak linkage in a ring when it is clear from context. Note that every walk in a weak linkage \mathcal{P} is an inner visitor, or an outer visitor, or a traversing walk. This partitions \mathcal{P} into $\mathcal{P}_{\text{in}}, \mathcal{P}_{\text{out}}, \mathcal{P}_{\text{traverse}}$. A weak linkage is *traversing* if it consists only of traversing walks. Assuming that $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ is rooted, we define the *winding number* of a traversing weak linkage \mathcal{P} as $\overline{\text{WindNum}}(\mathcal{P}) = \overline{\text{WindNum}}(P_1)$. Recall that any two walks in a weak linkage are non-crossing. Then as observed in [14, Observation 4.4],¹¹

$$|\overline{\text{WindNum}}(P_i) - \overline{\text{WindNum}}(\mathcal{P})| \leq 1 \quad \text{for all } i = 1, \dots, k.$$

The above definition is extended to any weak linkage \mathcal{P} in the ring as follows: if there is no walk in \mathcal{P} that traverses the ring, then $\overline{\text{WindNum}}(\mathcal{P}) = 0$, otherwise $\overline{\text{WindNum}}(\mathcal{P}) = \overline{\text{WindNum}}(\mathcal{P}_{\text{traverse}})$. Note that, two aligned weak linkages \mathcal{P} and \mathcal{Q} in the ring may have different winding numbers (with respect to any reference path). Replacing a linkage \mathcal{P} with an aligned linkage \mathcal{Q} having a “small” winding number will be the main focus of this section.

Lastly, we define a labeling of the edges based on the winding number of a walk (this relation is made explicit in the observation that follows).

Definition 7.4. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths, and H be the radial completion of G . Let α be a (not self-crossing) walk in H , and let β be a path in H such that α and β are edge disjoint. Let us fix (arbitrary) orientations of α and β , and define the left and right side of the path β with respect to its orientation. The labeling $\text{label}_\beta^\alpha$ of each ordered pair of consecutive edges, $(e, e') \in E_H(\alpha) \times E_H(\alpha)$ by $\{-1, 0, +1\}$ with respect to β , where e occurs before e' when traversing α according to its orientation is defined as follows.*

- The pair (e, e') is labeled $+1$ if e is on the left of β while e' is on the right of β .
- else, (e, e') is labeled -1 if e is on the right while e' is on the left of β ;
- otherwise e and e' are on the same side of β and (e, e') is labeled 0 .

Note that in the above labeling only pairs of consecutive edges may get a non-zero label, depending on how they cross the reference path. For the ease of notation, we extend the above labeling function to all ordered pairs of edges in α (including pairs of non-consecutive edges), by labeling them 0 . Then we have the following observation, when we restrict α to $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ and set β to be the reference path of this ring.

Observation 7.2. *Let α be a (not self-crossing) walk in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ with reference path η . Then $|\overline{\text{WindNum}}(\alpha, \eta)| = |\sum_{(e, e') \in E(\alpha) \times E(\alpha)} \text{label}_\eta^\alpha(e, e')|$.*

7.2 Rerouting in a Ring

We now address the question of rerouting a solution to reduce its winding number with respect to the backbone Steiner tree. As a solution is linkage in the graph G (i.e. a collection of vertex disjoint paths), we first show how to reroute linkages within a ring. In the later subsections, we will apply this to reroute a solution in the entire plane graph. We remark that from now onwards, our results are stated and proved only for linkages (rather than weak linkages). Further, define a *linkage of order k in a $\text{Ring}(I_{\text{in}}, I_{\text{out}})$* as a collection of k vertex-disjoint paths in G such that each of these paths belongs to $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ and its endpoints belong to $V(I_{\text{in}}) \cup V(I_{\text{out}})$. As before, we simply use the term ‘linkage’ when the ring is clear from context. We will use the

¹¹This inequality also follows from the second property of Proposition 7.1 by setting α to be the reference path, $\beta = P_1$ and $\gamma = P_i$ and noting that $\overline{\text{WindNum}}(\beta, \gamma) = 0$.

following proposition proved by Cygan et al. [14] using earlier results of Ding et al. [17]. Its statement has been rephrased to be compatible with our notation.

Proposition 7.2 (Lemma 4.8 in [14]). *Let $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ be a rooted ring in H and let \mathcal{P} and \mathcal{Q} be two traversing linkages of the same order in this ring. Then, there exists a traversing linkage \mathcal{P}' in this ring that is aligned with \mathcal{P} and such that $|\overline{\text{WindNum}}(\mathcal{P}') - \overline{\text{WindNum}}(\mathcal{Q})| \leq 6$.*

The formulation of [14] concerns directed paths in directed graphs and assumes a fixed pattern of in/out orientations of paths that is shared by the linkages \mathcal{P} , \mathcal{Q} and \mathcal{P}' . The undirected case (as expressed above) can be reduced to the directed one by replacing every undirected edge in the graph by two oppositely-oriented arcs with same endpoints, and asking for any orientation pattern (say, all paths should go from I_{in} to I_{out}). Moreover, the setting itself is somewhat more general, where rings and reference paths are defined by curves and (general) homotopy.

Rings with Concentric Cycles. Let $\mathcal{C} = (C_1, C_2, \dots, C_p)$ concentric sequence of cycles in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ (then, C_i is in the strict interior of C_{i+1} for $i \in \{1, 2, \dots, p-1\}$). If I_{in} is in the strict interior of C_1 and C_p is in the strict interior of I_{out} , then we say that \mathcal{C} is *encircling*. An encircling concentric sequence \mathcal{C} in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ is *tight* if every $C \in \mathcal{C}$ is a cycle in G , and there exists a path η in H traversing $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ such that the set of internal vertices of η contain exactly $|\mathcal{C}|$ vertices of $V(G)$, one on each cycle in \mathcal{C} . Let us fix one such encircling tight sequence in the ring $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ along with the path η witnessing the tightness. Then, we set the path η as the reference path of the ring. Here, we assume w.l.o.g. that η contains only the 0-th copy of each of the edges comprising it. Any paths or linkages that we subsequently consider will not use the 0-th copy of any edge, and hence their winding numbers (with respect to η) will be well-defined. This is because that they arise from G , and when we consider them in H , we choose a ‘non-0-th’ copy out of the $4n + 1$ copies of any (required) edge.

A linkage \mathcal{P} in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ is *minimal* with respect to \mathcal{C} if among the linkages aligned with \mathcal{P} , it minimizes the total number of edges traversed that do not lie on the cycles of \mathcal{C} . The following proposition is essentially Lemma 3.7 of [14].

Proposition 7.3. *Let G be a plane graph, and with radial completion H . Let $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ be a rooted ring in H . Suppose $|I_{\text{in}}|, |I_{\text{out}}| \leq \ell$, for some integer ℓ . Further, let $\mathcal{C} = (C_1, \dots, C_p)$ be an encircling tight concentric sequence of cycles in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$. Finally, let \mathcal{P} be a linkage in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ that is minimal with respect to \mathcal{C} . Then, every inner visitor of \mathcal{P} intersects less than 10ℓ of the first cycles in the sequence (C_1, \dots, C_p) , while every outer visitor of \mathcal{P} intersects less than 10ℓ of the last cycles in this sequence.*

A proof of this proposition can be obtained by first ordering the collection of inner and outer visitors by their ‘distance’ from the inner and outer interfaces, respectively, and the ‘containment’ relation between the cycles formed by them with the interfaces. This gives a partial order on the set of inner visitors and the set of outer visitors. Then if the proposition does not hold for \mathcal{P} , then the above ordering and containment relation can be used to reroute these paths along a suitable cycle. This will contradict the minimality of \mathcal{P} , since the rerouted linkage is aligned with it but uses strictly fewer edges outside of \mathcal{C} . The main result of this section can be now formulated as follows. (Its formulation and proof idea are based on Lemma 8.31 and Theorem 6.45 of [14].)

Lemma 7.1. *Let G be a plane graph with radial completion H . Let $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ be a ring in H . Suppose that $|I_{\text{in}}|, |I_{\text{out}}| \leq \ell$ for some integer ℓ , and that in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ there is an encircling tight concentric sequence of cycles \mathcal{C} of size larger than 40ℓ . Let η be a traversing path in the ring witnessing the tightness of \mathcal{C} , and fix η as the reference path. Finally, let $\mathcal{P} = \mathcal{P}_{\text{traverse}} \uplus \mathcal{P}_{\text{visitor}}$ be a linkage in G , where $\mathcal{P}_{\text{traverse}}$ is a traversing linkage comprising the paths of \mathcal{P} traversing $\text{Ring}(I_{\text{in}}, I_{\text{out}})$, while $\mathcal{P}_{\text{visitor}} = \mathcal{P} \setminus \mathcal{P}_{\text{traverse}}$ consists of the paths whose both*

endpoints lie in either $V(I_{\text{in}})$ or $V(I_{\text{out}})$. Further, suppose that \mathcal{P} is minimal with respect to \mathcal{C} . Then, for every traversing linkage \mathcal{Q} in G that is minimal with respect to \mathcal{C} such that every path in \mathcal{Q} is disjoint from η and $|\mathcal{Q}| \geq |\mathcal{P}_{\text{traverse}}|$, there is a traversing linkage $\mathcal{P}'_{\text{traverse}}$ in G such that

- (a) $\mathcal{P}'_{\text{traverse}}$ is aligned with $\mathcal{P}_{\text{traverse}}$,
- (b) the paths of $\mathcal{P}'_{\text{traverse}}$ are disjoint from the paths of $\mathcal{P}_{\text{visitor}}$, and
- (c) $|\overline{\text{WindNum}(\mathcal{P}'_{\text{traverse}})} - \overline{\text{WindNum}(\mathcal{Q})}| \leq 60\ell + 6$.

Proof. Let $\mathcal{C} = (C_1, \dots, C_p)$, where $p > 40\ell$. Recall that \mathcal{C} is a collection of cycles in G , and the path η that witnesses the tightness of \mathcal{C} contains $|\mathcal{C}|$ vertices of $V(G)$, one on each cycle of \mathcal{C} . Let v_i denote the vertex where η intersects the cycle $C_i \in \mathcal{C}$ for all $i \in \{1, 2, \dots, p\}$. Since \mathcal{P} is minimal with respect to \mathcal{C} , Proposition 7.3 implies that the paths in $\mathcal{P}_{\text{visitor}}$ do not intersect any of the cycles $C_{10\ell}, C_{10\ell+1}, \dots, C_{p-10\ell+1}$ (note that since $p > 40\ell$, this sequence of cycles is non-empty). Call a vertex $x \in V(\text{Ring}(I_{\text{in}}, I_{\text{out}}))$ in the ring *non-separated* if there exists a path from x to $C_{10\ell}$ whose set of internal vertices is disjoint from $\bigcup_{P \in \mathcal{P}_{\text{visitor}}} V(P)$. Otherwise, we say that the vertex x is *separated*. Observe that every path in $\mathcal{P}_{\text{traverse}}$ is disjoint from the paths in $\mathcal{P}_{\text{visitor}}$ and intersects $C_{10\ell}$, hence all vertices on the paths of $\mathcal{P}_{\text{traverse}}$ are non-separated. Let X denote the set of all non-separated vertices in the ring, and consider the graph $H[X]$. Observe that $H[X]$ is an induced subgraph of $\text{Ring}(I_{\text{in}}, I_{\text{out}})$, since it is obtained from $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ by deleting the separated vertices. Further, observe that $H[X]$ is a ring of H . Indeed, the inner interface of $H[X]$ is the cycle \widehat{I}_{in} obtained as follows: let \mathcal{P}_{in} be the set of inner visitors in \mathcal{P} ; then, \widehat{I}_{in} is the outer face of the plane graph $H[V(I_{\text{in}}) \cup \bigcup_{P \in \mathcal{P}_{\text{in}}} V(P)]$. It is easy to verify that all vertices on \widehat{I}_{in} are non-separated, and any vertex of $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ that lies in the strict interior of this cycle is separated. We then symmetrically obtain the outer interface \widehat{I}_{out} of $H[X]$ from the set \mathcal{P}_{out} of outer visitors of \mathcal{P} . Then, $H[X] = \text{Ring}(\widehat{I}_{\text{in}}, \widehat{I}_{\text{out}})$. Here, \widehat{I}_{in} is composed alternately of subpaths of I_{in} and inner visitors from $\mathcal{P}_{\text{visitor}}$, and symmetrically for \widehat{I}_{out} .

Note that the paths of $\mathcal{P}_{\text{traverse}}$ are completely contained in $\text{Ring}(\widehat{I}_{\text{in}}, \widehat{I}_{\text{out}})$ and they all traverse this ring. Thus, $\mathcal{P}_{\text{traverse}}$ can be regarded also as a traversing linkage in $\text{Ring}(\widehat{I}_{\text{in}}, \widehat{I}_{\text{out}})$. While $\mathcal{P}_{\text{traverse}}$ may have a different winding number in $\text{Ring}(\widehat{I}_{\text{in}}, \widehat{I}_{\text{out}})$ than in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$, the difference is “small” as we show below. (Note that the two winding numbers in the following claim are computed in two different rings.)

Claim 7.1. *Let P be a path in G that is disjoint from all paths in $\mathcal{P}_{\text{visitor}}$, such that P belongs to $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ and traverses it. Then, P also belongs to $\text{Ring}(\widehat{I}_{\text{in}}, \widehat{I}_{\text{out}})$, and $|\overline{\text{WindNum}(P, \eta)} - \overline{\text{WindNum}(P, \widehat{\eta})}| \leq 20\ell$.¹²*

Proof. Since P traverses $\text{Ring}(I_{\text{in}}, I_{\text{out}})$, it must intersect the cycle $C_{10\ell}$. Therefore, as P is disjoint from $\mathcal{P}_{\text{visitor}}$, all vertices in $V(P)$ are non-separated. Hence, P is present in $\text{Ring}(\widehat{I}_{\text{in}}, \widehat{I}_{\text{out}})$. Next, observe that there are at most 20ℓ vertices of G that are visited by η but not visited by $\widehat{\eta}$; indeed, these are vertices in the intersection of η with $I_{\text{in}}, I_{\text{out}}$ and the first and last $10\ell - 1$ cycles of \mathcal{C} . It follows that any path in G has at most 20ℓ more crossings with η than with $\widehat{\eta}$. Since each such crossing contributes $+1$ or -1 to the winding number of P with respect to η , the winding numbers of P with respect to η and $\widehat{\eta}$ differ by at most 20ℓ . \diamond

We now turn our attention to the linkage \mathcal{Q} . In essence, our goal is show that every path in \mathcal{Q} can be “trimmed” to a path traversing $\text{Ring}(\widehat{I}_{\text{in}}, \widehat{I}_{\text{out}})$ such that their winding numbers are not significantly different. First, however, we prove that the paths in \mathcal{Q} cannot “oscillate” too much in $\text{Ring}(I_{\text{in}}, I_{\text{out}})$, based on the supposition that \mathcal{Q} is minimal with respect to \mathcal{C} .

Claim 7.2. *Let $Q \in \mathcal{Q}$, and let $u \in V(Q)$ such that it also lies on an inner visitor from $\mathcal{P}_{\text{visitor}}$. Then, the prefix of Q between its endpoint on I_{in} and u does not intersect the cycle $C_{20\ell}$.*

¹²Here $\widehat{\eta}$ is the reference path of $\text{Ring}(\widehat{I}_{\text{in}}, \widehat{I}_{\text{out}})$, which is a subpath of η in this ring.




Figure 18: Illustration of Claim 7.2.

Proof. Suppose, for the sake of contradiction, that the considered prefix contains some vertex v that lies on $C_{20\ell}$. Since u lies on an inner visitor $P \in \mathcal{P}_{\text{visitor}}$ and Proposition 7.3 states that an inner visitor cannot intersect $C_{10\ell}$, we infer that on the infix of Q between v and u there exists a vertex that lies on the intersection of Q and $C_{10\ell}$. Let a be the first such vertex. Similarly, on the prefix of Q from its endpoint on I_{in} to v there exists a vertex that lies on the intersection of Q and $C_{10\ell}$. Let b be the last such vertex. Then the whole infix of Q between a and b does not intersect $C_{10\ell}$ internally (see Fig. 18), and hence, apart from endpoints, completely lies in the exterior of $C_{10\ell}$. Call this infix Q^* .

Now consider $\text{Ring}(C_{10\ell}, I_{\text{out}})$, the ring induced by I_{out} and $C_{10\ell}$. Moreover, consider the graph G' obtained from G by removing all vertices that are not in $\text{Ring}(C_{10\ell}, I_{\text{out}})$ and edges that are not in the strict interior of $\text{Ring}(C_{10\ell}, I_{\text{out}})$; in particular, the edges of $C_{10\ell}$ are removed, but the vertices are not. Note that G' is a subgraph of $\text{Ring}(C_{10\ell}, I_{\text{out}})$. Finally, let $\mathcal{C}' = \mathcal{C} \setminus \{C_1, C_2, \dots, C_{10\ell}\}$; then, \mathcal{C}' is an encircling tight sequence of concentric cycles in $\text{Ring}(C_{10\ell}, I_{\text{out}})$.

Let \mathcal{Q}' be the linkage in G obtained by restricting paths of \mathcal{Q} to G' . Here, a path in \mathcal{Q} may break into several paths in \mathcal{Q}' (that are its maximal subpaths contained in G'). Since \mathcal{Q} is minimal with respect to \mathcal{C} , it follows that \mathcal{Q}' is minimal with respect to \mathcal{C}' . Now, observe that Q^* belongs to \mathcal{Q}' , hence it is an inner visitor of $\text{Ring}(C_{10\ell}, I_{\text{out}})$. However, Q^* intersects the first $10\ell + 1$ concentric cycles $C_{10\ell+1}, \dots, C_{20\ell}$ in the family \mathcal{C}' , which contradicts Proposition 7.3. \diamond

Clearly, an analogous claim holds for outer visitors and the cycle $C_{p-20\ell+1}$. We now proceed to our main claim about the restriction of \mathcal{Q} to $\text{Ring}(\hat{I}_{\text{in}}, \hat{I}_{\text{out}})$.

Claim 7.3. *For every path $Q \in \mathcal{Q}$, there exists a subpath \hat{Q} of Q that traverses $\text{Ring}(\hat{I}_{\text{in}}, \hat{I}_{\text{out}})$ and such that $|\overline{\text{WindNum}}(\hat{Q}, \hat{\eta}) - \overline{\text{WindNum}}(Q, \eta)| \leq 40\ell$.*

Proof. We think of Q as oriented from its endpoint on I_{in} to its endpoint on I_{out} . Let a be the last vertex on Q that lies on \hat{I}_{in} and b be the first vertex on Q that lies on \hat{I}_{out} . Further, let Q_{in} be the prefix of Q from its start to a , and Q_{out} be the suffix of Q from b to its end. By Claim 7.2, Q_{in} is entirely contained in the ring induced by I_{in} and $C_{20\ell}$. By the claim analogous to Claim 7.2, Q_{out} is entirely contained in the ring induced by I_{out} and $C_{p-20\ell+1}$. Since $p > 40\ell$, it follows that Q_{in} and Q_{out} are disjoint, and in particular a appears before b on Q . Let \hat{Q} be the infix of Q between a and b . Then, \hat{Q} is a path in $\text{Ring}(\hat{I}_{\text{in}}, \hat{I}_{\text{out}})$ that traverses this ring, so it suffices to check that $|\overline{\text{WindNum}}(\hat{Q}, \hat{\eta}) - \overline{\text{WindNum}}(Q, \eta)| \leq 40\ell$.

Observe that every crossing of Q and η that is not a crossing of \hat{Q} and $\hat{\eta}$ has to occur on either Q_{in} or Q_{out} . However, Q_{in} and Q_{out} can have at most 40ℓ vertices in common with η , because these must be among the intersections of η with cycles $C_1, \dots, C_{20\ell}, C_{p-20\ell+1}, \dots, C_p$, of which there are 40ℓ . Each such crossing can contribute $+1$ or -1 to the difference between the winding numbers $\overline{\text{WindNum}}(\hat{Q}, \hat{\eta})$ and $\overline{\text{WindNum}}(Q, \eta)$, hence the difference between these winding numbers is at most 40ℓ . \diamond

For every path $Q \in \mathcal{Q}$, fix the path \widehat{Q} provided by Claim 7.3, and let $\widehat{\mathcal{Q}} \subseteq \{\widehat{Q} \mid Q \in \mathcal{Q}\}$ such that $|\widehat{\mathcal{Q}}| = |\mathcal{P}_{\text{traverse}}|$. Then, $\widehat{\mathcal{Q}}$ is a traversing linkage in $\text{Ring}(\widehat{I}_{\text{in}}, \widehat{I}_{\text{out}})$. Apply Proposition 7.2 to the linkages $\mathcal{P}_{\text{traverse}}$ and $\widehat{\mathcal{Q}}$ in $\text{Ring}(\widehat{I}_{\text{in}}, \widehat{I}_{\text{out}})$, yielding a linkage $\mathcal{P}'_{\text{traverse}}$ that is aligned from $\mathcal{P}_{\text{traverse}}$ and such that

$$|\overline{\text{WindNum}}(\mathcal{P}'_{\text{traverse}}, \widehat{\eta}) - \overline{\text{WindNum}}(\widehat{\mathcal{Q}}, \widehat{\eta})| \leq 6.$$

Clearly, by construction we have that the paths of $\mathcal{P}'_{\text{traverse}}$ are disjoint with the paths of $\mathcal{P}_{\text{visitor}}$. Furthermore, the paths in $\mathcal{P}'_{\text{traverse}}$ traverse $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ since they are aligned with $\mathcal{P}_{\text{traverse}}$ (i.e. they have the same endpoints). Finally, by Claim 7.3 we have

$$|\overline{\text{WindNum}}(\widehat{\mathcal{Q}}, \widehat{\eta}) - \overline{\text{WindNum}}(\mathcal{Q}, \eta)| \leq 40\ell,$$

and by Claim 7.1 (applied to paths in $\mathcal{P}'_{\text{traverse}}$) we have

$$|\overline{\text{WindNum}}(\mathcal{P}'_{\text{traverse}}, \eta) - \overline{\text{WindNum}}(\mathcal{P}'_{\text{traverse}}, \widehat{\eta})| \leq 20\ell.$$

By the above we conclude that

$$|\overline{\text{WindNum}}(\mathcal{P}'_{\text{traverse}}, \eta) - \overline{\text{WindNum}}(\mathcal{Q}, \eta)| \leq 60\ell + 6,$$

which completes the proof. \square

7.3 Rings of the Backbone Steiner Tree

Based on results in previous subsections, we proceed to show that if the given instance admits a solution, then it also admits a solution of small winding number. Recall the backbone Steiner tree R constructed in Section 6. Let $P = \text{path}_R(u, v)$ be a long maximal degree-2 path in R , where $u, v \in V_{=1}(R) \cup V_{\geq 3}(R)$, and assume without loss of generality (under the supposition that we are given a Yes instance) that the subtree of $R - V(P) - \{u, v\}$ containing v also contains the terminal $t^* \in T$ lying on the outer face of H . Recall the (minimal) separators $S_u = \text{Sep}_R(P, u)$ and $S_v = \text{Sep}_R(P, v)$ in H . Hence $H[S_u]$ and $H[S_v]$ form two cycles in H , and $H[S_u]$ is contained in the strict interior of $H[S_v]$. Further, recall that $|S_u|, |S_v| \leq \alpha_{\text{sep}}(k)$. Consider the ring induced by $H[S_u]$ and $H[S_v]$, i.e. $\text{Ring}(S_u, S_v) := \text{Ring}(H[S_u], H[S_v])$. Let $V(S_u, S_v)$ denote the set of all vertices (in $V(H)$) that lie in this ring, including those in S_u and S_v . Then, $\text{Ring}(S_u, S_v) = H[V(S_u, S_v)]$. Note that by definition it contains $\text{Sep}_R(P, u)$ and $\text{Sep}_R(P, v)$. Let $G_{u,v}$ denote the restriction of this graph to G , i.e. $G_{u,v} = G[V(G) \cap V(S_u, S_v)]$. Additionally, recall that there are two distinct vertices u' and v' in P that lie in S_u and S_v , respectively, such that $P = \text{path}_R(u, u') - \text{path}_R(u', v') - (v', v)$ (by Lemma 6.5 and Definition 6.5). Lastly, we remind that $A_{R,P,u}^*$ and $A_{R,P,v}^*$ are the two components of $R - V(\text{path}_R(u', v') - \{u', v'\})$ that contain u and v , respectively (Definition 6.5). Then, the following observation is immediate.

Observation 7.3. *The path $\text{path}_R(u', v')$ is contained in $\text{Ring}(S_u, S_v)$ and it traverses $\text{Ring}(S_u, S_v)$ from $u' \in S_u$ to $v' \in S_v$. Moreover, $A_{R,P,u}^* - \{u'\}$ is contained in the bounded region of $\mathbb{R}^2 - H[S_u]$, and similarly $A_{R,P,v}^* - \{v'\}$ is contained in the unbounded region of $\mathbb{R}^2 - H[S_v]$.*

Now, recall the encircling tight sequence of concentric cycles $\mathcal{C}(u, v)$ and the path witnessing its tightness, which we denote by $\eta(u, v)$ (constructed in Lemma 6.6). We assume that $\eta(u, v)$ consists of only the 0-th copies of the edges comprising it, and fix $\eta(u, v)$ as the reference path of the ring $\text{Ring}(S_u, S_v)$. Moreover, observe that the subpath $\text{path}_R(u', v')$ of R also traverses $\text{Ring}(S_u, S_v)$. We assume w.l.o.g. that $\text{path}_R(u', v')$ consists of only the 0-th copy of the edges comprising it. This will allow us to later define winding numbers with respect to $\text{path}_R(u', v')$ in $\text{Ring}(S_u, S_v)$. Note that $\eta(u, v)$ and $\text{path}_R(u', v')$ may be different paths with common edges, and further they may not be paths in G .

Let us first argue that we can consider the rings corresponding to each of the long degree-2 paths in R independently. To this end, consider another long maximal degree-2 path in R different from P , denoted by $\text{path}_R(\hat{u}, \hat{v})$, where $\hat{u}, \hat{v} \in V_{=1}(R) \cup V_{\geq 3}(R)$ and t^* lies in the subtree containing \hat{v} in $R - V(\text{path}_R(\hat{u}, \hat{v}) - \{\hat{u}, \hat{v}\})$. Let $S_{\hat{u}} = \text{Sep}_R(\text{path}_R(\hat{u}, \hat{v}), \hat{u})$ and $S_{\hat{v}} = \text{Sep}_R(\text{path}_R(\hat{u}, \hat{v}), \hat{v})$. The $\text{Ring}(S_{\hat{u}}, S_{\hat{v}})$ is symmetric to $\text{Ring}(S_u, S_v)$ above. The following corollary follows from Lemma 6.10.

Corollary 7.1. *The rings $\text{Ring}(S_u, S_v)$ and $\text{Ring}(S_{\hat{u}}, S_{\hat{v}})$ have no common vertices.*

Let t be the number of pairs of near vertices in $V_{=1}(R) \cup V_{\geq 3}(R)$ such that the path between them in R is long. As above, we have a ring corresponding to each of these paths. Then, we partition the plane graph H into $t + 1$ regions, one for each of the t rings of the long maximal degree-2 path in R , and the remainder of H that is not contained in any of these rings.

Lemma 7.2. *Let $\{u_1, v_1\}, \{u_2, v_2\}, \dots, \{u_t, v_t\}$ denote pairs of near vertices in $V_{=1}(R) \cup V_{\geq 3}(R)$ such that $\text{path}_R(u_i, v_i)$ is a long maximal degree-2 path in R for all $i \in \{1, 2, \dots, t\}$. Then, the corresponding rings $\text{Ring}(S_{u_1}, S_{v_1}), \text{Ring}(S_{u_2}, S_{v_2}), \dots, \text{Ring}(S_{u_t}, S_{v_t})$ are pairwise disjoint. Further, the number of vertices of R lying outside this collection of rings, $|V(R) \setminus \bigcup_{i=1}^t V(S_{u_i}, S_{v_i})|$, is upper bounded by $\alpha_{\text{nonRing}}(k) = 10^4 k \cdot 2^{ck}$.*

Proof. The first statement follows from Corollary 7.1 applied to each pair of rings. For the second statement, consider the vertices in $V(R) \setminus \bigcup_{i=1}^t V(S_{u_i}, S_{v_i})$. Each such vertex either belongs to a short degree-2 path between a pair of near vertices in $V_{=1}(R) \cup V_{\geq 3}(R)$, or it is a vertex that lies on a maximal degree-2 path of R at distance at most $\alpha_{\text{pat}}(k) = 100 \cdot 2^{ck}$ in R from $V_{=1}(R) \cup V_{\geq 3}(R)$. Recall that there are fewer than $4k$ vertices in $V_{=1}(R) \cup V_{\geq 3}(R)$ by Observation 6.1, and thus at most $4k - 2$ maximal degree-2 paths in R . Therefore, the number of vertices in $V_{=1}(R) \cup V_{\geq 3}(R)$ is upper bounded by $(4k - 2) \cdot \max\{2\alpha_{\text{pat}}(k), \alpha_{\text{long}}(k)\} \leq 10^4 k \cdot 26ck$. \square

7.4 Solutions with Small Winding Number

Having established all the required definitions in the previous subsections, we are ready to exhibit a solution with a small winding number. Consider a ring, say $\text{Ring}(S_u, S_v)$ with an encircling tight concentric sequence of cycles $\mathcal{C}(u, v)$ and reference path $\eta(u, v)$. Consider a linkage \mathcal{P} in G , and let $\mathcal{P}(u, v) = \mathcal{P}[V(S_u, S_v) \cap V(\mathcal{P})]$. Observe that $\mathcal{P}(u, v)$ is a linkage in $\text{Ring}(S_u, S_v)$, and its endpoints lie in $(S_u \cup S_v) \cap V(G)$. Therefore, $\mathcal{P}(u, v)$ is a flow from $S_u \cap V(G)$ to $S_v \cap V(G)$ in $G_{u,v}$. Assume w.l.o.g. that the paths in $\mathcal{P}(u, v)$ use the 1-st copy of each edge in H .

Definition 7.5. *Let $u, v \in V_{=1}(R) \cup V_{\geq 3}(R)$ be a pair of near vertices such that $\text{path}_R(u, v)$ is long maximal degree-2 path in R . Let \mathcal{P} be a path system in G . Then, winding number of \mathcal{P} in $\text{Ring}(S_u, S_v)$ is defined as $\text{WindNum}(\mathcal{P}, \text{Ring}(S_u, S_v)) = \max_{P \in \mathcal{P}(u, v)} |\overline{\text{WindNum}}(P, \text{path}_R(u', v'))|$*

We remark that the notation above differs from our earlier use of $\overline{\text{WindNum}}(\cdot)$ in the choice of the reference path in $\text{Ring}(S_u, S_v)$. For $\overline{\text{WindNum}}(\cdot)$, the path $\eta(u, v)$ was the reference path, whereas for $\text{WindNum}(\cdot, \text{Ring}(S_u, S_v))$, we choose $\text{path}_R(u', v')$ as the reference path.

We will now apply Lemma 7.1 in each ring of the form $\text{Ring}(S_{u_i}, S_{v_i})$ to obtain a solution of bounded winding number in that ring. Towards this, fix one pair $(u, v) := (u_i, v_i)$ for some $1 \leq i \leq t$. We argue that $\mathcal{C}(u, v)$ and $\text{Flow}_R(u, v)$ are suitable for the roles of \mathcal{C} and \mathcal{Q} in the premise of Lemma 7.1. Recall that $\text{Flow}_R(u, v)$ is a collection of vertex disjoint paths in the subgraph G_{u_i, v_i} of G , and assume w.l.o.g. that $\text{Flow}_R(u, v)$ uses only 2-nd copies of edges in H .

Observation 7.4. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths. Let R be a backbone Steiner tree. Let u, v be a pair of near vertices in $V_{=1}(R) \cup V_{\geq 3}(R)$ such that $\text{path}_R(u, v)$ is a long degree-2 path in R . Then, the following hold.*

(i) $\mathcal{C}(u, v)$ is a encircling tight sequence of concentric cycles in $\text{Ring}(S_u, S_v)$ where each cycle lies in $G_{u,v}$, S_u is in its strict interior and S_v is in its strict exterior. Further, $|\mathcal{C}(u, v)| \geq 40\alpha_{\text{sep}}(k)$ cycles.

(ii) $\text{Flow}_R(u, v)$ is a maximum flow between $S_u \cap V(G)$ and $S_v \cap V(G)$ in $G_{u,v}$ that is minimal with respect to $\mathcal{C}(u, v)$.

(iii) Each path $Q \in \text{Flow}_R(u, v)$ traverses $\text{Ring}(S_u, S_v)$, and $|\overline{\text{WindNum}}(\text{path}_R(u', v'), Q)| \leq 1$.

Proof. The first statement directly follows from the construction of $\mathcal{C}(u, v)$ (see Lemma 6.6). Similarly, the second statement directly follows from the construction of $\text{Flow}_R(u, v)$ using $\mathcal{C}(u, v)$ (see Observation 6.8). Additionally the construction implies that each path $Q \in \text{Flow}_R(u, v)$ traverses $\text{Ring}(S_u, S_v)$. For the second part of the third statement, first note any $Q \in \text{Flow}_R(u, v)$ and $\text{path}_R(u', v')$ are two edge-disjoint traversing paths in $\text{Ring}(S_u, S_v)$. Hence, $\overline{\text{WindNum}}(\text{path}_R(u', v'), P)$ is well defined. Since for any $Q \in \text{Flow}_R(u, v)$, there are at most two edges in $\text{path}_R(u', v')$ with only one endpoint in $V(Q)$ (by Corollary 6.1, the absolute value of the signed sum of crossing between these two paths is upper-bounded by 1, i.e. $|\overline{\text{WindNum}}(Q, \text{path}_R(u', v'))| \leq 1$. \square

Finally we are ready to prove that the existence of a solution of small winding number.

Lemma 7.3. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths. Let R be a backbone Steiner tree. Let $\{u_1, v_1\}, \{u_2, v_2\}, \dots, \{u_t, v_t\}$ be the pairs of near vertices in $V_{=1}(R) \cup V_{\geq 3}(R)$ such that $\text{path}_R(u_i, v_i)$ is a long maximal degree-2 path in R for all $i \in \{1, 2, \dots, t\}$. Let $\text{Ring}(S_{u_1}, S_{v_1}), \text{Ring}(S_{u_2}, S_{v_2}), \dots, \text{Ring}(S_{u_t}, S_{v_t})$ be the corresponding rings. Then, there is a solution \mathcal{P}^* to (G, S, T, g, k) such that $|\overline{\text{WindNum}}(\mathcal{P}^*, \text{Ring}(S_{u_i}, S_{v_i}))| \leq \alpha_{\text{winding}}(k)$ for all $i \in \{1, 2, \dots, t\}$, where $\alpha_{\text{winding}}(k) = 60\alpha_{\text{sep}}(k) + 11 < 300 \cdot 2^{ck}$.*

Proof. Consider a solution \mathcal{P} to (G, S, T, g, k) . Fix a pair $(u, v) := (u_i, v_i)$ for some $i \in \{1, 2, \dots, t\}$. Recall that $|S_u|, |S_v| \leq \ell$ where $\ell = \alpha_{\text{sep}}(k)$. Consider $\text{Ring}(S_u, S_v)$, and the linkage $\mathcal{P}(u, v)$ in G . Our goal is to modify $\mathcal{P}(u, v)$ to obtain another linkage $\mathcal{P}'(u, v)$ that is aligned with it and has a small winding number with respect to $\text{path}_R(u, v)$

Recall that by Observation 7.4, we have a encircling tight sequence of concentric cycles $\mathcal{C}(u, v)$ in $\text{Ring}(S_u, S_v)$ that contains at least 40ℓ cycles. Let $\eta(u, v)$ be the path in this ring witnessing the tightness of $\mathcal{C}(u, v)$. Further, recall the linkage $\text{Flow}_R(u, v)$ between S_u and S_v in the subgraph $G_{u,v}$ of G (the restriction of G to $\text{Ring}(S_u, S_v)$), and that $\text{Flow}_R(u, v)$ is minimal with respect to $\mathcal{C}(u, v)$. We can assume w.l.o.g. that $\mathcal{P}(u, v)$ is minimal with respect to $\mathcal{C}(u, v)$. Otherwise, there is another solution $\widehat{\mathcal{P}}$ such that it is identical to \mathcal{P} in $G - V(S_u, S_v)$ and $\widehat{\mathcal{P}}(u, v)$ is a minimal linkage with respect to $\mathcal{C}(u, v)$ (that is aligned with $\mathcal{P}(u, v)$). Then, we can consider $\widehat{\mathcal{P}}$ instead of \mathcal{P} . Let $\mathcal{P}_{\text{traverse}}(u, v)$ be the set of traversing paths in $\mathcal{P}(u, v)$. Since $\mathcal{P}_{\text{traverse}}(u, v)$ is a flow between S_u and S_v in $G_{u,v}$ and $\text{Flow}_R(u, v)$ is a maximum flow, clearly $|\text{Flow}_R(u, v)| \geq |\mathcal{P}_{\text{traverse}}(u, v)|$.

Now, apply Lemma 7.1 to $\mathcal{P}(u, v)$, $\mathcal{C}(u, v)$ and $\text{Flow}_R(u, v)$ in $\text{Ring}(S_u, S_v)$. We thus obtain a linkage $\mathcal{P}'_{\text{traverse}}(u, v)$ disjoint from $\mathcal{P}_{\text{visitor}}(u, v)$ that is aligned with $\mathcal{P}_{\text{traverse}}(u, v)$. Hence, $\mathcal{P}'(u, v) = \mathcal{P}_{\text{visitor}}(u, v) \cup \mathcal{P}'_{\text{traverse}}(u, v)$ is a linkage in G aligned with $\mathcal{P}(u, v)$. Assume w.l.o.g. that $\mathcal{P}'(u, v)$ uses the 3-rd copy of each edge in H . Let us now consider the winding number of $\mathcal{P}'_{\text{traverse}}(u, v)$ with respect to $\text{path}_R(u', v')$. By Lemma 7.1(c), $|\overline{\text{WindNum}}(\mathcal{P}'_{\text{traverse}}(u, v)) - \overline{\text{WindNum}}(\text{Flow}_R(u, v))| \leq 60\ell + 6$. Now, note that for any path $P' \in \mathcal{P}'_{\text{traverse}}(u, v)$, $|\overline{\text{WindNum}}(P') - \overline{\text{WindNum}}(\mathcal{P}'_{\text{traverse}}(u, v))| \leq 1$, (recall that the winding number of paths in a linkage differ by at most 1). Similarly, for any path $Q \in \text{Flow}_R(u, v)$, $|\overline{\text{WindNum}}(\text{Flow}_R(u, v)) - \overline{\text{WindNum}}(Q)| \leq 1$. Therefore, it follows that for any two paths $P' \in \mathcal{P}'_{\text{traverse}}(u, v)$ and $Q \in \text{Flow}_R(u, v)$ $|\overline{\text{WindNum}}(P') - \overline{\text{WindNum}}(Q)| \leq 60\ell + 8$. Recall that we chose $\eta(u, v)$ as the reference path of $\text{Ring}(S_u, S_v)$ in the above expression. Hence, we may rewrite it as $|\overline{\text{WindNum}}(P', \eta(u, v)) - \overline{\text{WindNum}}(Q, \eta(u, v))| \leq 60\ell + 8$. Note that P', Q and $\eta(u, v)$ are three edge-disjoint paths

traversing $\text{Ring}(S_u, S_v)$. Hence $\overline{\text{WindNum}}$ is well defined in $\text{Ring}(S_u, S_v)$ for any pair of them. By Proposition 7.1, $|\overline{\text{WindNum}}(P', Q)| \leq |\overline{\text{WindNum}}(P', \eta(u, v)) - \overline{\text{WindNum}}(Q, \eta(u, v))| + 1$. We have so far established that for any $P' \in \mathcal{P}'_{\text{traverse}}(u, v)$ and $Q \in \text{Flow}_R(u, v)$, $|\overline{\text{WindNum}}(P', Q)| \leq 60\ell + 9$. Now, consider $\text{path}_R(u', v')$ and recall that for any $Q \in \text{Flow}_R(u, v)$, $|\overline{\text{WindNum}}(\text{path}_R(u', v'), Q)| \leq 1$ by Observation 7.4. Furthermore $\text{path}_R(u', v')$ uses the 0-th copies of edges in H . Hence, for each $P' \in \mathcal{P}'_{\text{traverse}}(u, v)$ $\overline{\text{WindNum}}(P', \text{path}_R(u', v'))$ is well defined in $\text{Ring}(S_u, S_v)$, and by Proposition 7.1, $|\overline{\text{WindNum}}(P', \text{path}_R(u', v'))| \leq |\overline{\text{WindNum}}(P', Q) - \overline{\text{WindNum}}(\text{path}_R(u', v'), Q)| + 1 \leq 60\ell + 11$.

Finally, consider the paths in $\mathcal{P}'_{\text{visitor}} = \mathcal{P}_{\text{visitor}}$. For each $P \in \mathcal{P}_{\text{visitor}}$ the absolute value of the winding number of $\overline{\text{WindNum}}(P, \text{path}_R(u', v'))$ is bounded by 1 (by Observation 7.1). Hence, we conclude that $\text{WindNum}(\mathcal{P}', \text{Ring}(S_u, S_v)) \leq 60\alpha_{\text{sep}}(k) + 11 < 300 \cdot 2^{ck}$. \square

8 Pushing a Solution onto the Backbone Steiner Tree

In this section, we push a linkage that is a solution of small winding number onto the backbone Steiner tree to construct a “pushed weak linkage” with several properties that will make its reconstruction (in Section 9) possible. Let us recall that we have an instance (G, S, T, g, k) and H is the radial completion of G enriched with $4n + 1$ parallel copies of each edge. Then we construct a backbone Steiner tree R (Section 6), which uses the 0-th copy of each edge. Formally, a pushed weak linkage is defined as follows.

Definition 8.1 (Pushed Weak Linkage). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. A weak linkage \mathcal{W} in H is pushed (onto R) if $E(\mathcal{W}) \cap E(R) = \emptyset$ and every edge in $E(\mathcal{W})$ is parallel to some edge of R .*

In what follows, we first define the properties we would like the pushed weak linkage to satisfy. Additionally, we partition any weak linkage into segments that give rise to a potential function whose maintenance will be crucial while pushing a solution of small winding number onto R . Afterwards, we show how to push the solution, simplify it, and analyze the result. We remark that the simplification (after having pushed the solution) will be done in two stages.

8.1 Desired Properties and Potential of Weak Linkages

The main property we would like a pushed weak linkage to satisfy is to use only “few” parallel copies of any edge. This quantity is captured by the following definition of multiplicity, which we will eventually like to upper bound by a small function of k .

Definition 8.2 (Multiplicity of Weak Linkage). *Let H be a plane graph. Let \mathcal{W} be a weak linkage in H . Then, the multiplicity of \mathcal{W} is the maximum, across all edges e in H , of the number of edges parallel to e that belong to $E(\mathcal{W})$.*

Towards bounding the multiplicity of the pushed weak linkage we construct, and also an important ingredient on its own in the reconstruction in Section 9, we need to repeatedly eliminate U-turns in the weak linkage we deal with. Here, U-turns are defined as follows.

Definition 8.3 (U-Turn in Weak Linkage). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a weak linkage in H . Then, a U-turn in \mathcal{W} is a pair of parallel edges $\{e, e'\}$ visited consecutively by some walk $W \in \mathcal{W}$ such that the strict interior of the cycle formed by e and e' does not contain the first or last edge of any walk in \mathcal{W} . We say that \mathcal{W} is U-turn-free if it does not have any U-turn.*

Still, having a pushed weak linkage of low multiplicity and no U-turns does not suffice for faithful reconstruction due to ambiguity in which edge copies are being used. This ambiguity will be dealt with by the following definition.

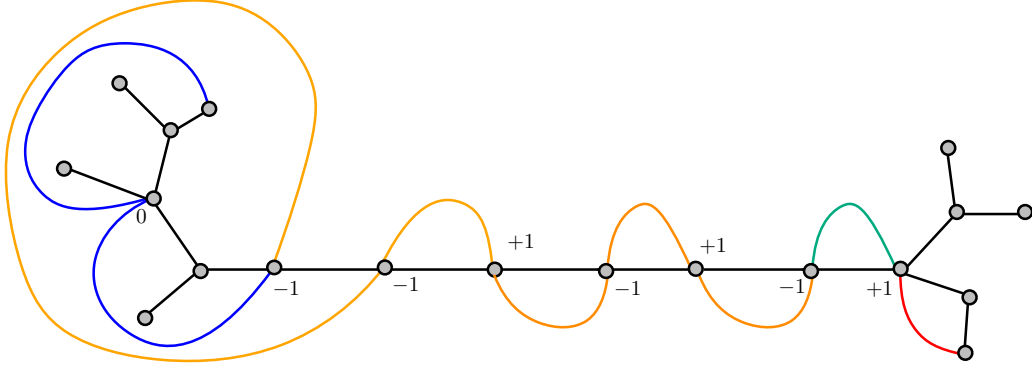


Figure 19: Segments, segment groups and their labeling.

Definition 8.4 (Canonical Weak Linkage). Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a weak linkage in H pushed onto R . Then, \mathcal{W} is canonical if (i) for every edge $e_i \in E(\mathcal{W})$, $i \geq 1$ (ii) if $e_i \in E(\mathcal{W})$, then all the parallel edges e_j , for $1 \leq j < i$, are also in $E(\mathcal{W})$.

For brevity, we say that a weak linkage \mathcal{W} in H is *simplified* if it is sensible, pushed onto R , canonical, U-turn-free and has multiplicity upper bounded by $\alpha_{\text{mul}}(k) := 2\alpha_{\text{potential}}(k)$ where $\alpha_{\text{potential}}(k) = 2^{\mathcal{O}(k)}$ will be defined precisely in Lemma 8.2. For the process that simplifies a pushed weak linkage, we will maintain a property that requires multiplicity at most $2n$ as well as a relaxation of canonicity. This property is defined as follows.

Definition 8.5 (Extremal Weak Linkage). Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a weak linkage in H that is pushed onto R . Then, \mathcal{W} is extremal if its multiplicity is at most $2n$ and for any two parallel edges $e_i, e_j \in E(\mathcal{W})$ where $i \geq 1$ and $j \leq -1$, we have $(i - 1) + |j + 1| \geq 2n$.

Additionally, we will maintain the following property.

Definition 8.6 (Outer-Terminal Weak Linkage). Let (G, S, T, g, k) be a nice instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a weak linkage in H . Then, \mathcal{W} is outer-terminal if it uses exactly one edge incident to t^* .

Segments, Segment Groups and Potential. To analyze the “complexity” of a weak linkage, we partition it into segments and segment groups, and then associate a potential function with it based on this partition. Intuitively, a segment of a walk is a maximal subwalk that does not cross R (see Fig. 19). Formally, it is defined as follows.

Definition 8.7 (Segment). Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let W be a walk in H that is edge-disjoint from R . A crossing of W with R is a crossing $(v, e, \hat{e}, e', \hat{e}')$ of W and some path in R .¹³ Then, a segment of W is a maximal subwalk of W that has no crossings with R . Let $\text{Seg}(W)$ denote the set¹⁴ of segments of W .

We remind that R only contains 0-copies of edges, hence we can ensure that we deal with walks that are edge-disjoint from R by avoiding the usage of 0-copies. Towards the definition of potential for a weak linkage, we group segments together as follows (see Fig. 19).

Definition 8.8 (Segment Group). Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let W be a walk in H that is edge-disjoint from R . A segment group of W is a maximal subwalk W' of W such that either (i) $\text{Seg}(W') \subseteq \text{Seg}(W)$ and all of the

¹³The path might not be a maximal degree-2 path, thus $(v, e, \hat{e}, e', \hat{e}')$ may concern a vertex $v \in V_{\geq 3}(R)$.

¹⁴Because we deal with walks that do not repeat edges, $\text{Seg}(W)$ is necessarily a set rather than a multiset.

endpoints of all of the segments in $\text{Seg}(W')$ are internal vertices of the same maximal degree-2 path of R , or (ii) $W' \in \text{Seg}(W)$ and the two endpoints of W' are not internal vertices of the same maximal degree-2 path in R .¹⁵ The set of segment groups of W is denoted by $\text{SegGro}(W)$.

Observe that the set of segments, as well as the set of segment groups, define a partition of a walk. We define the “potential” of a segment group based on its winding number in the ring that corresponds to its path (in case it is a long path where a ring is defined). To this end, recall the labeling function in Definition 7.4. Note that the labeling is defined for any two walks irrespective of the existence of a ring.

Definition 8.9 (Potential of Segment Group). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let W be a walk in H that is edge-disjoint from R and whose endpoints are in $V_{=1}(R)$. Let $W' \in \text{SegGro}(W)$. If $|\text{Seg}(W')| = 1$, then the potential of W' , denoted by $\text{Potential}(W')$, is defined to be 1. Otherwise, it is defined as follows.*

$$\text{Potential}(W') = 1 + \left| \sum_{(e,e') \in E(W^*) \times E(W^*)} \text{label}_P^{W'}(e, e') \right|,$$

where W^* is the walk obtained from W' by adding two edges to W' —the edge consecutive to the first edge of W' in W and the edge consecutive to the last edge of W' in W , and P is the maximal degree-2 path of R such that all of the endpoints of all of the segments in $\text{Seg}(W')$ are its internal vertices.

The potential of a segment group is well defined as we use the function `label` only for edges incident to internal vertices of the maximal degree-2 paths in R . For an example of a potential of a segment groups, see Fig. 19. Now, we generalize the notion of potential from segment groups to weak linkages as follows.

Definition 8.10 (Potential of Weak Linkage). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a weak linkage. Then, the potential of \mathcal{W} is*

$$\text{Potential}(\mathcal{W}) = \sum_{W' \in \text{SegGro}(\mathcal{W})} \text{Potential}(W'),$$

where $\text{SegGro}(\mathcal{W}) = \bigcup_{W \in \mathcal{W}} \text{SegGro}(W)$.

To upper bound the potential of a solution of a small winding number, we first upper bound the number of segment groups.

Lemma 8.1. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths, and R be a backbone Steiner tree. Then, there exists a solution \mathcal{P} to (G, S, T, g, k) such that $|\text{SegGro}(\mathcal{P})| \leq \alpha_{\text{segGro}}(k) := 10^5 k \cdot 2^{ck}$.*

Proof. First, notice that for every path $P \in \mathcal{P}$, every segment group of P whose endpoints are both internal vertices of some maximal degree-2 path of R has the following property: it is neither a prefix nor a suffix of P (when P is oriented arbitrarily, say, from its endpoint in S to its endpoint in T), and the segment groups that appear immediately before and after it necessarily satisfy that each of them does not have both of its endpoints being internal vertices of some maximal degree-2 path of R . Let s denote the number of segment groups of \mathcal{P} that do not have both of their endpoints being internal vertices of some maximal degree-2 path of R . Then, we have that $|\text{SegGro}(\mathcal{P})| \leq 2s - 1$, and therefore to complete the proof, it suffices to show that $s \leq \alpha_{\text{segGro}}(k)/2$.

¹⁵That is, the two endpoints of W' are internal vertices in different maximal degree-2 paths in R or at least one endpoint of W' is a vertex in $V_{=1}(R) \cup V_{\geq 3}(R)$.

Let $\{u_1, v_1\}, \{u_2, v_2\}, \dots, \{u_t, v_t\}$ denote the pairs of near vertices in $V_{=1}(R) \cup V_{\geq 3}(R)$ such that for each $i \in \{1, \dots, t\}$, $\text{path}_R(u_i, v_i)$ is a long maximal degree-2 path in R . Let $\text{Ring}(S_{u_1}, S_{v_1}), \text{Ring}(S_{u_2}, S_{v_2}), \dots, \text{Ring}(S_{u_t}, S_{v_t})$ be the corresponding rings. By Lemma 7.2, the number of vertices of R lying outside these rings, $|V(R) \setminus \bigcup_{i=1}^t V(S_{u_i}, S_{v_i})|$, is upper bounded by $\alpha_{\text{nonRing}}(k) = 10^4 k \cdot 2^{ck}$. We further classify each segment group S of \mathcal{P} that does not have both of its endpoints being internal vertices of some maximal degree-2 path of R as follows. (We remark that, by the definition of a segment group, S must consist of a single segment.)

- S has at least one endpoint in $V(R) \setminus \bigcup_{i=1}^t V(S_{u_i}, S_{v_i})$. Denote the number of such segment groups by s_1 .
- S has one endpoint in $V(S_{u_i}, S_{v_i})$ and another endpoint in $V(S_{u_j}, S_{v_j})$ for some $i, j \in \{1, \dots, t\}$ such that $i \neq j$. Denote the number of such segment groups by s_2 .

Then, $s_1 + s_2 = s$. Now, notice that the paths in \mathcal{P} are pairwise vertex-disjoint, and the collection of segment groups of each path $P \in \mathcal{P}$ forms a partition of P . Thus, because $|V(R) \setminus \bigcup_{i=1}^t V(S_{u_i}, S_{v_i})| \leq \alpha_{\text{nonRing}}(k)$ and each vertex in $V(R)$ is shared as an endpoint by at most two segment groups, we immediately derive that $s_1 \leq \alpha_{\text{nonRing}}(k)$. To bound s_2 , note that each segment group of the second type traverses at least one vertex in $S_{u_i} \cup S_{v_i}$ for some $i \in \{1, \dots, t\}$ (due to Lemma 6.12). By Lemma 6.4, for every $i \in \{1, \dots, t\}$, $|S_{u_i}|, |S_{v_i}| \leq \alpha_{\text{sep}}(k) = \frac{7}{2} \cdot 2^{ck} + 2$. Moreover, by Observation 6.1, $t < 4k$. Thus, we have that $s_2 \leq 4 \cdot 4k \cdot \alpha_{\text{sep}}(k)$, where the multiplication by 4 is done because two segment groups can share an endpoint and each maximal degree-2 path is associated with two separators. From this, we conclude that $s \leq 10^4 k \cdot 2^{ck} + 16k(\frac{7}{2} \cdot 2^{ck} + 2) \leq \alpha_{\text{segGro}}(k)/2$. \square

Now, based on Observation 7.2 and Lemmas 7.3 and 8.1, we derive the existence of a solution with low potential (if there exists a solution).

Lemma 8.2. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths, and R be a backbone Steiner tree. Then, there exists a solution \mathcal{P} to (G, S, T, g, k) such that $\text{Potential}(\mathcal{P}) \leq \alpha_{\text{potential}}(k) := (10^4 \cdot 4^{ck} + 1) \cdot \alpha_{\text{segGro}}(k)$.*

Proof. Let \mathcal{Q} be a solution with the property in Lemma 7.3. Let \mathcal{S} denote the segment groups in $\text{SegGro}(\mathcal{Q})$ whose both endpoints belong to the same maximal degree-2 path of R . For each segment group $S \in \mathcal{S}$, denote $\text{labPot}(S) = |\sum_{(e, e') \in E(S) \times E(S)} \text{label}_{P_S}^S(e, e')|$ where P_S is the path in \mathcal{Q} that has S as a segment group. Furthermore, denote $M = \max_{S \in \mathcal{S}} \text{labPot}(S)$. Now, by the definition of potential,

$$\text{Potential}(\mathcal{Q}) = |\text{SegGro}(\mathcal{Q})| + \sum_{S \in \mathcal{S}} \text{labPot}(S) \leq (M + 1)|\text{SegGro}(\mathcal{Q})|.$$

By Lemma 8.1, $|\text{SegGro}(\mathcal{Q})| \leq \alpha_{\text{segGro}}(k)$. Thus, to complete the proof, it suffices to show that $M \leq 10^4 \cdot 4^{ck}$. To this end, consider some segment group $S \in \mathcal{S}$. Notice that $\text{labPot}(S)$ is upper bounded by the number of crossings of S with P_S . Thus, if P_S is short, it follows that $\text{labPot}(S) < \alpha_{\text{long}}(k) = 10^4 \cdot 2^{ck}$.

Now, suppose that P_S is long, and let $\text{Ring}(S_u, S_v)$ be the ring that corresponds to P_S . By the choice of \mathcal{Q} , $|\text{WindNum}(\mathcal{Q}, \text{Ring}(S_u, S_v))| \leq \alpha_{\text{winding}}(k) < 300 \cdot 2^{ck}$. Let $\hat{\mathcal{A}}$ be collection of maximal subpaths of S that are fully contained within $\text{Ring}(S_u, S_v)$, and let \hat{P}_S be the maximal subpath of P_S that is fully contained within $\text{Ring}(S_u, S_v)$. Then, by Observation 7.2,

$$\text{labPot}(S) \leq |V(P_S) \setminus V(\hat{P}_S)| + \sum_{\hat{A} \in \hat{\mathcal{A}}} |\overline{\text{WindNum}}(\hat{A}, \hat{P}_S)|.$$

By the definition of WindNum , we have that $|\overline{\text{WindNum}}(\hat{A}, \hat{P}_S)| \leq |\text{WindNum}(\mathcal{Q}, \text{Ring}(S_u, S_v))|$ for each $\hat{A} \in \hat{\mathcal{A}}$. Moreover, by Lemma 6.4, $|\hat{\mathcal{A}}| \leq |S_u| + |S_v| \leq 2\alpha_{\text{sep}}(k) = 7 \cdot 2^{ck} + 4$. Additionally,

by Lemmas 6.5 and 6.11, we have that $|V(P_S) \setminus V(\widehat{P}_S)| \leq 2\alpha_{\text{pat}}(k) = 200 \cdot 2^{ck}$. From this,

$$\text{labPot}(S) \leq 200 \cdot 2^{ck} + (7 \cdot 2^{ck} + 4) \cdot 300 \cdot 2^{ck} \leq 10^4 \cdot 4^{ck}.$$

Thus, because the choice of S was arbitrary, we conclude that $M \leq 10^4 \cdot 4^{ck}$. \square

8.2 Pushing a Solution onto R

Let us now describe the process of pushing a solution onto R . To simplify this process, we define two “non-atomic” operations that encompass sequences of atomic operations in discrete homotopy. We remind that we only deal with walks that do not repeat edges.

Definition 8.11 (Non-Atomic Operations in Discrete Homotopy). *Let G be a triangulated plane graph with a weak linkage \mathcal{W} , and C be a cycle¹⁶ in G . Let $W \in \mathcal{W}$.*

- **Cycle Move.** *Applicable to (W, C) if there exists a subpath Q of C such that (i) Q is a subpath of W , (ii) $1 \leq |E(Q)| \leq |E(C)| - 1$, (iii) no edge in $E(C) \setminus E(Q)$ belongs to any walk in \mathcal{W} , and (iii) no edge drawn in the strict interior of C belongs to any walk in \mathcal{W} . Then, the cycle move operation replaces Q in W by the unique subpath of C between the endpoints of Q that is edge-disjoint from Q .*
- **Cycle Pull.** *Applicable to (W, C) if (i) C is a subwalk Q of W , and (ii) no edge drawn in the strict interior of C belongs to any walk in \mathcal{W} . Then, the cycle pull operation replaces Q in W by a single occurrence of the first vertex in Q .*

We now prove that the operations above are compositions of atomic operations.

Lemma 8.3. *Let G be a triangulated plane graph with a weak linkage \mathcal{W} , and C be a cycle in G . Let $W \in \mathcal{W}$ with a non-atomic operation applicable to (W, C) . Then, the result of the application is a weak linkage that is discretely homotopic to \mathcal{W} .*

Proof. We prove the claim by induction on the number of faces of G in the interior of C . In the basis, where C encloses only one face, then the cycle move and cycle pull operations are precisely the face move and face pull operations, respectively, and therefore the claim holds. Now, suppose that C encloses $i \geq 2$ faces and that the claim is correct for cycles that enclose at most $i - 1$ faces. We consider several cases as follows.

First, suppose that C has a path P fully drawn in its interior whose endpoints are two (distinct) vertices $u, v \in V(C)$, and whose internal vertices and all of its edges do not belong to C . (We remark that P might consist of a single edge, and that edge might be parallel to some edge of C .) Now, notice that P partitions the interior of C into the interior of two cycles C_1 and C_2 that share only P in common as follows: one cycle C_1 consists of one subpath of C between u and v and the path P , and the other cycle C_2 consists of the second subpath of C between u and v and the path P . Notice that C_1 encloses less faces than C , and so does C_2 . At least one of these two cycles, say, C_1 , contains at least one edge of Q . Then, the cycle move operation is applicable to (W, C_1) . Indeed, let \widehat{Q} be the subpath of Q that is a subpath of C , and notice that $E(P) \subseteq E(C_1) \subseteq E(C) \cup E(P)$ and $E(P) \cap E(\mathcal{W}) = \emptyset$ (because the cycle move/pull operation is applicable to (W, C)). Therefore, \widehat{Q} is a subpath of W , $1 \leq |E(\widehat{Q})| \leq |E(C_1)| - 1$, and no edge in $E(C_1) \setminus E(\widehat{Q})$ belongs to any walk in \mathcal{W} . Moreover, because C_1 belongs to the interior (including the boundary) of C , no edge drawn in the strict interior of C belongs to any walk in \mathcal{W} . Now, notice that after the application of the cycle move operation for (W, C_1) , C_2 also has at least one edge used by the walk W' into which W was modified—in particular, $E(P) \subseteq E(W')$. Moreover, consider the subpath (or subwalk that is a cycle) Q' of W' that results from the replacement of \widehat{Q} in Q by the subpath of C_1 between the endpoints of Q' that

¹⁶A pair of parallel edges is considered to be a cycle.




Figure 20: A Sequence, its projecting cycle and a shrinking cycle.

does not belong to W . Then, Q' traverses some subpath (possibly empty) of C_1 or C_2 , then traverses P , and next traverses some other subpath of C_1 or C_2 . So, the restriction of Q' to C_2 is a non-empty path or cycle Q^* that is a subwalk of W' . Furthermore, because C_2 is drawn in the interior of C and the cycle move/pull operation is applicable to (W, C) , we have that no edge of $E(C_2) \setminus E(Q^*)$ or the strict interior of C_2 belongs to $E(W)$. Thus, the cycle move/pull operation is applicable to (W', C_2) . Now, the result of the application of this operation is precisely the result of the application of the original cycle move or pull operation applicable to (W, C) . To see this, observe that the edges of $E(C) \setminus E(W)$ that occur in C_1 along with $E(P)$ have replaced the edges of $E(C) \cap E(W)$ that occur in C_1 in the first operation, and the edges of $E(C) \setminus E(W)$ that occur in C_2 have replaced the edges of $E(C) \cap E(W)$ that occur in C_1 along with $E(P)$ in the second operation. Thus, by the inductive hypothesis with respect to (W, C_1) and (W', C_2) , and because discrete homotopy is transitive, the claim follows.

Thus, it remains to prove that C has a path P fully drawn in its interior whose endpoints are two (distinct) vertices $u, v \in V(C)$, and whose internal vertices and all of its edges do not belong to C . In case C has a chord (that is, an edge in G between two vertices of C that does not belong to C), then the chord is such a path P . Therefore, we now suppose that this is not the case. Then, C does not contain in its interior an edge parallel to an edge of C . In turn, because G is triangulated, when we consider some face f in the interior of C that contains an edge e of C , this face must be a triangle. Moreover, the vertex of f that is not incident to e cannot belong to C , since otherwise we obtain a chord in C . Thus, the subpath (that consists of two edges) of f between the endpoints of e that does not contain e is a path P with the above mentioned properties. \square

In the process of pushing a solution onto R , we push parts of the solution one-by-one. We refer to these parts as sequences, defined as follows (see Fig. 20).

Definition 8.12 (Sequence). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let W be a walk. Then, a sequence of W is a maximal subwalk of W whose internal vertices (if any exists) do not belong to R and which contains at least one edge that is not parallel to an edge of R . The set of sequences of W is denoted by $\text{Seq}(W)$. For a weak linkage \mathcal{W} , the set of sequences of \mathcal{W} is defined as $\text{Seq}(\mathcal{W}) = \bigcup_{W' \in \mathcal{W}} \text{Seq}(W')$.*

Notice that the set of sequences of a walk does not necessarily form a partition of the walk because the walk can traverse edges parallel to the edges of R and these edges do not belong to any sequence. Moreover, for sensible weak linkages, the endpoints of every sequence belong to R . To deal only with sequences that are paths, we need the following definition.

Definition 8.13 (Well-Behaved Weak Linkage). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. A weak linkage \mathcal{W} is well-behaved if every sequence in $\text{Seq}(\mathcal{W})$ is a path or a cycle.*

When we will push sequences one-by-one, we ensure that the current sequence to be pushed can be handled by the cycle move operation in Definition 8.11. To this end, we define the notion of an innermost sequence, based on another notion called a projecting cycle (see Fig. 20) We remark that this cycle will not necessarily be the one on which we apply a cycle move operation, since this cycle might contain in its interior edges of some walks of the weak linkage.

Definition 8.14 (Projecting Cycle). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a sensible well-behaved weak linkage, and $S \in \text{Seq}(\mathcal{W})$. The projecting cycle of S is the cycle C formed by S and the subpath P of R between the endpoints of S . Additionally, $\text{Volume}(S)$ denotes the number of faces enclosed by the projecting cycle of S .*

Now, we define the notion of an innermost sequence.

Definition 8.15 (Innermost Sequence). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a sensible well-behaved weak linkage, and $S \in \text{Seq}(\mathcal{W})$. Then, S is innermost if every edge in $E(\mathcal{W})$ that is drawn in the interior of its projecting cycle is parallel to some edge of R .*

We now argue that, unless the set of sequences is empty, there must exist an innermost one.

Lemma 8.4. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a sensible well-behaved weak linkage such that $\text{Seq}(\mathcal{W}) \neq \emptyset$. Then, there exists an innermost sequence in $\text{Seq}(\mathcal{W})$.*

Proof. Let $S \in \text{Seq}(\mathcal{W})$ be a sequence that minimizes $\text{Volume}(S)$. We claim that S is innermost. Suppose, by way of contradiction, that this claim is false. Then, there exists an edge $e \in E(\mathcal{W})$ that is drawn in the interior of the projecting cycle of S and is not parallel to any edge of R . Thus, e belongs to some sequence $S' \in \text{Seq}(\mathcal{W})$. Because \mathcal{W} is well-behaved, S and S' are vertex disjoint. This implies that the projecting cycle of S' is contained in the interior of the projecting cycle of S . Because H is triangulated, this means that $\text{Volume}(S') < \text{Volume}(S)$, which is a contradiction to the choice of S . \square

When we push the sequence onto R , we need to ensure that we have enough copies of each edge of R to do so. To this end, we need the following definition.

Definition 8.16 (Shallow Weak Linkage). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a sensible well-behaved weak linkage. Then, \mathcal{W} is shallow if for every edge $e_0 \in E(R)$, the following condition holds. Let ℓ (resp. h) be the number of sequences $S \in \text{Seq}(\mathcal{W})$ whose projecting cycle encloses e_1 (resp. e_{-1}). Then, e_i is not used by \mathcal{W} for every $i \in \{-n, -n+1, \dots, -n+\ell-1\} \cup \{0\} \cup \{n-h+1, n-h+2, \dots, n\}$.*

To ensure that we make only cycle moves/pulls as in Definition 8.11, we do not necessarily push the sequence at once, but gradually shrink the area enclosed by its projecting cycle.¹⁷

Definition 8.17 (Shrinking Cycle). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a sensible well-behaved weak linkage, and $S \in \text{Seq}(\mathcal{W})$ with an endpoint $v \in V(R)$. Then, a cycle C in H is a shrinking cycle for (S, v) if it has no edge of R in its interior and it can be partitioned into three paths where the first has at least one edge and the last has at most one edge: (i) a subpath P_1 of S with v as an endpoint; (ii) a subpath P_2 from the other endpoint u of P_1 to a vertex on R that consists only of edges drawn in the strict interior of the projecting cycle of S and no vertex of S apart from u ; (iii) a subpath P_3 that has v as an endpoint and whose edge (if one exists) is either not parallel to any edge of R or it is the i -th copy of an edge parallel to some edge of R for some $i \in \{-n+\ell-1, n-h+1\}$, where ℓ and h are as in Definition 8.16.*




Figure 21: An illustration of Lemma 8.5

With respect to shrinking cycles, we prove two claims. First, we assert their existence.

Lemma 8.5. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a sensible well-behaved weak linkage, and $S \in \text{Seq}(\mathcal{W})$ with an endpoint $v \in V(R)$. Then, there exists a shrinking cycle for (S, v) .*

Proof. Let e be an edge in S incident to v (if there are two such edges, when S is a cycle, pick one of them arbitrarily), and denote the other endpoint of e by u . Because H is triangulated, e belongs to the boundary B of a face f of H in the interior of C such that B is a cycle that consists of only two or three edges. If B does not contain any vertex of $V(R) \cup V(S)$ besides u and v , then it is clearly a shrinking cycle (see Fig. 21). Thus, we now suppose that B is a cycle on three edges whose third vertex, w , belongs to $V(R) \cup V(S)$. If $w \in V(S)$, then the cycle that consists of the subpath of S from v to w and the edge in $E(B)$ between v and w is also clearly a shrinking cycle (see Fig. 21). Thus, we now also suppose that $w \in V(R)$.

We further distinguish between two cases. First, suppose that w is not adjacent to v on R . In this case, B does not enclose any edge of R as well as any edge parallel to an edge of R . Moreover, B can be partitioned into P_1, P_2 and P_3 that are each a single edge, where P_1 consists of the edge in B between v and u , P_2 consists of the edge in B between u and w , and P_3 consists of the edge in B between w and v , thereby complying with Definition 8.17. Thus, B is a shrinking cycle for (S, v) . Now, suppose that w is adjacent to v on R . Then, define P_1, P_2 and P_3 similarly to before except that to P_3 , we do not take the edge of B between v and w but its parallel i -th copy where $i \in \{-n + \ell - 1, n - h + 1\}$ such that ℓ and h are as in Definition 8.16. The choice of whether $i = -n + \ell - 1$ or $i = n - h + 1$ is made so that the cycle B' consisting of P_1, P_2 and P_3 does not enclose any edge of R . (Such a choice necessarily exists, see Fig. 21). \square

Now, we prove that making a cycle move/pull operation using a shrinking cycle is valid and maintains some properties of weak linkages required for our analysis.

Lemma 8.6. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a sensible, well-behaved, shallow and outer-terminal weak linkage, and $S \in \text{Seq}(\mathcal{W})$ be innermost with an endpoint $v \in V(R) \setminus \{t^*\}$. Let C be a shrinking cycle for (S, v) that encloses as many faces as possible. Then, the cycle move/pull operation is applicable to (W, C) where $W \in \mathcal{W}$ is the walk having S as a sequence. Furthermore, the resulting weak linkage \mathcal{W}' is sensible, well-behaved, shallow and outer-terminal having the same potential as \mathcal{W} , and $\sum_{\hat{S} \in \text{Seq}(\mathcal{W}')} \text{Volume}(\hat{S}) < \sum_{\hat{S} \in \text{Seq}(\mathcal{W})} \text{Volume}(\hat{S})$.*

Proof. We first argue that the cycle move/pull operation is applicable to (W, C) . Let P_1, P_2 and P_3 be the partition of C in Definition 8.17. Note that because S is innermost, its projecting cycle does not contain any edge in $E(\mathcal{W})$ that is not parallel to some edge of R , and hence so does C as it is drawn in the interior (including the boundary) of the projecting cycle. Furthermore, the only edges parallel to R that C can contain are those parallel to the edge e_i of P_3 whose

¹⁷Instead, we could have also always pushed a sequence at once by defining moves and pulls for closed walks, which we find somewhat more complicated to analyze formally.

subscripts have absolute value larger than $|i|$. However, none of these edges belong to \mathcal{W} because $i \in \{-n + \ell - 1, n - h + 1\}$ where ℓ and h are as in Definition 8.16, and \mathcal{W} is shallow. Lastly, note that P_1 is either S (which might be a cycle) or a subpath of S , and hence it is a subwalk of W . Thus, the cycle move or pull (depending on whether P_1 is a cycle) is applicable to (W, C) . Furthermore, the new walk W' that results from the application is the modification of W that replaces P_1 by the path consisting of P_2 and P_3 .

Because \mathcal{W} is sensible and the endpoints of no walk in \mathcal{W} are changed in \mathcal{W}' , we have that \mathcal{W}' is sensible as well. Moreover, the vertices of P_2 are not used by any walk in \mathcal{W} apart from W' and only in its subwalk that traverses P_2 , and therefore, as \mathcal{W} is well-behaved, so is \mathcal{W}' . Additionally, note that \mathcal{W} is shallow and that each edge belongs to at most as many projecting cycles of sequences in \mathcal{W}' as it does in \mathcal{W} . Thus, if P_3 does not contain an edge (the only edge parallel to an edge of R that might be used by \mathcal{W}' but not by \mathcal{W} is the edge e_i of P_3 , if it exists), it is immediate that \mathcal{W}' is shallow. Now, suppose that e_i exists. Let $b \in \{-1, 1\}$ have the same sign as i . Recall that $i \in \{-n + \ell - 1, n - h + 1\}$ where ℓ and h are as in Definition 8.16, thus to conclude that \mathcal{W}' is shallow, we only need to argue that e_b belongs to the interior of fewer projecting cycles of sequences in \mathcal{W}' as it does in \mathcal{W} . However, this holds since the only difference between the sequences of \mathcal{W} and \mathcal{W}' is that the sequence S occurs in \mathcal{W} (and contains e_b in the interior of its projecting cycle), but is transformed into (one or two) other sequences in \mathcal{W}' , and these new sequences, by the definition of W' , no longer contain e_b in their projecting cycles. In this context, also note that the projecting cycles of the (one or two) new sequences enclose disjoint areas contained in the area enclosed by the projecting cycle of S , and the projecting cycles of the new sequences do not enclose the faces enclosed by C , but the projecting cycle of S does enclose them. Thus, $\sum_{\hat{S} \in \text{Seq}(\mathcal{W}')} \text{Volume}(\hat{S}) < \sum_{\hat{S} \in \text{Seq}(\mathcal{W})} \text{Volume}(\hat{S})$.

It remains to show that \mathcal{W}' is outer-terminal and that has the same potential as \mathcal{W} . The second claim is immediate since \mathcal{W} and \mathcal{W}' have precisely the same crossings with R . For the first claim, note that since \mathcal{W} is outer-terminal, it uses exactly one edge incident to t^* . The only vertex of R that can possibly be incident to more edges in $E(\mathcal{W}')$ than in $E(\mathcal{W})$ is the other endpoint, say, w , of the edge of P_3 in the case where P_3 contains an edge. So, suppose that P_3 does contain an edge and that $w = t^*$, else we are done. Since t^* is a leaf of R that belongs to the boundary of the outer-face of H , it cannot be enclosed in the strict interior of the projecting cycle of S and therefore it must be a vertex of S . However, this together with the maximality of the number of cycles enclosed by the shrinking cycle C implies that the C is equal to the projecting cycle of S . Thus, by the definition of W' , the only difference between the edges incident to t^* in \mathcal{W} compared to \mathcal{W}' is that in \mathcal{W} it is incident to an edge of S , while in \mathcal{W}' it is incident to the edge of P_3 . In particular, this means that \mathcal{W}' has exactly one edge incident to t^* and therefore it is outer-terminal. \square

Having Lemmas 8.2, 8.4, 8.5 and 8.6 at hand, we are ready to push a solution onto R . Since this part is only required to be existential rather than algorithmic, we give a simpler proof by contradiction rather than an explicit process to push the solution. Notice that once the solution has already been pushed, rather than using the notion of shallowness, we only demand to have multiplicity at most $2n$.

Lemma 8.7. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths, and R be a backbone Steiner tree. Then, there exists a sensible outer-terminal weak linkage \mathcal{W} in H that is pushed onto R , has multiplicity at most $2n$, discretely homotopic in H to some solution of (G, S, T, g, k) and $\text{Potential}(\mathcal{W}) \leq \alpha_{\text{potential}}(k)$.*

Proof. By Lemma 8.2, there exists a solution \mathcal{P} to (G, S, T, g, k) such that $\text{Potential}(\mathcal{P}) \leq \alpha_{\text{potential}}(k)$. Because the paths in \mathcal{P} are pairwise vertex-disjoint and \mathcal{P} has a path where t^* is an endpoint, it is clear that \mathcal{P} is a sensible, well-behaved, shallow and outer-terminal weak linkage. Since \mathcal{P} is discretely homotopic to itself, it is well defined to let \mathcal{W} a weak linkage that,

among all sensible, well-behaved, shallow and outer-terminal weak linkages that are discretely homotopic to \mathcal{P} , minimizes $\sum_{S \in \text{Seq}(\mathcal{W})} \text{Volume}(S)$. Notice that shallowness is a stronger demand than having multiplicity at most $2n$, and that being pushed onto R is equivalent to having an empty set of sequences. Thus, to conclude the proof, it suffices to argue that $\text{Seq}(\mathcal{W}) = \emptyset$.

Suppose, by way of contradiction, that $\text{Seq}(\mathcal{W}) \neq \emptyset$. Then, by Lemmas 8.4 and 8.5, there exist an innermost sequence $S \in \text{Seq}(\mathcal{W})$ and a shrinking cycle C for (S, v) where we pick v as an endpoint of S that is not t^* (because \mathcal{W} is outer-terminal, not both endpoints of S can be t^*), and we pick a shrinking cycle enclosing as many faces as possible. By Lemma 8.6, the cycle move/pull operation is applicable to (W, C) where $W \in \mathcal{W}$ is the walk having S as a sequence. Furthermore, the resulting weak linkage \mathcal{W}' is sensible, well-behaved, shallow and outer-terminal having the same potential as \mathcal{W} , and $\sum_{\hat{S} \in \text{Seq}(\mathcal{W}')} \text{Volume}(\hat{S}) < \sum_{\hat{S} \in \text{Seq}(\mathcal{W})} \text{Volume}(\hat{S})$. Since discrete homotopy is an equivalence relation, \mathcal{W}' is discretely homotopic to \mathcal{P} . However, this contradicts the choice of \mathcal{W} . \square

8.3 Bounding the Total Number of Segments

Having pushed the solution onto R , we further need to make the resulting weak linkage simplified, which requires to make it have low multiplicity, be U-turn-free and canonical. We first show that we can focus only on the first two properties, as being canonical can be easily derived using cycle move operations on cycles consisting of two parallel edges.

Lemma 8.8. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a weak linkage in H that is sensible and pushed onto R , whose multiplicity is at most $2n$. Then, there exists a weak linkage \mathcal{W}' that is sensible, pushed onto R , canonical, discretely homotopic to \mathcal{W} , and whose multiplicity is upper bounded by the multiplicity of \mathcal{W} .*

Proof. Let us first argue that, there is a weak linkage \mathcal{W}' that is sensible, pushed onto R , discretely homotopic to \mathcal{W} , and whose multiplicity is upper bounded by the multiplicity of \mathcal{W} such that for all edges $e_i \in E(\mathcal{W}')$ $i \geq i$. In other words, \mathcal{W}' contains only edges of positive subscript. Consider all weak linkages in H that are sensible, pushed onto R , discretely homotopic to \mathcal{W} , and whose multiplicity is upper bounded by the multiplicity of \mathcal{W} . Among these weak linkages, let \mathcal{W}' be one such that the sum of the subscripts of the edge copies in $E(\mathcal{W}')$ is maximized. Let us argue that for every $e_i \in E(\mathcal{W})$, $i \geq 1$ and for every $j > i$, $e_i \in E(\mathcal{W})$. Suppose not, and there exists an edge e_i that is used by \mathcal{W}' such that (i) $i \leq 0$ (in fact, $i = 0$ is not possible as \mathcal{W}' is pushed onto R), or (ii) there exists an edge e_j (parallel to e_i) for some $1 \leq j < i$ that is not used by $E(\mathcal{W}')$. Since \mathcal{W}' has multiplicity at most $2n$, the satisfaction of the first condition implies the satisfaction of the second one, thus there exists such an edge e_j . Let e_t be the edge (parallel to e_j and e_i) of largest j that is used by $E(\mathcal{W}')$. Moreover, let C be the cycle (which might be the boundary of a single face) that consists of two edges: e_j and e_t . By the choice of e_t , the strict interior of C does not contain any edge of $E(\mathcal{W}')$. Thus, the cycle move operation is applicable to (W, C) where W is the walk in \mathcal{W}' that uses e_t . Let \mathcal{W}^* be the result of the application of this operation. Then, the only difference between \mathcal{W}^* and \mathcal{W}' is the replacement of e_t by e_j .

Because \mathcal{W}^* is discretely homotopic to \mathcal{W}' , \mathcal{W}' is discretely homotopic to \mathcal{W} and discrete homotopy is transitive, we derive that \mathcal{W}^* is discrete. Moreover, the endpoints of the walks in \mathcal{W}' were not changed when the cycle move operation was applied. Thus, because \mathcal{W}' is sensible, so is \mathcal{W}^* . Moreover, it is clear that \mathcal{W}^* is pushed onto R and has the same multiplicity as \mathcal{W}' . However, as $j > t$, the sum of the subscripts of the edge copies in $E(\mathcal{W}^*)$ is larger than that of $E(\mathcal{W}')$, which contradicts the choice of \mathcal{W}' .

Hence, there exist weak linkages \mathcal{W}' that are sensible, pushed onto R , discretely homotopic to \mathcal{W} , whose multiplicity is upper bounded by the multiplicity of \mathcal{W} , and for all edges $e_i \in E(\mathcal{W}')$ $i \geq i$. Consider the collection of all such weak linkages, and let \mathcal{W}^* be the one maximizing

$w(E(\mathcal{W}^*)) = \sum_{e \in E(\mathcal{W}^*)} w(e)$ where Let us begin by defining a weight function $w : E(H) \rightarrow \mathbb{Z}$ on the parallel copies of edges in H as follows.

$$w(e) = \begin{cases} -2n & e \text{ is not parallel to any edge in } E(R) \\ -2n & e = e_i \text{ is parallel to an edge in } E(R) \text{ and } i \leq 0 \\ 2n - i & e = e_i \text{ is parallel to an edge in } E(R) \text{ and } i \geq 1 \end{cases}$$

We claim that \mathcal{W}^* is canonical, i.e. for every edge $e_i \in E(\mathcal{W}^*)$ the subscript $i \geq 1$, and for every parallel edge e_j where $i \leq j < i$, $e_j \in E(\mathcal{W})$. The first property is ensured by the choice of \mathcal{W}^* . For the second property, we argue as before. Suppose not, then choose i and j such that $i - j$ is minimized. Then clearly $j = i - 1$, since any parallel copy e_t with $j < t < i$ is either in $E(\mathcal{W}^*)$ contradicting the choice of i , or not in $E(\mathcal{W}^*)$ contradicting the choice if j . Therefore, the edges e_i and e_j form a cycle C such that the interior of C contains no edge of any walk in \mathcal{W}^* . Let $W \in \mathcal{W}^*$ be the walk containing e_i , and observe that the cycle move operation is applicable to (W, C) . Let \widehat{W} be the result of this operation. Then observe that $w(E(\widehat{W})) > w(E(\mathcal{W}^*))$, since $w(e_j) > w(e_i)$ and $E(\widehat{W}) \setminus \{e_j\} = E(\mathcal{W}^*) \setminus \{e_i\}$. And, \widehat{W} is discretely homotopic to \mathcal{W}' which is in turn discretely homotopic to \mathcal{W} , as before we can argue that \widehat{W} contradicts the choice of \mathcal{W}^* . Hence, \mathcal{W}^* must be canonical. \square

In case we are interested only in extremality rather than canonicity, we can use the following lemma that does not increase potential. The proof is very similar to the proof of Lemma 8.9, except that now we can “move edges in either direction”, and hence avoid creating new crossings.

Lemma 8.9. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a weak linkage in H that is sensible and pushed onto R , whose multiplicity is at most $2n$. Then, there exists a weak linkage \mathcal{W}' that is sensible, pushed onto R , extremal, discretely homotopic to \mathcal{W} , and whose potential is upper bounded by the potential of \mathcal{W} .*

Proof. Consider all weak linkages in H that are sensible, pushed onto R , discretely homotopic to \mathcal{W} , and whose potential and multiplicity are upper bounded by the potential and multiplicity, respectively, of \mathcal{W} . Among these weak linkages, let \mathcal{W}' be one such that the sum of the absolute values of the subscripts of the edge copies in $E(\mathcal{W}')$ is maximized. We claim that \mathcal{W}' is extremal, which will prove the lemma. To this end, suppose by way of contradiction that \mathcal{W}' is not extremal. Thus, there exist an edge $e_i, e_j \in \mathcal{W}'$ where $i \geq 1, j \leq -1$ and $(i - 1) + |j + 1| \leq 2n - 1$. Because the multiplicity of \mathcal{W}' is at most $2n$, this means that there exists an edge e_p (parallel to e_i and e_j) for some $p > i \geq 1$ that is not used by $E(\mathcal{W}')$. Let e_t be the edge (parallel to e_j and e_i) of largest subscript smaller than p that is used by $E(\mathcal{W}')$. Moreover, let C be the cycle (which might be the boundary of a single face) that consists of two edges: e_p and e_t . By the choice of e_t , the strict interior of C does not contain any edge of $E(\mathcal{W}')$. Thus, the cycle move operation is applicable to (W, C) where W is the walk in \mathcal{W}' that uses e_t . Let \mathcal{W}^* be the result of the application of this operation. Then, the only difference between \mathcal{W}^* and \mathcal{W}' is the replacement of e_t by e_p .

Because \mathcal{W}^* is discretely homotopic to \mathcal{W}' , \mathcal{W}' is discretely homotopic to \mathcal{W} and discrete homotopy is transitive, we derive that \mathcal{W}^* is discrete Moreover, the endpoints of the walks in \mathcal{W}' were not changed when the cycle move operation was applied. Thus, because \mathcal{W}' is sensible, so is \mathcal{W}^* . Moreover, it is clear that \mathcal{W}^* is pushed onto R , because \mathcal{W}^* and \mathcal{W}' cross R exactly in the same vertices and in the same direction (indeed, we have only replaced one edge of positive subscript by another parallel edge of positive subscript), they have the same potential. However, as $p > t \geq 1$, the sum of the absolute values of the subscripts of the edge copies in $E(\mathcal{W}^*)$ is larger than that of $E(\mathcal{W}')$, which contradicts the choice of \mathcal{W}' . \square

To achieve the properties of having low multiplicity and being U-turn-free, we perform two stages. In the first stage, that is the focus of this subsection, we make modifications that bound



Figure 22: Special U-Turns.

the total number of segments (rather than only the number of segments groups). The second stage, where we conclude the two properties, will be performed in the next subsection. The first stage in itself is partitioned into two phases as follows.

Phase I: Eliminating Special U-Turns. We eliminate some of the U-turns, but not all of them. Specifically, the elimination of some U-turns may result in too major changes in the segment groups, and hence we only deal with them after we bound the total number of segments, in which case classification into segment groups becomes immaterial. The U-turns we eliminate now are defined as follows (see Fig. 22).

Definition 8.18 (Special U-Turn). Let (G, S, T, g, k) be an instance of *Planar Disjoint Paths*, and R be a Steiner tree. Let \mathcal{W} be weak linkage that is pushed onto R . Let $U = \{e_i, e_j\}$ be a U-turn. Then, U is special if at least one among the following conditions hold: (i) i and j have the same sign, i.e. they are on the same side of R ; (ii) both endpoints of e_i and e_j do not belong to $V_{=1}(R) \cup V_{\geq 3}(R)$.

We eliminate special U-turns one-by-one, where the U-turn chosen to eliminate at each step is an innermost one, defined as follows.

Definition 8.19 (Innermost U-Turn). Let (G, S, T, g, k) be an instance of *Planar Disjoint Paths*, and R be a Steiner tree. Let \mathcal{W} be weak linkage that is pushed onto R . Let $U = \{e_i, e_j\}$ be a U-turn. Then, U is innermost if there does not exist a parallel edge $e_\ell \in E(\mathcal{W})$ such that $\min\{i, j\} \leq \ell \leq \max\{i, j\}$. We say that U is crossing if the signs of i and j are different, i.e. e_i and e_j are on opposite sides of R .

We argue that if there is a (special) U-turn, then there is also an innermost (special) one.

Lemma 8.10. Let (G, S, T, g, k) be an instance of *Planar Disjoint Paths*, and R be a Steiner tree. Let \mathcal{W} a weak linkage pushed onto R with at least one U-turn $U = \{e_i, e_j\}$. Then, \mathcal{W} has at least one innermost U-turn $U' = \{e_x, e_y\}$ whose edges lies in the interior (including the boundary) of the cycle C formed by e_i and e_j .

Proof. Denote the endpoints of e_i and e_j by u and v . Among all U-turns whose edges lies in the interior (including the boundary) of the cycle C formed by e_i and e_j (because U satisfies these conditions, there exists at least one such U-turn), let $U' = \{e_x, e_y\}$ be one whose edges e_x and e_y form a cycle C' that contains minimum number of edges of H in its interior. Let W' be the walk in \mathcal{W} that traverses e_x and e_y consecutively. Without loss of generality, suppose that when we traverse W' so that we visit e_x and then e_y , we first visit u , then v , and then u again.

We claim that U' is innermost. To this end, suppose by way of contradiction that U' is not innermost. Thus, by Definition 8.19, this means that C' contains an edge e_ℓ in its strict interior that belongs to some walk $\widehat{W} \in \mathcal{W}$ (possibly $\widehat{W} = W'$). Because U' is a U-turn, e_ℓ is neither the first nor the last edge of \widehat{W} . Thus, when we traverse \widehat{W} so that when we visit e_ℓ , we first visit u and then v , we next visit an edge e' . Because \mathcal{W} is a weak linkage, this edge must belong to the strict interior of C' (because otherwise we obtain that $(v, e_x, e_y, e_\ell,)$ is a crossing or an edge is used more than once). However, this implies that e' is parallel to the edges e_ℓ, e_x, e_y , and $\widehat{U} = \{e_\ell, e'\}$ is a U-turn whose edges lies in the interior (including the boundary) of the cycle C and which forms a cycle \widehat{C} that contains fewer edge of H than C' in its interior. This is a contradiction to the choice of U' . \square




Figure 23: U-Turn with e_i and e_j on the same side

We now prove that an innermost U-turn corresponds to a cycle on which we can perform the cycle pull operation, and consider the result of its application.

Lemma 8.11. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a sensible, outer-terminal, extremal weak linkage that is pushed onto R , and let $U = \{e_i, e_j\}$ be an innermost U-turn. Let W be the walk in \mathcal{W} that uses e_i and e_j , and C be the cycle in H that consists of e_i and e_j . Then, the cycle pull operation is applicable to (W, C) . Furthermore, the resulting weak linkage \mathcal{W}' is sensible, outer-terminal, extremal, pushed onto R , having fewer edges than \mathcal{W} , $|\text{Seg}(\mathcal{W}')| \leq |\text{Seg}(\mathcal{W})|$, and if U is special, then also its potential is upper bounded by the potential of \mathcal{W} .*

Proof. Because U is innermost, there does not exist an edge in the strict interior of C that belongs to $E(\mathcal{W})$. Therefore, the cycle pull operation is applicable to (W, C) . The only difference between \mathcal{W}' and \mathcal{W} is that \mathcal{W}' does not use the edge e_i and e_j and hence the vertex, say, v , that W visits between them. Therefore, because \mathcal{W} be a sensible, outer-terminal, extremal and pushed onto R , so is \mathcal{W}' . Moreover, the walks in \mathcal{W}' have the same endpoints as their corresponding walks in \mathcal{W} , and thus because \mathcal{W} is sensible, so is \mathcal{W}' . Let u be the other endpoints of the edges e_i and e_j , and let W' be the walk in \mathcal{W}' that resulted from W . Observe that W' has at most as many crossings with R as W has—indeed, if the elimination of e_i and e_j created a new crossing at u (this is the only new crossing that may be created), then W crosses R between e_i and e_j and this crossing does not occur in W' . Thus, $|\text{Seg}(\mathcal{W}')| \leq |\text{Seg}(\mathcal{W})|$.

Now, suppose that U is special. Then, at least one among the following conditions holds: (i) i and j have the same sign; (ii) $u, v \notin V_{=1}(R) \cup V_{\geq 3}(R)$. We first consider the case where i and j have the same sign, say positive (without loss of generality; see Fig. 23). Let $e'_x, \widehat{e}_y \in E(\mathcal{W})$ be such that e'_x and e_i are consecutive in W (if such an edge e'_x exists) and denote the segment that contains e'_x by S_x , and \widehat{e}_y and e_j are consecutive in W (if such an edge \widehat{e}_y exists) and denote the segment that contains \widehat{e}_y by S_y . Possibly some edges among e'_x, \widehat{e}_y and e_i are parallel. In case e'_x does not exist (see Fig. 23(a)), then $u \in V_{=1}(R)$ and therefore e_i and e_j belong to a segment that is in a singleton segment group. When we remove e_i and e_j , either this segment shrinks (and remains in a singleton group) or it is removed completely together with its segment group. If the segment shrinks, the potential clearly remains unchanged, and otherwise the reduction of segment groups makes the potential decrease by 1 (the potential of the consecutive segment group remains unchanged as it is a singleton segment group because it has an endpoint in $V_{=1}(R)$). The case where \widehat{e}_y does not exist is symmetric, thus we now assume that both e'_x and \widehat{e}_y exist. In case both x and y are on the same side as e_i and e_j (see Fig. 23(b)), then the removal of e_i and e_j only shrinks the segment $S_x = S_y$ where all of the four edges e_i, e_j, e'_x and \widehat{e}_y lie, and thus does not change the potential. Similarly, if e'_x is on the same side as e_i, e_j and \widehat{e}_y is on the opposite side (see Fig. 23(c)), then we only shrink the segment where e'_x, e_i and e_j lie, and rather than crossing from e_j to \widehat{e}_y , we cross from e'_x to \widehat{e}_y (which have the same label as both cross from the positive side to the negative side). The case where e'_x is on the opposite side of e_i, e_j and \widehat{e}_y is on the same side is not possible, since then W crosses itself.

Now, consider the case where both e'_x and \widehat{e}_y are both on the opposite side of e_i, e_j (see Fig. 24) Then, e_i and e_j form a complete segment, which we call S_{ij} , and the segments S_x, S_{ij} and S_y are different. Notice that the two crossings with R , one consisting of e'_x with e_i , and

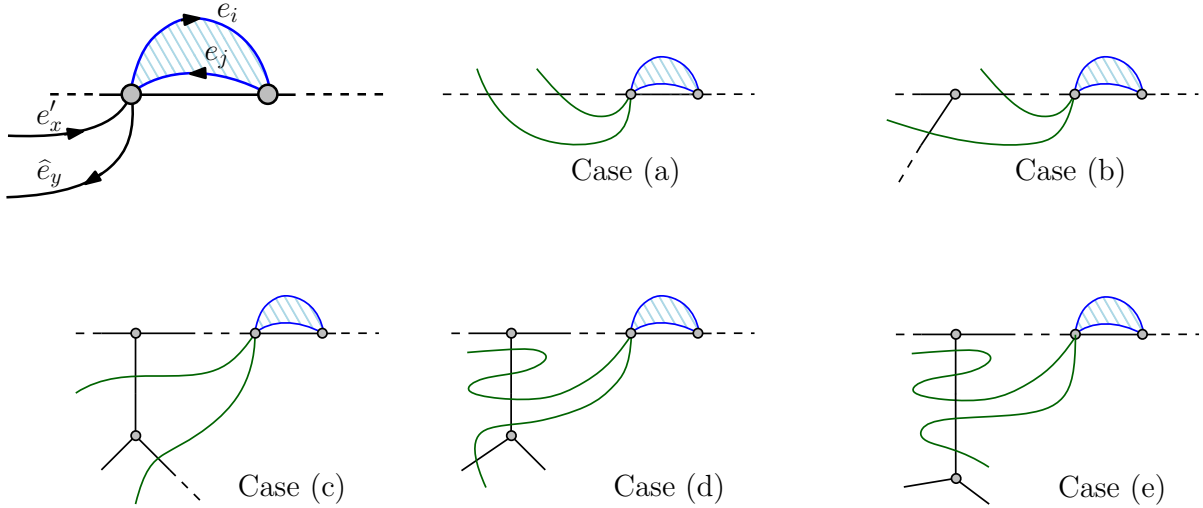


Figure 24: U-Turn with e_i, e_j on one side and e'_x, \hat{e}_y is on the opposite side.

the other consisting of e_j with \hat{e}_y , cross in different directions. If both S_x, S_{ij} and S_y belong to the same segment group (see Fig. 24(a)), then the removal of S_{ij} removes the contributions of its two crossings (mentioned above), which sum to 0 as their direction is opposite. Then, the potential remains unchanged. Now, suppose that exactly one among S_x and S_y belongs to the same group as S_{ij} . Without loss of generality, suppose that it is S_x (the other case is symmetric; see Fig. 24(b)). Then, when we remove S_{ij} , the segments S_x and S_y merge into one segment that has endpoints in different maximal degree-2 paths in R (or on vertices of degree other than 2) and hence forms its own group. This group replaces the singleton group of \mathcal{W} that contained only S_y . Furthermore, the labeling of all crossings remain the same (as well as all other associations into segment groups), apart from the two crossings consisting of e'_x with e_i , and of e_j with \hat{e}_y , which are both eliminated but have previously contributed together 0 (as they cross in opposite directions). We remark that the size of the segment group that previously contained S_x might become 1 or completely removed, but this does not increase potential. Thus, overall the potential does not increase.

Lastly, suppose that S_x, S_{ij} and S_y belong to different segment groups. Then, S_x and S_y had endpoints in different maximal degree-2 paths in R (or on vertices of degree other than 2), hence each one among S_x, S_{ij} and S_y belonged to a singleton segment group of \mathcal{W} . The removal of S_{ij} eliminates all of these three segment groups, which results in a decrease of 3 in the potential. However, now S_x and S_y belong to the same segment group. If they form a singleton segment group (see Fig. 24(c)), the overall the potential decreases by 2. Else, they join an existing segment group, and we have several subcases as follows. In the first subcase, suppose that they join only the group that contains the segment \tilde{S}_x of \mathcal{W} consecutive to S_x (in the walk in \mathcal{W} to which S_x, S_{ij} and S_y belong; see Fig. 24(d)). Then, the crossing at the endpoint of S_y is now contributing (1 or -1) to the sum of labels in the potential, and if \tilde{S}_x was in a singleton group in \mathcal{W} , then so are its crossings. Overall, this results in a contribution of at most 3, so in total the potential does not increase. The subcase where they join only the group that contains the segment \tilde{S}_y of \mathcal{W} consecutive to S_y is symmetric. Now, consider the subcase where they join both of these groups and hence merge them (see Fig. 24(e)). In this subcase, we have four new crossings that may contribute to the sum of labels, but we have also merged two groups which makes the potential decreased by at least 1, so overall the potential does not increase.

Now, suppose that only case (ii) holds. That is, $u, v \notin V_{=1}(R) \cup V_{\geq 3}(R)$, and i and j have the different signs (see Fig. 25). Without loss of generality, suppose that $i \geq 1$ and $j \leq -1$. Because $u, v \notin V_{=1}(R) \cup V_{\geq 3}(R)$ and \mathcal{W} is sensible, e'_x and \hat{e}_y exist. The case where e'_x is on the opposite side of e_i and \hat{e}_y is on the same side as e_i cannot occur since then W crosses itself.




Figure 25: U-Turn with e_i on one side and e_j on the opposite side.

Thus, we are left with three cases: (a) e'_x is on the same side as e_i and \widehat{e}_y is on the opposite side of e_i ; (b) both e'_x and \widehat{e}_y are on the opposite side of e_i ; and (c) both e'_x, \widehat{e}_y are on the same side as e_i . The cases (b) and (c) are symmetric, therefore we will only consider cases (a) and (b). In case (a), e'_x and e_i belong to one segment, and e_j and \widehat{e}_y belong to a different segment (see Fig. 25(a)). The removal of e_i and e_j only shirks these two segments by one edge each, and does not change the labeling of the crossings at their endpoints—previously, we crossed from e_i to e_j , and now we cross from e'_x to \widehat{e}_y , which are both crossings from the side of e_i to the opposite side. Thus, the potential does not increase.

Lastly, consider case (b) (see Fig. 25(b)). In this case, e_i belongs to a segment S_i containing only e_i , and e_j belongs to S_y . Further, when crossing from e_x to e_i , we cross from the opposite side of e_i to the same side, and when we cross from e_i to e_j , we cross from the side of e_i to the opposite side. Additionally, notice that the elimination of e_i and e_j results in the elimination of S_i , and the merge of S_x and S_y with e_j removed. We consider several subcases as follows. In the subcase where S_x, S_i and S_y belong to the same group (see Fig. 25(c)), then the only possible effect with respect to the potential of this group is the cancellation of the two crossings (of e'_x with e_i and of e_i with e_j), but these two crossings together contribute 0 to the sum of labels because they cross in opposite directions. Possibly the size of the segment group shrunk to 1, but this does not increase potential. Thus, in this subcase, the potential does not increase.

Now, consider the subcase where S_x and S_i are in the same segment group, and S_y is in a different segment group (see Fig. 25(d)). Then, S_y is in a singleton segment group because its endpoints belong to different maximal degree-2 path of R (or vertices of degree other than 2). In \mathcal{W}' , the segment group that resulted from the merge of S_x and S_y is also a singleton segment group. Furthermore, the segment group of \mathcal{W} that contained S_x and S_i does not change in terms of its labeled sum since the crossings at the endpoints of S_i crossed in opposite directions. Possibly the size of the segment group shrunk to 1, but this does not increase potential. Thus, in this subcase, the potential does not increase. Next, we note that the analysis of the subcase where S_i and S_y are in the same segment group, and S_x is in a different segment group, is symmetric. Lastly, suppose that S_x, S_i and S_y belong to different segment groups (see Fig. 25(e)). The analysis of this case is the same as the analysis of the last subcase of case (i) (i.e., the subcase where S_x, S_{ij} and S_y belong to the same segment group, where now we have S_i instead of S_{ij}). \square

We are now ready to assert that all special U-turns can be eliminated as follows.




Figure 26: A swollen segment and the cycles in its move-through tuple.

Lemma 8.12. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths, and R be a backbone Steiner tree. Then, there exists a sensible, outer-terminal, extremal weak linkage \mathcal{W} in H that has no special U-turns, is pushed onto R and discretely homotopic in H to some solution of (G, S, T, g, k) and $\text{Potential}(\mathcal{W}) \leq \alpha_{\text{potential}}(k)$.*

Proof. By Lemma 8.7, there exists a sensible outer-terminal weak linkage in H that is pushed onto R , has multiplicity at most $2n$, discretely homotopic in H to some solution of (G, S, T, g, k) and has potential at most $\alpha_{\text{potential}}(k)$. By Lemma 8.9 and because discrete homotopy is an equivalence relation, there also exists such a weak linkage \mathcal{W}' that is extremal. Thus, there exists a weak linkage \mathcal{W} that among all sensible, outer-terminal, extremal weak linkages in H that are pushed onto R , discretely homotopic in H to some solution of (G, S, T, g, k) and satisfy $\text{Potential}(\mathcal{W}) \leq \alpha_{\text{potential}}(k)$, the weak linkage \mathcal{W} is one that minimizes the number of edges that it uses. To conclude the proof, it suffices to argue that \mathcal{W} has no special U-turns.

Suppose, by way of contradiction, that \mathcal{W} has at least one special U-turn. Then, by Lemma 8.10, \mathcal{W} has an innermost special U-turn $U = \{e_i, e_j\}$. Let W be the walk in \mathcal{W} that uses e_i and e_j , and C be the cycle in H that consists of e_i and e_j . Then, by Lemma 8.11, the cycle pull operation is applicable to (W, C) . Furthermore, by Lemma 8.11, the resulting weak linkage \mathcal{W}' is sensible, outer-terminal, extremal, pushed onto R , has fewer edges than \mathcal{W} , and its potential is upper bounded by the potential of \mathcal{W} . Since discrete homotopy is an equivalence relation, \mathcal{W}' is discretely homotopic to some solution of (G, S, T, g, k) . However, this is a contradiction to the choice of \mathcal{W} . \square

Phase II: Eliminating Swollen Segments. The goal of the second phase is to eliminate the existence crossings with opposing “signs” for each segment and thereby, as the potential is bounded, bound the number of segments (rather than only the number of segment groups). We remark that one can show, even without this step, that the multiplicity is bounded, however this complicates the analysis. Towards this, we eliminate “swollen” segments (see Fig. 26).

Definition 8.20 (Swollen Segment). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be weak linkage that is pushed onto R . Consider segment $S \in \text{Seg}(\mathcal{W})$ for some $W \in \mathcal{W}$ such that S does not contain two the extreme edges of W . Let e and e' be the extreme edges of S (possibly $e = e'$), and let \hat{e} and \hat{e}' be the edges of $E(W) \setminus E(S)$ that are consecutive to e and e' on W , respectively. Then, S is swollen if its endpoints are internal vertices of the same maximal degree-2 path P of R and $\text{label}_P^W(e, e') \neq \text{label}_P^W(\hat{e}, \hat{e}')$.*

We show that due to the first phase, when we deal with outer-terminal weak linkages, the swollen segments have a “clean appearance” as stated in the following lemma.

Lemma 8.13. *Let (G, S, T, g, k) be a nice instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be an outer-terminal weak linkage that is pushed onto R and has no special U-turns, and $S \in \text{Seg}(\mathcal{W})$ be swollen. Then, S is parallel to a subpath of a maximal degree-2 path of R .*

Proof. Let e_i and e'_j be the first and last edges of S , and denote the endpoints of S by u and v where u is an endpoint of e_i . Because S is a swollen segment, both these edges are on the same side of $\text{path}_R(u, v)$. Let P be the unique subpath in R between u and v . Because \mathcal{W} does not have any special U-turn, we have that one of the following cases occurs (see Fig. 27): (i) S traverses a path that starts at e_i , consists of edges parallel to P and ends at e'_j ; (ii) S traverses a path that starts at e_i , consists of edges parallel to P but does not end at e'_j , and hence (to reach e'_j without having U-turns) S traverses at least two copies (on opposite sides) of every

edge of R ; (iii) the first edge that S traverses after e_i is not parallel to an edge of P , and hence (to reach e'_j without having U-turns) S traverses at least two copies (on opposite sides) of every edge of R except possibly for the edges of P . In the first case, we are done. In the other two cases, we have that $E(\mathcal{W})$ contains more than one copy of the edge incident to t^* in R , which contradicts the assumption that \mathcal{W} is outer-terminal. \square

The segment chosen to move at each step is an innermost one, formally defined as follows.

Definition 8.21 (Innermost Swollen Segment). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be weak linkage that is pushed onto R . Let $S \in \text{Seg}(\mathcal{W})$ be swollen. Then, S is innermost if there do not exist parallel edges $e_i \in E(S)$ and $e_j \in E(\mathcal{W}) \setminus E(S)$ such that i and j have the same sign and $|j| < |i|$.*

We now argue that if there is a swollen segment, then there is also an innermost one.

Lemma 8.14. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} an outer-terminal weak linkage that has no special U-turns and is pushed onto R , such that $\text{Seg}(\mathcal{W})$ contains at least one swollen segment. Then, $\text{Seg}(\mathcal{W})$ contains at least one innermost swollen segment.*

Proof. Let S be a swollen segment of \mathcal{W} such that the sum of the absolute values of the indices of the edge copies it uses is minimized. By Lemma 8.13, S is parallel to a subpath P of a maximal degree-2 path of R . Thus, because S is a segment, all the edge copies it uses are on the same side. We claim that S is innermost. Suppose, by way of contradiction, that this claim is false. Thus there exist parallel edges $e_i \in E(S)$ and $e_j \in E(\mathcal{W}) \setminus E(S)$ such that i and j have the same sign and $|j| < |i|$. Let S' be the segment of \mathcal{W}' to which e_j belongs. Because \mathcal{W} has no special U-turns and because weak linkages contain neither crossings nor repeated edges, it follows that S' is parallel to a subpath Q of P and consists only of edge copies whose indices strictly smaller absolute value than the edges of P they are parallel to. However, this implies that S' is a swollen segment of \mathcal{W} such that the the sum of the absolute value of the indices of the edge copies it uses is smaller than the sum of S . This contradicts the choice of S . \square

Given an innermost swollen segment whose copies have, on one side of R , we would like to move the segment to “the other side” of R . We know that these copies will be free in case we handle an extremal weak linkage. We now define a tuple of cycles on which we will perform move operations (see Fig. 26). The fact that this notion is well-defined (in the sense that the indices ℓ in the definition exist) will be argued in the lemma that follows it.

Definition 8.22 (Move-Through Tuple). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} an outer-terminal extremal weak linkage that has no special U-turns and is pushed onto R , and let $S \in \text{Seg}(\mathcal{W})$ be an innermost swollen segment. Let $e_{i_1}^1, e_{i_2}^2, \dots, e_{i_t}^t$, where $t = |E(S)|$, be the edges of S in the order occurred when S is traversed from one endpoint to another.¹⁸ Then, the move-through tuple of S is $T = (C_1, \dots, C_t)$ where for every $j \in \{1, \dots, t\}$, C_j is a cycle that consists of two parallel edges: $e_{i_j}^j$ and e_ℓ^j where ℓ is the index of sign opposite to i_j that has the largest absolute value such that all indices r of the same sign as ℓ and whose absolute value is upper bounded by $|\ell|$ satisfy that $e_r^j \notin E(\mathcal{W})$.*

The application of T is done by applying the cycle move operation to (W, C_i) for i from 1 to t (in this order) where W is the walk that contains S as a segment.¹⁹

Now, we prove that application of a move-through tuple is valid.

¹⁸To avoid ambiguity in the context of this definition, suppose that we have a fixed choice (e.g., lexicographic) of which endpoint is traversed first.

¹⁹Note that W changes in each application, thus by W we mean the current walk with the same endpoints as the original walk in \mathcal{W} that had S as a segment.




Figure 27: Illustration of Lemma 8.13.

Lemma 8.15. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} an outer-terminal extremal weak linkage that has no special U-turns and is pushed onto R , and let $S \in \text{Seg}(\mathcal{W})$ be an innermost swollen segment. Then, the move-through tuple $T = (C_1, \dots, C_t)$ of S is well-defined, and the application of T is valid (that is, the cycle move operation is applicable to (W, C_i) when it is done).*

Proof. Let $e_{i_1}^1, e_{i_2}^2, \dots, e_{i_t}^t$, where $t = |E(S)|$, be the edges of S in the order occurred when S is traversed from one endpoint to another. To assert that T is valid, consider some $j \in \{1, \dots, t\}$. We need to show that e_b^j , where $b \in \{-1, 1\}$ has sign opposite to the sign of i_j , does not belong to $E(\mathcal{W})$. Indeed, then there also exists an index ℓ_j of sign opposite to i_j that has the largest absolute value such that all indices r of the same sign as ℓ_j and whose absolute value is upper bounded by $|\ell_j|$ satisfy that $e_r^j \notin E(\mathcal{W})$. To this end, suppose by contradiction that $e_b^j \in E(\mathcal{W})$. Without loss of generality, suppose that $b = -1$ (the other case is symmetric). Then, we have $(i_j - 1) + |-b + 1| = i_j - 1 \leq 2n - 1$, which contradicts the assumption that \mathcal{W} is extremal. Thus, T is valid.

Notice that for every $j \in \{1, \dots, t\}$, all of the edges parallel to $e_{i_j}^j$ (and $e_{\ell_j}^j$) whose index is of the same sign as i_j and has absolute value is smaller than $|i_j|$ do not belong to $E(\mathcal{W})$ (because S is innermost). Additionally for every $j \in \{1, \dots, t\}$, all of the edges parallel to $e_{i_j}^j$ (and $e_{\ell_j}^j$) whose index is of sign opposite to i_j and has absolute value is smaller or equal to than $|\ell_j|$ do not belong to $E(\mathcal{W})$ (by the choice of ℓ_j). Thus, for each $j \in \{1, \dots, t\}$ the interior of C_j does not contain any edge of \mathcal{W} , and in particular the cycle move operation is applicable to it. When we apply the cycle move operation to some cycle C_j , it replaces $e_{i_j}^j$ by $e_{i_j}^j$. By Lemma 8.14, these replacements are done on edges not parallel to one another—that is, for every pair of distinct $j, j' \in \{1, \dots, t\}$, the edges $e_{i_j}^j$ and $e_{i_{j'}}^{j'}$ are not parallel. Thus, the application of one cycle move operation in the application of T does not effect the applicability of any other cycle move operation in the application of T . Therefore, the application of T is valid. \square

Now, we consider the properties of the weak linkage that results from the application of a move-through tuple.

Lemma 8.16. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a sensible, outer-terminal, extremal weak linkage that has no special U-turns and is pushed onto R , and let $S \in \text{Seg}(\mathcal{W})$ be an innermost swollen segment. Let T be the move-through tuple of S , and let \mathcal{W}' be the weak linkage that results from the application of T . Then, \mathcal{W}' is a sensible, outer-terminal, extremal weak linkage that has no special U-turns and is pushed onto R , whose potential is upper bounded by the potential of \mathcal{W} and which has fewer segments than \mathcal{W} .*

Proof. Let u and v be the endpoints of S . Because S is swollen, u and v are internal vertices of the same maximal degree-2 path P of R . Thus, because \mathcal{W} is sensible, there exists a segment S_1 and a segment S_2 that the walk $W \in \mathcal{W}$ that has S as a segment traverses immediately before visiting the endpoint u of S , and immediately after visiting the endpoint v of S , respectively. (Here, we supposed without loss of generality that W is traversed from one endpoint to another such that the endpoint u of S is visited before the endpoint v of S .) By Lemma 8.14, S is parallel to the subpath Q of P with endpoints u and v , and without loss of generality, we suppose that it uses the positive copies of the edges of Q . Then, the application of T replaces each one of these positive copies by a negative copy. Thus, the segments S_1, S, S_2 are removed and replaced by

one segment S^* that consists of S_1 , the new negative copies of the edges of Q , and S_2 . Hence, \mathcal{W}' has fewer (by 2) segments than \mathcal{W} . Therefore, because \mathcal{W} is sensible, outer-terminal, has no special U-turn and is pushed onto R , it is clear that \mathcal{W}' also has these properties. Now, we show that \mathcal{W}' also has the property that it is extremal. Clearly, because \mathcal{W} is extremal, the multiplicity of \mathcal{W}' is also upper bounded by $2n$, and for any two parallel edges $e_i, e_j \in E(\mathcal{W}')$ that are not parallel to an edge of Q , where $i \geq 1$ and $j \leq -1$, we have $(i - 1) + |j + 1| \geq 2n$. Now, consider some two edges $e_i, e_j \in E(\mathcal{W}')$ that are parallel to an edge e_0 of Q , where $i \geq 1$ and $j \leq -1$. Let e_t be the edge of S parallel to e_i and e_j , and let e_r be the edge with whom e_t is replaced in the application of t (thus, $e_t \in E(\mathcal{W}) \setminus E(\mathcal{W}')$ and $e_r \in E(\mathcal{W}') \setminus E(\mathcal{W})$). Without loss of generality, suppose that $t \geq 1$. Then, by the applicability of T , we have that $i \geq t + 1$ and $j \leq r$. Thus, $(i - 1) + |j + 1| \geq t + |r + 1| = (t - 1) + |(r - 1) + 1|$. Furthermore, by the definition of a move-through tuple, either $e_{r-1} \in E(\mathcal{W})$ or $r = -2n$. In the first case, because \mathcal{W} is extremal, we obtain that $(t - 1) + |(r - 1) + 1| \geq 2n$, and in the second case we obtain that $(t - 1) + |(r - 1) + 1| \geq 2n$ as well (because $t \geq 1$).

It remains to prove that the potential of \mathcal{W}' is upper bounded by the potential of \mathcal{W} . For this purpose, we consider several cases as follows. First, suppose that S_1, S and S_2 belong to the same segment group in \mathcal{W} , then the only effect on the potential that might increase it is the removal of the two crossings at the endpoints of S . However, by the definition of a swollen segment, these crossings have opposite labels, and hence their removal does not effect the potential. (Possibly the segment group that contained S_1, S and S_2 has shrunk to a singleton group with respect to \mathcal{W}' , but this does not increase the potential.) Now, suppose that only one among S_1 and S_2 is in the same segment group as S in \mathcal{W} , and without loss of generality, suppose that it is S_1 . Then, in \mathcal{W}' the segment S^* belongs to a singleton group (as its endpoints belong to different maximal degree-2 paths of R or it has an endpoint in $V_{=1}(R) \cup V_{\geq 3}(R)$). Moreover, in \mathcal{W} the segment S_2 belongs to a singleton segment group. Thus, one singleton segment group has been replaced by another, and the size of existing segment groups might have shrunk. Since the only change in terms of crossing is that the crossing at the endpoints of S were eliminated, and as in the previous case, this does not effect the potential, we conclude that the potential does not increase.

Lastly, we consider the case where S_1, S and S_2 belong to different segment groups in $\text{SegGro}(\mathcal{W})$. These three groups are singleton groups — S_1 and S_2 have endpoints that belong to different maximal degree-2 paths of R (or an endpoint in $V_{=1}(R) \cup V_{\geq 3}(R)$), and S lies in between them. In \mathcal{W}' , these three groups are eliminated, which results in a decrease of 3 in the potential. If S^* forms a singleton segment group, the overall the potential decreases by 2. Else, S^* joins an existing segment group, and we have several subcases as follows. In the first subcase, suppose that S^* joins only the group that contains the segment \tilde{S}_1 of \mathcal{W} consecutive to S_1 in \mathcal{W} (that is not S). Then, the crossing at the endpoint of S_2 is now contributing (1 or -1) to the sum of labels in the potential of the group, and if \tilde{S}_1 was in a singleton group in \mathcal{W} , then so are its two crossings. Overall, this results in a contribution of at most 3, so in total the potential does not increase. The subcase where S^* joins only the group that contains the segment \tilde{S}_2 of \mathcal{W} consecutive to S_2 is symmetric. Now, consider the subcase where S^* joins both of these groups and hence we merge them with respect to \mathcal{W}' . In this subcase, we have four new crossings that may contribute to the sum of labels, but we have also merged two groups which makes the potential decreased by at least 1, so overall the potential does not increase. \square

Lastly, we assert that all swollen segments can be eliminated as follows.

Lemma 8.17. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths, and R be a backbone Steiner tree. Then, there exists a sensible, outer-terminal, extremal weak linkage \mathcal{W} in H that is pushed onto R , has no special U-turns and swollen segments, is discretely homotopic in H to some solution of (G, S, T, g, k) and $\text{Potential}(\mathcal{W}) \leq \alpha_{\text{potential}}(k)$.*

Proof. By Lemma 8.12, there exists a sensible, outer-terminal, extremal weak linkage in H

that has no special U-turns, is pushed onto R , discretely homotopic in H to some solution of (G, S, T, g, k) and whose potential is upper bounded by $\alpha_{\text{potential}}(k)$. Among all such weak linkages, let \mathcal{W} be one with minimum number of segments. To conclude the proof, it suffices to argue that \mathcal{W} has no swollen segments.

Suppose, by way of contradiction, that \mathcal{W} has at least one swollen segment. Then, by Lemma 8.14, \mathcal{W} has an innermost swollen segment S . By Lemma 8.15, the move-through tuple T of S is well-defined, and its application is valid. Furthermore, by Lemma 8.16, the resulting weak linkage \mathcal{W}' is sensible, outer-terminal, extremal, pushed onto R , has no special U-turns and fewer segments than \mathcal{W} , and its potential is upper bounded by the potential of \mathcal{W} . Since discrete homotopy is an equivalence relation, \mathcal{W}' is discretely homotopic to some solution of (G, S, T, g, k) . However, this is a contradiction to the choice of \mathcal{W} . \square

Lastly, we prove that having eliminated all swollen segments indeed implies that the total number of segments is small.

Lemma 8.18. *Let (G, S, T, g, k) be a nice instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be a sensible, outer-terminal, extremal weak linkage that is pushed onto R and has no special U-turns and swollen segments. Then, $|\text{Seg}(\mathcal{W})| \leq \text{Potential}(\mathcal{W})$.*

Proof. To prove that $|\text{Seg}(\mathcal{W})| \leq \text{Potential}(\mathcal{W})$, it suffices to show that for every segment group $W \in \text{SegGro}(\mathcal{W})$, we have that $|\text{Seg}(W)| \leq \text{Potential}(W)$. For segment groups of size 1, this inequality is immediate. Thus, we now consider a segment group $W \in \text{SegGro}(\mathcal{W})$ of size at least 2. Then,

$$\text{Potential}(W) = 1 + \left| \sum_{(e, e') \in E(W) \times E(W)} \text{label}_P^W(e, e') \right|,$$

where P is the maximal degree-2 path of R such that all of the endpoints of all of the segments in $\text{Seg}(W)$ are its internal vertices. Because there do not exist swollen segments, we have that label_P^W assigns either only non-negative values (0 or 1) or only non-positive values (0 or -1). Without loss of generality, suppose that it assigns only non-negative values. Now, notice that every pair of edges consecutively visited by W that belong to different segments of W is assigned 1 (because it creates a crossing with P). However, the number of segments of \mathcal{W} is upper bounded by one plus the number of such pairs of edges. Thus, $|\text{Seg}(W)| - 1 \leq \sum_{(e, e') \in E(W) \times E(W)} \text{label}_P^W(e, e')$. From this, we conclude that $|\text{Seg}(W)| \leq \text{Potential}(W)$. \square

8.4 Completion of the Simplification

The purpose of this section is to prove the following lemma.

Lemma 8.19. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths, and R be a backbone Steiner tree. Then, there exists a simplified weak linkage \mathcal{W} in H that is discretely homotopic in H to some solution of (G, S, T, g, k) .*

To this end, we first eliminate all remaining (non-special) U-turns based on Lemmas 8.11, 8.17 and 8.18, similarly to the proof of Lemma 8.12.

Lemma 8.20. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths, and R be a backbone Steiner tree. Then, there exists a sensible, outer-terminal, U-turn-free weak linkage \mathcal{W} in H that is pushed onto R , discretely homotopic in H to some solution of (G, S, T, g, k) , and $|\text{Seg}(\mathcal{W})| \leq \alpha_{\text{potential}}(k)$.*

Proof. By Lemma 8.17, there exists a sensible, outer-terminal, extremal weak linkage in H that is pushed onto R , has no special U-turns and swollen segments, is discretely homotopic in H to some solution of (G, S, T, g, k) and whose potential is upper bounded by $\alpha_{\text{potential}}(k)$. By Lemma 8.18, its number of segments is also upper bounded $\alpha_{\text{potential}}(k)$. Among all such weak

linkages that are sensible, outer-terminal, pushed onto R and whose number of segments is upper bounded $\alpha_{\text{potential}}(k)$, let \mathcal{W} be one with minimum number of edges. To conclude the proof, it suffices to argue that \mathcal{W} is U-turn-free.

Suppose, by way of contradiction, that \mathcal{W} has at least one U-turn. Then, by Lemma 8.10, \mathcal{W} has an innermost U-turn $U = \{e_i, e_j\}$. Let W be the walk in \mathcal{W} that uses e_i and e_j , and C be the cycle in H that consists of e_i and e_j . Then, by Lemma 8.11, the cycle pull operation is applicable to (W, C) . Furthermore, by Lemma 8.11, the resulting weak linkage \mathcal{W}' is sensible, outer-terminal, pushed onto R , has fewer edges than \mathcal{W} , and its number of segments is upper bounded by the number of segments of \mathcal{W} . Since discrete homotopy is an equivalence relation, \mathcal{W}' is discretely homotopic to some solution of (G, S, T, g, k) . However, this is a contradiction to the choice of \mathcal{W} . \square

Now, we prove that having no U-turns implies that each segment can use only two parallel copies of every edge.

Lemma 8.21. *Let (G, S, T, g, k) be a nice instance of Planar Disjoint Paths, and R be a Steiner tree. Let \mathcal{W} be an outer-terminal, U-turn-free weak linkage that is pushed onto R . Then, each segment $S \in \text{Seg}(\mathcal{W})$ uses at most two copies of every edge in $E(R)$.*

Proof. Consider some segment $S \in \text{Seg}(\mathcal{W})$. Suppose, by way of contradiction, that there exists some edge $e_0 \in E(R)$ such that S contains at least three edges parallel to e_0 (but it cannot contain e_0 as it is pushed onto R). Then, without loss of generality, suppose that it contains two copies with positive subscript, e_i and e_j , and let S' be the subwalk of S having these edge copies as the extreme edges. Then, because S is U-turn free, when we traverse S' from e_i to e_j , then we must visit the positive and negative copy of every other edge in $E(R)$ exactly once. However, this means that $E(\mathcal{W})$ contains more than one copy of the edge incident to t^* in R , which contradicts the assumption that \mathcal{W} is outer-terminal. \square

Having established Lemmas 8.8, 8.20 and 8.21, we are ready to prove Lemma 8.19.

Proof of Lemma 8.19. By Lemma 8.20, there exists a sensible, outer-terminal, U-turn-free weak linkage \mathcal{W}' in H that is pushed onto R , discretely homotopic in H to some solution of (G, S, T, g, k) , and whose number of segments is upper bounded by $\alpha_{\text{potential}}(k)$. By Lemma 8.21, the multiplicity of \mathcal{W}' is upper bounded by $2\alpha_{\text{potential}}(k) = \alpha_{\text{mul}}(k)$. By Lemma 8.8, there exists a weak linkage \mathcal{W} that is sensible, pushed onto R , canonical, discretely homotopic to \mathcal{W}' , and whose multiplicity is upper bounded by the multiplicity of \mathcal{W}' . Thus, \mathcal{W} is simplified. Moreover, since discrete homotopy is an equivalence relation, \mathcal{W} is discretely homotopic to some solution of (G, S, T, g, k) . \square

9 Reconstruction of Pushed Weak Linkages from Templates

In this section, based on the guarantee of Lemma 8.19, we only attempt to reconstruct simplified weak linkages. Towards this, we introduce the notion of a template (based on another notion called a pairing). Roughly speaking, a template indicates how many parallel copies of each edge incident to a vertex in $V_{=1}(R) \cup V_{\geq 3}(R)$ are used by the walks in the simplified weak linkage \mathcal{W} under consideration, and how many times, for each pair (e, e') of non-parallel edges sharing a vertex, the walks in \mathcal{W} traverse from a copy of e to a copy of e' . Observe that a template does not indicate which edge copy is used by each walk, but only specifies certain numbers. Nevertheless, we will show that this is sufficient for faithful reconstruction of simplified weak linkages. The great advantage of templates, proved later, is that there are only few of them.

9.1 Generic Templates and Templates of Simplified Weak Linkages

We begin with the definition of the notion of a pairing, which will form the basis of a template. Let $V^*(R) = V_{=1}(R) \cup V_{\geq 3}(R) \cup V_2^*(R)$ where $V_2^*(R) = \{v \in V_{=2}(R) \mid \exists u \in V_{=1}(R) \cup V_{\geq 3}(R) \text{ such that } \{u, v\} \in E(R)\}$. Observe that $|V_2^*(R)| \leq 2(|V_{=1}(R)| + |V_{\geq 3}(R)| - 1) \leq 8k$, by Observation 6.1. Therefore, $|V^*(R)| \leq 12k$. Let $E^*(R)$ denote the set of edges in $E(R)$ that are incident on a vertex of $V^*(R)$, and observe that $|E^*(R)| \leq 24k$ (since R is a tree).

Definition 9.1 (Pairing). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths with a Steiner tree R . For a vertex $v \in V_{\geq 3}(R)$, a pairing at v is a set pairing_v of unordered pairs of distinct edges in R incident to v . For a vertex $v \in V_2^*(R)$, a pairing at v is a collection of pairs (possibly non-distinct) edges in $E_R(v)$. And, for a vertex $v \in V_{=1}(R)$, it is the empty set or singleton set of the pair where the (unique) edge incident to v in R occurs twice. A collection $\{\text{pairing}_u\}_{u \in V^*(R)}$, where pairing_u is a pairing at u for every vertex $u \in V^*(R)$, is called a pairing.*

As we will see later, simplified weak linkages can only give rise to a specific type of pairings, which we call non-crossing pairings.

Definition 9.2 (Non-Crossing Pairing). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths. Let R be a Steiner tree. Consider a vertex $v \in V^*(R)$, and let e^1, e^2, \dots, e^r be the edges in $E(R)$ incident to v in clockwise order where the first edge e^1 is chosen arbitrarily. A pairing pairing_v at v is non-crossing if there do not exist two pairs (e^i, e^j) and (e^x, e^y) in pairing_v , where $i < j$ and $x < y$, such that $i < x < j < y$ or $x < i < y < j$. More generally, a pairing $\{\text{pairing}_u\}_{u \in V^*(R)}$ is non-crossing if, for every $u \in V^*(R)$, the pairing pairing_u is non-crossing.*

We now show that a non-crossing pairing can contain only $\mathcal{O}(k)$ pairs, which is better than a trivial bound of $\mathcal{O}(k^2)$. This bound will be required to attain a running time of $2^{\mathcal{O}(k^2)} n^{\mathcal{O}(1)}$ rather than $2^{\mathcal{O}(k^3)} n^{\mathcal{O}(1)}$.

Lemma 9.1. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths. Let R be a Steiner tree. Let $\{\text{pairing}_v\}_{v \in V^*(R)}$ be a non-crossing pairing. Then, $|\bigcup_{v \in V^*(R)} \text{pairing}_v| \leq \alpha_{\text{npair}}(k) := 48k$.*

Proof. Towards the bound on $|\bigcup_{v \in V^*(R)} \text{pairing}_v|$, we first obtain a bound on each individual set pairing_v . To this end, consider some vertex $v \in V_{\geq 3}(R)$, and let e^0, e^1, \dots, e^{r-1} be the edges in $E(R)$ incident to v in clockwise order. Consider the undirected graph C on vertex set $\{u_{e^i} \mid i \in \{0, 1, \dots, r-1\}\}$ and edge set $\{\{u_{e^i}, u_{e^{(i+1) \bmod r}}\} \mid i \in \{0, 1, \dots, r-1\}\} \cup \{\{u_{e^i}, u_{e^j}\} \mid (e^i, e^j) \in \text{pairing}_v\}$. Now, notice that C is an outerplanar graph (Fig. 28). To see this, draw the vertices of C on a circle on the plane, so that the curves on the cycle that connect them correspond to the drawing of the edges in $\{\{u_{e^i}, u_{e^{(i+1) \bmod r}}\} \mid i \in \{0, 1, \dots, r-1\}\}$. Now, for each edge in $\{\{u_{e^i}, u_{e^j}\} \mid (e^i, e^j) \in \text{pairing}_v\}$, draw a straight line segment inside the circle that connects u_{e^i} and u_{e^j} . The condition that asserts that pairing_v is non-crossing ensures that no two lines segments among those drawn previously intersect (except for at their endpoints). As an outerplanar graph on q vertices can have at most $2q - 3$ edges, we have that $|E(C)| < 2|V(C)| = 2r$. Because $|\text{pairing}_v| \leq |E(C)|$, we have that $|\text{pairing}_v| \leq 2r$. For a vertex in $v \in V_2^*(R)$, since it has only two edges incident on it, $|\text{pairing}_v| \leq 3$. Finally, for $v \in V_{=1}(R)$, $|\text{pairing}_v| \leq 1$.

Thus, for every vertex $v \in V^*(R)$, $|\text{pairing}_v|$ is bounded by twice the degree of v in R . Since $|V^*(R)| \leq 12k$, the sum of the degrees in R of the vertices in $V^*(R)$ is upper bounded by $24k$. From this, we conclude that $|\bigcup_{v \in V^*(R)} \text{pairing}_v| \leq \alpha_{\text{npair}}(k)$. \square

Now, we associate a pairing with a pushed weak linkage.

Definition 9.3 (Pairing of a Weak Linkage). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths. Let R be a Steiner tree, and let \mathcal{W} be a weak linkage pushed onto R . For a*




Figure 28: Illustration of the outerplanar graph in Lemma 9.1.

vertex $v \in V^*(R)$, the pairing of \mathcal{W} at v is the set that contains every pair of edges (e, e') in R that are incident to v and such that there exists at least one walk in \mathcal{W} where e_i and e'_j occur consecutively, where e_i and e'_j are parallel copies of e and e' , respectively. More generally, the pairing of \mathcal{W} is the collection $\{\text{pairing}_u\}_{u \in V^*(R)}$, where pairing_u is the pairing of \mathcal{W} at u for every vertex $u \in V^*(R)$.

Apart from a pairing, to be able to reconstruct a simplified weak linkage, we need additional information in the form of an assignment of numbers to the pairs in the pairing. To this end, we have the definition of a template.

Definition 9.4 (Template). Let (G, S, T, g, k) be an instance of Planar Disjoint Paths. Let R be a Steiner tree. Let pairing_v be a pairing at some vertex $v \in V^*(R)$. A template for pairing_v is a function $\text{template}_v : \text{pairing}_v \rightarrow \mathbb{N}$. If the maximum integer assigned by template_v is upper bounded by N , for some $N \in \mathbb{N}$, then it is also called an N -template. More generally, a template (resp. N -template) of a pairing $\{\text{pairing}_u\}_{u \in V^*(R)}$ is a collection $\{\text{template}_u\}_{u \in V^*(R)}$, where template_u is a template (resp. N -template) for pairing_u for every vertex $u \in V^*(R)$.

We proceed to associate a template with a weak linkage.

Definition 9.5 (Template of Weak Linkage). Let (G, S, T, g, k) be an instance of Planar Disjoint Paths. Let R be a Steiner tree, and let \mathcal{W} be a weak linkage pushed onto R . Let $\{\text{pairing}_u\}_{u \in V^*(R)}$ be the pairing of \mathcal{W} . For a vertex $v \in V^*(R)$, the template of \mathcal{W} at v is the function $\text{template}_v : \text{pairing}_v \rightarrow \mathbb{N}$ such that for every $(e, e') \in \text{pairing}_v$, we have

$$\text{template}_v((e, e')) = |\{ \{\hat{e}, \tilde{e}'\} \mid \hat{e} \text{ is parallel to } e, \tilde{e}' \text{ is parallel to } e', \\ \exists W \in \mathcal{W} \text{ s.t. } W \text{ traverses } \hat{e} \text{ and } \tilde{e}' \text{ consecutively} \}|.$$

More generally, the template of \mathcal{W} is the collection $\{\text{template}_u\}_{u \in V^*(R)}$, where template_u is the template of \mathcal{W} at u for every vertex $u \in V^*(R)$.

Now, we claim that the pairing of a pushed weak linkage non-crossing.

Lemma 9.2. Let (G, S, T, g, k) be an instance of Planar Disjoint Paths. Let R be a Steiner tree. Then, the pairing of any weak linkage \mathcal{W} pushed onto R is non-crossing.

Proof. Let $\{\text{pairing}_v\}_{v \in V^*(R)}$ be the pairing of \mathcal{W} . Suppose, by way of contradiction, that $\{\text{pairing}_v\}_{v \in V^*(R)}$ is crossing. If $v \in V_{=1}(R) \cup V_2^*(R)$, then pairing_v is trivially non-crossing. Thus, there exists a vertex $v \in V_{\geq 3}(R)$ such that pairing_v is crossing. Let e^1, e^2, \dots, e^r be the edges in $E(R)$ incident to v in clockwise order. Because pairing_v is crossing, there exist two pairs (e^i, e^j) and (e^x, e^y) in pairing_v , where $i < j$ and $x < y$, such that $i < x < j < y$ or $x < i < y < j$. By the definition of pairing_v , this means that there exist walks $\widehat{W}, \overline{W} \in \mathcal{W}$ (possibly $\widehat{W} = \overline{W}$) and edges $\hat{e}^i, \hat{e}^j, \bar{e}^x$ and \bar{e}^y that are parallel to e^i, e^j, e^x and e^y , respectively, such that \widehat{W} traverses \hat{e}^i and \hat{e}^j consecutively, and \overline{W} traverses \bar{e}^x and \bar{e}^y consecutively. However, because $i < x < j < y$ or $x < i < y < j$, and because parallel edges incident to each vertex appear consecutively in its cyclic order, we derive that $(v, \hat{e}^i, \hat{e}^j, \bar{e}^x, \bar{e}^y)$ is a crossing of \widehat{W} and \overline{W} . \square

From Lemmas 9.1 and 9.2, we obtain the following corollary.

Corollary 9.1. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths with a backbone Steiner tree R . Let \mathcal{W} be a weak linkage pushed onto R with a pairing $\{\text{pairing}_v\}_{v \in V^*(R)}$. Then, $|\bigcup_{v \in V^*(R)} \text{pairing}_v| \leq \alpha_{\text{npair}}(k)$.*

Additionally, we claim that we can focus our attention on pushed flows whose templates are $\alpha_{\text{weight}}(k)$ -templates.

Lemma 9.3. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths with a backbone Steiner tree R . Then, there exists a simplified weak linkage that is discretely homotopic in H to some solution of (G, S, T, g, k) , whose template is an $\alpha_{\text{mul}}(k)$ -template.*

Proof. By Lemma 8.19, there exists a simplified weak linkage \mathcal{W} that is discretely homotopic in H to some solution of (G, S, T, g, k) . Because \mathcal{W} is simplified, its multiplicity is upper bounded by $\alpha_{\text{mul}}(k)$. Let $\{\text{pairing}_v\}_{v \in V^*(R)}$ and $\{\text{template}_v\}_{v \in V^*(R)}$ be the pairing and template of \mathcal{W} , respectively. Consider some vertex $v \in V^*(R)$ and pair $(e, e') \in \text{pairing}_v$. To complete the proof, we need to show that $\text{template}_v((e, e')) \leq \alpha_{\text{mul}}(k)$. By the definition of a template,

$$\text{template}_v((e, e')) = |\{\{\hat{e}, \hat{e}'\} \mid \hat{e} \text{ is parallel to } e, \hat{e}' \text{ is parallel to } e', \\ \exists W \in \mathcal{W} \text{ s.t. } W \text{ traverses } \hat{e} \text{ and } \hat{e}' \text{ consecutively}\}|.$$

Thus, because every the walks in \mathcal{W} are edge-disjoint and each walk in \mathcal{W} visits distinct edges (by the definition of a weak linkage), $\text{template}_v((e, e'))$ is upper bounded by the number of edges parallel to e that belong to $E(\mathcal{W})$. Thus, by the definition of the multiplicity of a linkage, we conclude that $\text{template}_v((e, e')) \leq \alpha_{\text{mul}}(k)$. \square

In light of Corollary 9.1 and Lemma 9.3, we define the set of all pairings and templates in which we will be interested as follows.

Definition 9.6 (The Set ALL). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths. Let R be a Steiner tree. The set ALL contains every collection $\{(\text{pairing}_v, \text{template}_v)\}_{v \in V^*(R)}$ where $\{\text{pairing}_v\}_{v \in V^*(R)}$ is a non-crossing pairing that satisfies $|\bigcup_{v \in V^*(R)} \text{pairing}_v| \leq \alpha_{\text{npair}}(k)$, and $\{\text{template}_v\}_{v \in V^*(R)}$ is an $\alpha_{\text{mul}}(k)$ -template for $\{\text{pairing}_v\}_{v \in V^*(R)}$.*

From Corollary 9.1 and Lemma 9.3, we have the following result.

Corollary 9.2. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths. Let R be a backbone Steiner tree. Then, there exists a simplified weak linkage that is discretely homotopic in H to some solution of (G, S, T, g, k) and satisfies the following property: There exists $\{(\text{pairing}_v, \text{template}_v)\}_{v \in V^*(R)} \in \text{ALL}$ such that $\{\text{pairing}_v\}_{v \in V^*(R)}$ is the pairing of \mathcal{W} , and $\{\text{template}_v\}_{v \in V^*(R)}$ is the template of \mathcal{W} .*

Because we only deal with pairings having just $\mathcal{O}(k)$ pairs and upper bound the largest integer assigned by templates by $2^{\mathcal{O}(k)}$, the set ALL is “small” as asserted by the following lemma.

Lemma 9.4. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths. Let R be a Steiner tree. Then, $|\text{ALL}| = 2^{\mathcal{O}(k^2)}$. Moreover, ALL can be computed in time $2^{\mathcal{O}(k^2)}$.*

Proof. First, we upper bound the number of pairings $\{\text{pairing}_v\}_{v \in V^*(R)}$ that satisfy $|\bigcup_{v \in V^*(R)} \text{pairing}_v| \leq \alpha_{\text{npair}}(k)$. By Observation 6.1, the number of edges in $E^*(R)$ is at most $24k$, and hence the number of pairs of edges in $E^*(R)$ is at most $(24k)^2$. Thus, the number of choices for $\bigcup_{v \in V^*(R)} \text{pairing}_v$ is at most $\binom{(24k)^2}{\alpha_{\text{npair}}(k)}$. Note that each pair of edges in this union belongs to pairing_v for at most one vertex $v \in V^*(R)$. From this, we conclude that the number of pairings $\{\text{pairing}_v\}_{v \in V^*(R)}$ that satisfy $|\bigcup_{v \in V^*(R)} \text{pairing}_v| \leq \alpha_{\text{npair}}(k)$ is at most $\binom{(24k)^2}{\alpha_{\text{npair}}(k)} \cdot 2^{\alpha_{\text{npair}}(k)} = 2^{\mathcal{O}(k \log k)}$ (as $\alpha_{\text{npair}}(k) = \mathcal{O}(k)$).

Now, fix some pairing $\{\text{pairing}_v\}_{v \in V^*(R)}$ that satisfies $|\bigcup_{v \in V^*(R)} \text{pairing}_v| \leq \alpha_{\text{npair}}(k)$. We test if this pairing is a non-crossing pairing, at each vertex $v \in V^*(R)$, by testing all possible 4-tuples of edges in $E_R(v)$. This takes $k^{\mathcal{O}(1)}$ time in total. Then, the number of $\alpha_{\text{mul}}(k)$ -templates for $\{\text{pairing}_v\}_{v \in V^*(R)}$ is upper bounded by $(\alpha_{\text{mul}}(k))^{\alpha_{\text{npair}}(k)} = 2^{\mathcal{O}(k^2)}$ (as $\alpha_{\text{mul}}(k) = 2^{\mathcal{O}(k)}$ and $\alpha_{\text{npair}}(k) = \mathcal{O}(k)$). Thus, we have that $|\text{ALL}| = 2^{\mathcal{O}(k^2)}$. It should also be clear that these arguments, by simple enumeration, imply that ALL can be computed in time $2^{\mathcal{O}(k^2)}$. \square

Extension of Pairings and Templates. To describe the reconstruction of simplified weak linkages from their pairings and templates, we must extend them from $V^*(R)$ to all of $V(R)$. Intuitively, this extension is based on the observation that \mathcal{W} is U-turn free and sensible. Therefore, if a walk in \mathcal{W} , visits a maximal degree-2 path in R , then it must traverse the entirety of this path. Hence, the pairings and templates at any internal vertex of a degree-2 path can be directly obtained from the endpoint vertices of the path. We begin by identifying which collections of pairings and templates can be extended.

Definition 9.7. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and let R be a backbone Steiner tree. And let $\mathcal{A} = \{(\text{pairing}_v, \text{template}_v)\}_{v \in V^*(R)}$. Then \mathcal{A} is extensible to all of $V(R)$ if the following conditions are true for every maximal degree-2 path in R .*

- *Let $u, v \in V_2^*(R)$ such that they lie on the same maximal degree-2 path of R . Consider the subpath $\text{path}_R(u, v)$ with endpoints u and v , and let e^u and e^v be the edges in $E(\text{path}_R(u, v))$ incident on u and v , respectively. Suppose that e^u and e^v are distinct edges. Then $(e, e^u) \in \text{pairing}_u$ if and only if $(e', e^v) \in \text{pairing}_v$, where $E_R(u) = \{e, e^u\}$ and $E_R(v) = \{e', e^v\}$.*
- *Assuming that the above condition is true, furthermore $\text{template}_u(e, e^u) = \text{template}_v(e', e^v)$.*

The following lemma shows that pairings and templates of simplified weak linkages are extensible to all of $V(R)$.

Lemma 9.5. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R be a backbone Steiner tree. Let \mathcal{W} be a simplified weak linkage with pairing and template $\mathcal{A} = \{(\text{pairing}_v, \text{template}_v)\}_{v \in V^*(R)}$. Then, \mathcal{A} is extensible to all of $V(R)$.*

Proof. Let $u, v \in V_2^*(R)$ such that they lie on the same maximal degree-2 path of R . Consider the subpath $\text{path}_R(u, v)$ with endpoints u and v , and let e^u and e^v be the edges in $E(\text{path}_R(u, v))$ incident on u and v , respectively. Suppose that e^u and e^v are distinct edges. Then $\text{path}_R(u, v)$ must have an internal vertex. Let $V(\text{path}_R(u, v)) = \{u = w_0, w_1, w_2, \dots, w_p, w_{p+1} = v\}$ and let $E(\text{path}_R(u, v)) = \{e^i = \{w_i, w_{i+1}\} \mid 0 \leq i \leq p\}$ where $e^0 = e^u$ and $e^p = e^v$. Consider an internal vertex w_i of $\text{path}_R(u, v)$ and note that $E_R(w_i) = \{e^{i-1}, e^i\}$. Observe that, as \mathcal{W} is U-turn free, $w^i \notin N_R(V_{=1}(R))$ and the endpoints of every the walk in \mathcal{W} lies in $V_{=1}(R)$, there is no walk in \mathcal{W} that visits two parallel copies of an edge in $E_R(w_i)$ consecutively, as that would constitute a U-turn. Therefore, any walk in \mathcal{W} that visits $e_{j_{i-1}}^{i-1}$ must also visit $e_{j_i}^i$, where $e_{j_{i-1}}^{i-1}$ and $e_{j_i}^i$ are parallel copies of e^{i-1} and e^i respectively. This holds for all vertices w_1, w_2, \dots, w_p . Let $E^*(\text{path}_R(u, v))$ be the collection of all the parallel copies of every edge in $E(\text{path}_R(u, v))$. Then, for a walk $W \in \mathcal{W}$, let $\widehat{W}_1, \widehat{W}_2, \dots, \widehat{W}_t$ be the maximal subwalks of W restricted to $E^*(\text{path}_R(u, v))$. Then each \widehat{W}_i is a path from u to v along the parallel copies of e^0, e^1, \dots, e_p . Hence, if $(e, e^u) = (e, e_0) \in \text{pairing}_u$ if and only if $(e', e^v) = (e', e_p) \in \text{pairing}_v$ where e and e' are the edges in $E(R) \setminus E(\text{path}_R(u, v))$ that are incident on u and v , respectively. Further, observe that, each time a walk $W \in \mathcal{W}$ visits a parallel copy of $e^u = e^0$ (immediately after visiting a parallel copy of e) it must traverse a parallel copy of $e^1, e^2, \dots, e^p = e^v$ consecutively (and then immediately visit a parallel copy of e'), and vice-versa. Therefore by definition, $\text{template}_u(e, e^u) = \text{template}_{w_1}(e^u, e^1) = \text{template}_{w_2}(e^1, e^2) = \dots = \text{template}_v(e_{p-1}, e^v) = \text{template}_v(e', e^v)$. \square

We now have the following definition.

Definition 9.8 (Extension of Pairings and Templates). *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths. Let R be a backbone Steiner tree. Let $\mathcal{A} = \{(\text{pairing}_v, \text{template}_v)\}_{v \in V^*(R)}$. Then, the extension of \mathcal{A} to $V(R)$ is the collection $\widehat{\mathcal{A}} = \{(\widehat{\text{pairing}}_v, \widehat{\text{template}}_v)\}_{v \in V(R)}$ such that:*

- *If \mathcal{A} is not extensible (Definition 9.7), then $\widehat{\mathcal{A}}$ is invalid.*
- *Otherwise, \mathcal{A} is extensible and we have two cases.*
 - *If $v \in V_{=1}(R) \cup V_{\geq 3}(R)$, then $\widehat{\text{pairing}}_v = \text{pairing}_v$ and $\widehat{\text{template}}_v = \text{template}_v$.*
 - *Otherwise, $v \in V_{=2}(R)$ and let $u, w \in V_{=1}(R) \cup V_{\geq 3}(R)$ such that $v \in V(\text{path}_R(u, w))$. Let e_u and e_w be the two edges in $\text{path}_R(u, w)$ incident on u and w , respectively. If $(e, e_u) \notin \text{pairing}_u$ for any $e \in E_R(u)$, then $\text{pairing}_v = \emptyset$. Otherwise, there is some $e \in E_R(u)$ such that $(e, e_u) \in \text{pairing}_u$. Then $\text{pairing}_v = \{(e', e'')\}$ where e' and e'' are the two edges in $E(R)$ incident on v . Further, $\text{template}_v(e', e'') = \text{template}_u(e, e_u)$.*

Let $\widehat{\text{ALL}}$ denote the collection of extensions of all the pairings and templates in ALL . Then we have the following corollary of Lemma 9.4.

Lemma 9.6. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths. Let R be a Steiner tree. Then, $|\widehat{\text{ALL}}| = 2^{\mathcal{O}(k^2)}$. Moreover, $\widehat{\text{ALL}}$ can be computed in time $2^{\mathcal{O}(k^2)}n$.*

Proof. Given $\mathcal{A} = \{(\text{pairing}_v, \text{template}_v)\}_{v \in V^*(R)} \in \text{ALL}$, we apply Definition 9.8 to obtain the extension $\widehat{\mathcal{A}}$. Note that, for every vertex $v \in V^*(R)$, we have that $|\text{pairing}_v| = \mathcal{O}(k)$ and the bitsize of the numbers assigned by template_v are bounded by $2^{\mathcal{O}(k)}$. Further, since $|V^*(R)| \leq 12k$, we can test the conditions in Definition 9.8 in $k^{\mathcal{O}(1)}$ time. Finally we can construct $\widehat{\mathcal{A}} = \{(\widehat{\text{pairing}}_v, \widehat{\text{template}}_v)\}_{v \in V(R)}$ in total time $2^{\mathcal{O}(k)}n$, as $|V(R)| \leq n$. Since $|\text{ALL}| = 2^{\mathcal{O}(k^2)}$ and it can be enumerated in time $2^{\mathcal{O}(k^2)}$, it follows that $|\widehat{\text{ALL}}| = 2^{\mathcal{O}(k^2)}$ and it can be enumerated in $2^{\mathcal{O}(k^2)}n$ time. \square

In the rest of this section, we only require pairings and templates that are extended to all of $V(R)$. For convenience, we abuse the notation to denote the extension of a collection of pairings and templates, $\mathcal{A} \in \text{ALL}$ to all of $V(R)$, by $\{\text{pairing}_v\}_{v \in V(R)}$ and $\{\text{template}_v\}_{v \in V(R)}$, respectively. The following corollary follows from the definition of $\widehat{\text{ALL}}$ and Corollary 9.2.

Corollary 9.3. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths. Let R be a backbone Steiner tree. Then, there exists a simplified weak linkage that is discretely homotopic in H to some solution of (G, S, T, g, k) and satisfies the following property: There exists $\{(\widehat{\text{pairing}}_v, \widehat{\text{template}}_v)\}_{v \in V(R)} \in \widehat{\text{ALL}}$ such that $\{\widehat{\text{pairing}}_v\}_{v \in V(R)}$ is the pairing of \mathcal{W} , $\{\widehat{\text{template}}_v\}_{v \in V(R)}$ is the template of \mathcal{W} .*

Stitching of Weak Linkages. Let us now introduce the notion of a stitching, which gives a localized view of a weak linkage pushed onto R at each vertex in $V(R)$. Intuitively, the stitching at a vertex $v \in V(R)$ is function on the set of edges incident on v , that maps each edge to the next (or the previous) edge in a weak linkage. Note that, only a subset of the edges incident on v may participate in a weak linkage. Therefore, we introduce the following notation: an edge is mapped to \perp to indicate that it is not part of weak linkage. Also recall that, for a vertex $v \in V(R)$, order_v is an enumeration of the edges in $\widehat{E}_R(v)$ in either clockwise or anticlockwise order, where $\widehat{E}_R(v) = \{e \in E_H(v) \mid e \text{ is parallel to an edge } e' \in E(R)\}$.

Definition 9.9 (Stitching at a Vertex). *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths, and let R be a backbone Steiner tree. For a vertex $v \in V(R)$, a function $f_v : E_H(v) \rightarrow E_H(v) \cup \perp$ is a stitching at v if it satisfies the following conditions.*

- For any edge $e \in E_H(v) \setminus \widehat{E}_R(v)$, $f_v(e) = \perp$.
- For a pair of (possibly non-distinct) edges $e, e' \in \widehat{E}_R(v)$, $f_v(e) = e'$ if and only if $f_v(e') = e$.
- If $v \in S \cup T$, then there is exactly one edge such that $f_v(e) = e$. Otherwise, there is no such edge.
- If $e_1, e_2, e_3, e_4 \in \widehat{E}_R(v)$ such that $f_v(e_1) = e_2$ and $f_v(e_3) = e_4$. Then $\{e_1, e_2\}$ and $\{e_3, e_4\}$ are disjoint and non-crossing in order _{v} .²⁰

Let $\{f_v\}_{v \in V(R)}$ be a collection of functions such that f_v is a stitching at v for each $v \in V(R)$. Then, this collection is called a stitching if for every edge $e = \{u, v\} \in E_H(R)$, $f_u(e) = \perp$ if and only if $f_v(e) = \perp$.

Let us now describe the stitching of a weak linkage that is pushed onto R .

Definition 9.10 (Stitching of a Weak Linkage Pushed onto R). Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths, and let R be a backbone Steiner tree. Let \mathcal{W} be a weak linkage pushed onto R . Then we define the stitching of \mathcal{W} as the collection of functions $\{\text{stitch}_v\}_{v \in V(R)}$, where $\text{stitch}_v : E_H(v) \rightarrow E_H(v) \cup \perp$ satisfies the following.

- If there is $W \in \mathcal{W}$ where $e = \{v, w\}$ is the first edge of W , then $v \in S \cup T$ and $\text{stitch}_v(e) = e$.
- If there is $W \in \mathcal{W}$ where $e = \{u, v\}$ is the last edge of W , then $v \in S \cup T$ and $\text{stitch}_v(e) = e$.
- If there is a walk $W \in \mathcal{W}$ such that $e, e' \in E_H(v)$ are consecutive edges with a common endpoint $v \in V(R)$, then $\text{stitch}_v(e) = e'$ and $\text{stitch}_v(e') = e$.
- If $e \in E_H(v)$ is not part of any walk in \mathcal{W} , then $\text{stitch}_v(e) = \perp$.

It is easy to verify that $\{\text{stitch}_v\}_{v \in V(R)}$ is indeed a stitching. Let us make a few more observations on the properties of this stitching.

Observation 9.1. Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths, and let R be a backbone Steiner tree. Let \mathcal{W} be a weak linkage pushed onto R and let $\{\text{stitch}_v\}_{v \in V(R)}$ be the stitching of \mathcal{W} . Let $\{\text{pairing}_v\}_{v \in V(R)}$ and $\{\text{template}_v\}_{v \in V(R)}$ be the pairing and template of v , respectively. Then the following holds.

- Let $e_i, e'_j \in E_H(v)$, then $\text{stitch}_v(e_i) = e'_j$ if and only if $\text{stitch}_v(e'_j) = e_i$.
- Let $e, e' \in E_R(v)$. Then $(e, e') \in \text{pairing}_v$ if and only if there is a pair e_i, e'_j of edges in $E_H(R)$, where e_i is parallel to e and e'_j is parallel to e' , such that $\text{stitch}_v(e_i) = e'_j$ and $\text{stitch}_v(e'_j) = e_i$. Further, the number of pairs of parallel edges is equal to $\text{template}_v(e, e')$.
- If e_i, e'_j and e_p^*, \widehat{e}_q are pairs of edges in $E_H(v)$ such that $\text{stitch}_v(e_i) = e'_j$ and $\text{stitch}_v(e_p^*) = \widehat{e}_q$, then the pairs e_i, e'_j and e_p^*, \widehat{e}_q are non-crossing in order _{v} .
- If the multiplicity of \mathcal{W} is upperbounded by ℓ , then for each edge $e = \{u, v\} \in E(R)$, $|\{e' \in E(H) \mid e' \text{ is parallel to } e \text{ and } \text{stitch}_v(e') \neq \perp\}| \leq k \cdot \ell$.

²⁰That is, in a clockwise (or anticlockwise) enumeration of $\widehat{E}_R(v)$ starting from e_1 , these edges occur as either e_1, e_2, e_3, e_4 or e_1, e_3, e_4, e_2 , where without loss of generality we assume that e_3 occurs before e_4 in this ordering.

9.2 Translating a Template Into a Stitching

Given (i) an instance $I = (G, S, T, g, k)$ of Planar Disjoint Paths, (ii) a backbone Steiner tree R , and (iii) a collection $\{(\text{pairing}_v, \text{template}_v)\}_{v \in V(R)} \in \widehat{\text{ALL}}$, our current objective is to either determine that $\{(\text{pairing}_v, \text{template}_v)\}_{v \in V(R)}$ is invalid or construct a multiplicity function ℓ and a stitching $\{f_v\}_{v \in V(R)}$ to reconstruct the a weak linkage. The cases where we determine that $\{(\text{pairing}_v, \text{template}_v)\}_{v \in V(R)}$ is invalid will be (some of the) cases where there exists no simplified weak linkage whose pairing and template are $\{\text{pairing}_v\}_{v \in V(R)}$ and $\{\text{template}_v\}_{v \in V(R)}$, respectively. Let us begin with the notion of multiplicity function ℓ of a collection of pairings and templates, as follows.

Definition 9.11 (Multiplicity Function). *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths, and R be a Steiner tree. Let $\mathcal{A} = \{(\text{pairing}_v, \text{template}_v)\}_{v \in V(R)} \in \widehat{\text{ALL}}$. For every vertex $v \in V^*(R)$, let ℓ_v be the function that assigns $\sum_{e':(e,e') \in \text{pairing}_v} \text{template}_v((e, e'))$ to every edge $e \in E(R)$ incident to v . If one of the following conditions is satisfied, then the multiplicity function extracted from \mathcal{A} is invalid.*

1. *There exists an edge $e = \{u, v\}$ such that $u, v \in V^*(R)$ and $\ell_u(e) \neq \ell_v(e)$.*
2. *There exists a terminal $v \in S \cup T$ such that $\text{pairing}_v = \emptyset$.*

Otherwise, the multiplicity function extracted from \mathcal{A} is valid and it is the function $\ell : E_{1,3+}(R) \rightarrow \mathbb{N}_0$ such that for each $e \in E_{1,3+}(R)$, $\ell(e) = \ell_v(e)$ where v is an endpoint of e in $V^(R)$.²¹*

Let \mathcal{W} be a weak linkage pushed onto R . The *multiplicity function* of a \mathcal{W} is defined as the multiplicity function ℓ extracted from \mathcal{A} , the pairings and templates of \mathcal{W} . It is clear that the multiplicity of \mathcal{W} is $\max_{e \in E(R)} \ell(e)$.

Observation 9.2. *Let (G, S, T, g, k) be an instance of Planar Disjoint Paths with a simplified weak linkage \mathcal{W} , and let ℓ is the multiplicity function of \mathcal{W} . For any $e \in E(R)$, $\ell(e) \leq \alpha_{\text{mul}}(k)$.*

Having extracted a multiplicity function, we turn to extract a stitching. Towards this, recall the embedding of H with respect to R , and the resulting enumeration of edges around vertices in $V(R)$ (see Section 6.6). Let us now describe the stitching extraction at a terminal vertex.

Definition 9.12 (Stitching Extraction at Terminals). *Let (G, S, T, g, k) be a nice instance of Planar Disjoint Paths. Let R be a Steiner tree. Consider a collection $\mathcal{A} = \{(\text{pairing}_v, \text{template}_v)\}_{v \in V(R)} \in \widehat{\text{ALL}}$. Let ℓ be the multiplicity function extracted from \mathcal{A} , and suppose that ℓ is valid. Let $v \in S \cup T$, and let e^* be the unique edge in $E(R)$ incident to v . If $\ell(e^*)$ is an even number, then the local stitching extracted from \mathcal{A} at v is invalid. Otherwise, the stitching extracted from \mathcal{A} at v is valid and it is the involution $f_v : E_H(v) \rightarrow E_H(v) \cup \perp$ defined as follows.*

$$f_v(e) = \begin{cases} e_{\ell(e)+1-i}^* & \text{if } e = e_i^* \text{ and } 1 \leq i \leq \ell(e) \\ \perp & \text{otherwise.} \end{cases}$$

Next, we describe how to extract a stitching at a vertex $v \in V_{=2}(R) \cup V_{\geq 3}(R)$.

Definition 9.13 (Stitching Extraction at Non-Terminals). *Let (G, S, T, g, k) be a nice instance of Planar Disjoint Paths. Let R be a backbone Steiner tree. Consider a collection $\mathcal{A} = \{(\text{pairing}_v, \text{template}_v)\}_{v \in V(R)} \in \widehat{\text{ALL}}$. Let ℓ be the multiplicity function extracted from \mathcal{A} , and suppose that ℓ is valid. Let $v \in V_{=2}(R) \cup V_{\geq 3}(R)$, and suppose that let e^1, e^2, \dots, e^r denote the arcs in $E(R)$ incident to v enumerated as per order_v starting from e^1 . Then the define a function $f_v : E_H(v) \rightarrow E_H(v) \cup \perp$ as follows.*

²¹The choice of the endpoint when both belong to $V^*(R)$ is immaterial by the definition of invalidity.

- For each $(e, e') \in \text{pairing}_v$ such that $\text{template}_v(e, e') > 0$, where e occurs before e' in order_v , let $\text{inner}(e, e') = \{e^* \in E_R(v) \mid e^* \text{ occurs between } e \text{ and } e' \text{ in } \text{order}_v\}$, and $\text{outer}(e, e') = \{e^* \in E_R(v) \mid \text{either } e^* \text{ occurs before } e \text{ or occurs after } e' \text{ in } \text{order}_v\}$.
- Then, for each $i \in \{1, \dots, \text{template}_v(e, e')\}$, let $f_v(e_{i+x}) = e'_{y-i}$ and $f_v(e'_{y-i}) = e_{i+x}$ where

$$x = \sum_{e^* \in \text{outer}(e, e')} \text{template}_v(e, e^*),$$

and

$$y = 1 + \text{template}_v(e, e') + \sum_{e^* \in \text{inner}(e, e')} \text{template}_v(e, e^*)$$

- For all other edges in $E_H(v)$, define $f_v(e) = \perp$.

If the assignment f_v is fixed point free, then f_v is the stitching extracted from \mathcal{A} at v , which is said to be valid. Otherwise, it is invalid

Lastly, based on Definitions 9.12 and 9.13, we extract the stitching as follows,

Definition 9.14 (Stitching Extraction). Let (G, S, T, g, k) be a nice instance of *Planar Disjoint Paths*. Let R be a Steiner tree. Consider a collection $\mathcal{A} = \{(\text{pairing}_v, \text{template}_v)\}_{v \in V(R)} \in \widehat{\text{ALL}}$. For each $v \in V(R)$, let f_v be the stitching extracted from \mathcal{A} at v . Then the stitching extracted from \mathcal{A} is invalid if it satisfies one of the following conditions.

- There is a vertex $v \in V_{=1}(R)$ such that the stitching extracted from \mathcal{A} at v is invalid.
- There is an edge $e = \{u, v\} \in E(H)$ parallel to an edge in $E(R)$ such that $\text{stitch}_u(e) = \perp$ and $\text{stitch}_v(e) \neq \perp$.

Otherwise, the stitching extracted from \mathcal{A} is valid and defined as the collection $\{f_v\}_{v \in V(R)}$ where f_v is the stitching extracted from \mathcal{A} at v for every $v \in V(R)$.

Less obviously, we also show that in case we are given a collection in $\widehat{\text{ALL}}$ that corresponds to weak linkage, not only is the stitching extracted from that collection valid, but also, most crucially, it is the stitching we were originally given (under the assumption that the pair of flow and stitching we deal with is simplified). In other words, we are able to faithfully reconstruct a stitching from the template of weak linkage. The implicit assumption in this lemma that $\{(\text{pairing}_v, \text{template}_v)\}_{v \in V(R)}$ belongs to $\widehat{\text{ALL}}$ is supported by Corollary 9.3.

Lemma 9.7. Let (G, S, T, g, k) be a good Yes-instance of *Planar Disjoint Paths*, and let R be a backbone Steiner tree. Let \mathcal{W} be a simplified weak linkage in H and let ℓ be the multiplicity function of \mathcal{W} . Consider the collection $\mathcal{A} = \{(\text{pairing}_v, \text{template}_v)\}_{v \in V(R)} \in \widehat{\text{ALL}}$, such that it is the collection of pairings and templates of \mathcal{W} . Let $\{f_v\}_{v \in V(R)}$ be the stitching extracted from \mathcal{A} . Then for every vertex $v \in V_{=1}(R)$, $\text{stitch}_v(e) = f_v(e)$ for every edge $e \in E_H(v)$.

Proof. Let $E^*(v) = \{e_i^* \mid 1 \leq i \leq \ell(e^*)\}$ where e_i^* denotes the i -th parallel copy of e^* , in the enumeration in order_v . Observe that, since \mathcal{W} is simplified, $E^*(v)$ is exactly the set of edges from $E_H(v)$ that appear in \mathcal{W} . Hence, for any edge $e \in E_H(v)$, if $e \notin E^*(v)$ then $\text{stitch}_v(e) = \perp$.

Let us now consider the edges in $E^*(v)$. Since \mathcal{W} is a weak linkage, exactly one walk, say W_1 , that has the vertex v as an endpoint. Any other walk in \mathcal{W} contains an even number of edges from $E_H(v)$, and since \mathcal{W} is pushed onto R , these edges are all parallel copies of e^* . Hence, the walks in \mathcal{W} contain an odd number of parallel copies of e^* in total, i.e. $\ell(e^*)$ is an odd number. Since v is an endpoint of $W_1 \in \mathcal{W}$, there is exactly one edge in $e_z^* \in E^*(v)$ such that $\text{stitch}_v(e_z^*) = e_z^*$.

We claim that $z = \frac{\ell(e^*)+1}{2}$. Towards this, let us argue that for any edge e_i^* , where $i < z$, if $\text{stitch}_v(e_i^*) = e_j^*$ then $j > z$. Suppose not, and without loss of generality assume that $i < j < z$. Let us choose i such that $|j - i|$ is minimized, and note that $j \neq i$. Consider the collection of edges e_p^* such that $i < p < j$. If this collection is empty, i.e. $j = i + 1$, then observe that the edges (e_i, e_j) form a U-turn, since $\text{stitch}_v(e_i) = e_j$ only if they were consecutive edges of some walk in \mathcal{W} and there is no edge in the strict interior of the cycle formed by the parallel edges e_i and e_j . Otherwise this collection is non-empty, then observe that if $\text{stitch}_v(e_p^*) = e_q^*$ then $i < q < j$. Indeed, if this were not the case then the pairs e_i^*, e_j^* and e_p^*, e_q^* are crossing at v , since they occur as $e_i^* < e_p^* < e_j^* < e_q^*$ in order_v . This contradicts the weak linkage \mathcal{W} is non-crossing. Otherwise, $i < q < j$ and hence $|q - p| < |j - i|$. But this contradicts the choice of i . Hence, for every $i < z$, $\text{stitch}_v(e_i^*) = e_j^*$ where $j > z$. A symmetric argument holds for the other case, i.e. if $i > z$ then $\text{stitch}_v(e_i^*) = e_j^*$ where $j < z$. Therefore, we can conclude that $z = \frac{\ell(e^*)+1}{2}$, and hence $\text{stitch}_v(e_z^*) = e_{\ell(e^*)+1-z}^*$.

Let us now consider the other edges in $E^*(v)$. Suppose that there exist integers $1 \leq i, p \leq \ell(e^*)$ such that $\text{stitch}_v(e_i^*) = e_j^*$, $\text{stitch}_v(e_p^*) = e_q^*$ such that $i < p < z$ and $z < j < q$. Then it is clear that the pairs e_i^*, e_j^* and (e_p^*, e_q^*) are crossing at v , which is a contradiction. Therefore, if $i < p < z$ then $z < q < j$, and this holds for every choice of i and p . A symmetric argument holds in the other direction, i.e. if $i > p > z$ and $\text{stitch}_v(e_i^*) = e_j^*$ and $\text{stitch}_v(e_p^*) = e_q^*$, then $j < q < z$. Now we claim that for any $i \in \{1, 2, \dots, \ell(e^*)\}$, if $\text{stitch}_v(e_i^*) = e_j^*$ then $j = \ell(e^*) + 1 - i$. Suppose not, and consider the case $i < z$, and further let $j < \ell(e^*) + 1 - i$. Then observe that, for any edge $e_p^* \in \{e_{i+1}^*, \dots, e_{z-1}^*\}$, $\text{stitch}_v(e_p^*) \in \{e_{z+1}^*, \dots, e^{*j-1}\}$. But $|\{e_{i+1}^*, \dots, e_{z-1}^*\}|$ is strictly larger than $|\{e_{z+1}^*, \dots, e^{*j-1}\}|$, which is a contradiction to the definition of stitch_v . Hence, $j \geq \ell(e^*) + 1 - i$. A symmetric argument implies that $j \leq \ell(e^*) + 1 - i$. Therefore, for any $i < z$, $\text{stitch}_v(e_i^*) = e_{\ell(e^*)+1-i}^*$. We can similarly argue that for $i > z$ $\text{stitch}_v(e_i^*) = e_{\ell(e^*)+1-i}^*$. Since we have already shown that $\text{stitch}_v(e_z^*) = e_z^*$, this concludes the proof of this lemma. \square

Lemma 9.8. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths, and let R be a backbone Steiner tree. Let \mathcal{W} be a simplified weak linkage in H and let ℓ be the multiplicity function of \mathcal{W} . Consider the collection $\mathcal{A} = \{(\text{pairing}_v, \text{template}_v)\}_{v \in V(R)} \in \widehat{\text{ALL}}$, such that it is the collection of pairings and templates of \mathcal{W} . Let $\{f_v\}_{v \in V(R)}$ be the stitching extracted from \mathcal{A} . Then, for every vertex $v \in V_{=2}(R) \cup V_{\geq 3}(R)$, $\text{stitch}_v(e) = f_v(e)$ for all edges $e \in E_H(v)$.*

Proof. Let ℓ be the multiplicity function of the simplified weak linkage \mathcal{W} . Then, as \mathcal{W} is canonical, for each edge $e \in E_R(v)$ with a parallel copy e_i , $\text{stitch}_v(e_i) \neq \perp$ if and only if $i \in \{1, 2, \dots, \ell(e^*)\}$. Since \mathcal{W} is a sensible and $v \notin V_1(R)$, it cannot be the endpoint of any walk in \mathcal{W} . Hence, any walk contains an even number of edges from $E_H(v)$, and further any such edge is a parallel copy of an edge in $E_R(v) = \{e^1, e^2, \dots, e^r\}$, where these edges are enumerated according to order_v . Note that, the collections of parallel copies of these edges also occur in the same manner in order_v . We present our arguments in three steps.

Claim 9.1. *Consider a pair of edges $(e, e') \in \text{pairing}_v$, such that $\text{template}_v(e, e') > 0$. Then stitch_v maps each edge in $\{e_{(x_{e,e'}+1)}, \dots, e_{(x_{e,e'}+\text{template}_v(e,e'))}\}$ to some edge in $\{e'_1, e'_2, \dots, e'_{\ell(e')}$, and vice versa.*

Proof. Suppose not, and consider the case where e occur before e' in order_v . Consider a parallel copy of e , say $e_i \in \{e_{(x_{e,e'}+1)}, \dots, e_{(x_{e,e'}+\text{template}_v(e,e'))}\}$ such that $\text{stitch}_v(e_i) = \widehat{e}_j$, where $\widehat{e} \in E_R(v)$ and \widehat{e}_j is the j -th parallel copy of \widehat{e} . Let us choose e' (with respect to e) so that the $x_{e,e'}$ is minimized, and then choose e_i such that i is minimized. Here, note that $i > x_{e,e'}$. Let us argue that $\widehat{e} = e'$. Suppose not, and note that $(e, \widehat{e}) \in \text{pairing}_v$ and $\text{template}_v(e, \widehat{e}) > 0$. Then we have three cases depending on the position of these edges in order_v , either $e < \widehat{e} < e'$, or $\widehat{e} < e < e'$, or $e < e' < \widehat{e}$. Consider the first case, and note that every parallel copy of \widehat{e} occurs before all parallel copies of e' and after all parallel copies of e in order_v . We claim that for any

$e_p \in \{e_{i+1}, \dots, e_{\ell(e)}\}$, $\text{stitch}_v(e_p) \notin \{e'_1, \dots, e'_{\ell(e')}\}$. If this claim were false, then observe that, as $e < \widehat{e} < e'$, we have $e_i < e_p < \widehat{e}_j < \text{stitch}_v(e_p)$ in order_v . Hence e_i, \widehat{e}_j and $e_p, \text{stitch}_v(e_p)$ are crossing pairs at v , in the weak linkage \mathcal{W} , which is a contradiction. On the other hand, if $\text{stitch}_v(e_p) \notin \{e'_1, \dots, e'_{\ell(e')}\}$ for any $e_p \in \{e_{i+1}, \dots, e_{\ell(e)}\}$, then we claim that stitch_v maps strictly fewer than $\text{template}_v(e, e')$ edges from $\{e_1, \dots, e_{\ell(e)}\}$ to $\{e'_1, \dots, e'_{\ell(e')}\}$. Indeed, we choose e' such that $x_{e, e'}$ is minimized, and hence the edges in $\{e_1, \dots, e_{x_{e, e'}}\}$ are not mapped to any edge in $\{e'_1, \dots, e'_{\ell(e')}\}$. And since, no edge in $\{e_i, e_{i+1}, \dots, e_{\ell(e)}\}$ maps to $\{e'_1, \dots, e'_{\ell(e')}\}$, only the edges in $\{e_{(x_{e, e'}+1)}, \dots, e_{i-1}\}$ remain, which is strictly fewer than $\text{template}_v(e, e')$. But this contradicts the definition of $\text{template}_v(e, e')$. Hence, it cannot be the case that $e < \widehat{e} < e'$ in order_v . Next, consider the case when $e < e' < \widehat{e}$. Note that $\widehat{e} \in \text{outer}(e, e')$, and by definition $x_{e, \widehat{e}} < x_{e, e'}$. Since we choose e' to minimize $x_{e, e'}$, and we didn't choose $e' = \widehat{e}$, stitch_v maps the edges in $\{e_{(x_{e, \widehat{e}}+1)}, \dots, e_{(x_{e, \widehat{e}}+\text{template}_v(e, \widehat{e}))}\}$ to $\text{template}_v(e, \widehat{e})$ edges in $\{\widehat{e}_1, \dots, \widehat{e}_{\ell(\widehat{e})}\}$. Therefore, if $\text{stitch}_v(e_i) = \widehat{e}_j$, then there are $\text{template}_v(e, \widehat{e}) + 1$ parallel copies of e that are mapped to parallel copies \widehat{e} , which is a contradiction. Hence it is not possible that $e < e' < \widehat{e}$. The last case, $\widehat{e} < e < e'$ is similar to the previous case, since $\widehat{e} \in \text{outer}(e, e')$ in this case as well. Hence, we conclude that if $\text{stitch}_v(e_i) = \widehat{e}_j$ then $\widehat{e} = e'$. Therefore, when e occurs before e' in order_v , stitch_v maps each edge in $\{e_{(x_{e, e'}+1)}, \dots, e_{(x_{e, e'}+\text{template}_v(e, e'))}\}$ to some edge in $\{e'_1, e'_2, \dots, e'_{\ell(e')}\}$.

By a symmetric argument, we obtain that for any edge in $\{e'_{(y_{e, e'}-\text{template}_v(e, e'))}, \dots, e'_{(y_{e, e'}-1)}\}$ maps to an edge in $\{e_1, e_2, \dots, e_{\ell(e)}\}$.²² \diamond

We now proceed to further restrain the mapping of edges to the ranges determined by $x_{e, e'}$, $y_{e, e'}$ and $\text{template}_v(e, e')$.

Claim 9.2. *Consider a pair $(e, e') \in \text{pairing}_v$ such that $\text{template}_v(e, e') > 0$. Then, stitch_v maps $\{e_{(x_{e, e'}+1)}, \dots, e_{(x_{e, e'}+\text{template}_v(e, e'))}\}$ to $\{e'_{(y_{e, e'}-\text{template}_v(e, e'))}, \dots, e'_{(y_{e, e'}-1)}\}$ and vice-versa²³.*

Proof. Suppose not, and without loss of generality assume that e occurs before e' in order_v . Then consider the case when there is an edge $e_i \in \{e_{(x_{e, e'}+1)}, \dots, e_{(x_{e, e'}+\text{template}_v(e, e'))}\}$ such that $\text{stitch}_v(e_i) = e'_j$, where either $j > y_{e, e'} - 1$ or $j < y_{e, e'} - \text{template}_v(e, e')$. Then consider the collection $\{e'_j\} \cup \{e'_{(y_{e, e'}-\text{template}_v(e, e'))}, \dots, e'_{(y_{e, e'}-1)}\}$, and observe that each edge in this collection is mapped to a distinct edge in $\{e_1, e_2, \dots, e_{\ell(e)}\}$. But then there are $\text{template}_v(e, e') + 1$ edges in $\{e'_1, \dots, e'_{\ell(e')}\}$ that map to an edge in $\{e_1, \dots, e_{\ell(e)}\}$ under stitch_v . This is a contradiction to the definition of $\text{template}_v(e, e')$. \diamond

Finally, we show that stitch_v is equal to f_v .

Claim 9.3. *Consider a pair $(e, e') \in \text{pairing}_v$ such that $\text{template}_v(e, e') > 0$. Then for each $i \in \{1, 2, \dots, \text{template}_v(e, e')\}$, $\text{stitch}_v(e_{x_{e, e'}+i}) = e'_{y_{e, e'}-i}$ and $\text{stitch}_v(e'_{y_{e, e'}-i}) = e_{x_{e, e'}+i}$.*

Proof. Suppose not and consider the case when $\text{stitch}_v(e_{x_{e, e'}+i}) = e'_j$ where $j \neq y_{e, e'} - i$. Note that $e'_j \in \{e'_{(y_{e, e'}-\text{template}_v(e, e'))}, \dots, e'_{(y_{e, e'}-1)}\}$ by previous arguments. Consider the case when $j > y_{e, e'} - i$. We claim that, the edges in $\{e'_{(y_{e, e'}-\text{template}_v(e, e'))}, \dots, e'_{(j-1)}\}$ must map to the edges in $\{e_{(x_{e, e'}+i+1)}, \dots, e_{(x_{e, e'}+\text{template}_v(e, e'))}\}$. If not, then consider an edge $e'_p \in \{e'_{(y_{e, e'}-\text{template}_v(e, e'))}, \dots, e'_{(j-1)}\}$ such that $\text{stitch}_v(e'_p) = e_q$ where $q < x_{e, e'} + i$. Then consider the pairs $e_{x_{e, e'}+i}, e_j$ and e_p, e_q in order_v , and observe that $e_q < e_{x_{e, e'}+i} < e'_p < e'_j$ in order_v . Then these pairs of edges are crossing at v , which is a contradiction to the fact that \mathcal{W} is weak linkage. On the other hand, $\left| \{e'_{(y_{e, e'}-\text{template}_v(e, e'))}, \dots, e'_{(j-1)}\} \right|$ is strictly larger

²²Note that, this is equivalent to the case when e' occurs before e in order_v . Here, we obtain a contradiction by choosing e (with respect to e') that maximizes $y_{e, e'}$, and then choosing the maximum i such that $y_{e, e'} - \text{template}_v(e, e') \leq i \leq y_{e, e'}$ and $\text{stitch}_v(e_i) \notin \{e_1, \dots, e_{\ell(e)}\}$.

²³Note that, by definition of stitch_v , this immediately implies the other direction.

than $\left| \{e_{(x_{e,e'}+i+1)}, \dots, e_{(x_{e,e'}+\text{template}_v(e,e'))}\} \right|$, which is again a contradiction, since all edges in $\{e'_1, \dots, e'_{\ell(e')}\}$ are mapped to distinct edges by stitch_v , and they are not mapped to \perp . By symmetric arguments, the case when $j < y_{e,e'} - i$ also leads to a contradiction. \diamond

Now, by considering all pairs in pairing_v and applying the above claims, we obtain that $\text{stitch}_v = f_v$ for all $v \in V_{=2}(R) \cup V_{\geq \epsilon}(R)$. This concludes the proof of this lemma. \square

The following lemma is a corollary of Lemma 9.7 and Lemma 9.8

Lemma 9.9. *Let (G, S, T, g, k) be a nice instance of Planar Disjoint Paths. Let R be a Steiner tree. Let \mathcal{W} be a simplified weak linkage, and let $\mathcal{A} = \{(\text{pairing}_v, \text{template}_v)\}_{v \in V(R)} \in \widehat{\text{ALL}}$ be pairing and template of \mathcal{W} . Let $\{f_v\}_{v \in V(R)}$ be the stitching extracted from \mathcal{A} . Then $f_v = \text{stitch}_v$ for every vertex $v \in V(R)$.*

Now, we consider the computational aspect of the definitions considered so far in this section.

Lemma 9.10. *Let (D, S, T, g, k) be a nice instance of Directed Planar Disjoint Paths. Let R be a Steiner tree. Let $\mathcal{A} = \{(\text{pairing}_v, \text{template}_v)\}_{v \in V(R)} \in \widehat{\text{ALL}}$. Then, the multiplicity function extracted from \mathcal{A} can be computed in time $k^{\mathcal{O}(1)}n$, and the stitching extracted from \mathcal{A} can be computed in time $2^{\mathcal{O}(k)}n$.*

Proof. First we consider the computation of the multiplicity function $\ell = \{\ell_v\}_{v \in V(R)}$ extracted from \mathcal{A} according to Definition 9.11. Note that, for every vertex $v \in V(R)$, we have that $|\text{pairing}_v| = \mathcal{O}(k)$ and the bitsize of the numbers assigned by template_v are bounded by $2^{\mathcal{O}(k)}$. Therefore, ℓ_v can be computed in $2^{\mathcal{O}(k)}$ time for each $v \in V(R)$, taking a total of $2^{\mathcal{O}(k)}n$ time. Now, note that for any vertex $v \in V(R)$, it holds that $|\ell_v(e)| = 2^{\mathcal{O}(k)}$ for any edge $e \in E_H(v)$ (because template_v is a $2^{\mathcal{O}(k)}$ -template).

Let $\{f_v\}_{v \in V(R)}$ be the stitching extracted from \mathcal{A} by Definition 9.12. Observe that when describing the stitching f_v extracted at a vertex $v \in V(R)$, we only need to describe it for the parallel copies of edges in $E(R)$, and then only for the parallel copies $\{e_1, e_2, \dots, e_{\ell(v)}\}$ of $e \in E(R)$, where ℓ is the multiplicity function extracted from \mathcal{A} . For all other edges and parallel copies, the stitching maps them to \perp . Since $\ell(e) \leq \alpha_{\text{mul}}(k)$, and the tree R has at most $2k$ leaves, the stitching at each vertex can be described by a collection of $\mathcal{O}(k \cdot \alpha_{\text{mul}}) = 2^{\mathcal{O}(k)}$ pairs of edges in $E_H(v) \times E_H(v)$. Further, by the construction described in Definitions 9.12 and 9.13, the stitching f_v at each vertex $v \in V(R)$ can be constructed in time $2^{\mathcal{O}(k)}$ time. Therefore, the collection $\{f_v\}_{v \in V(R)}$ can be constructed in $2^{\mathcal{O}(k)}n$ time. Finally, we need to test if this collection is a valid stitching, as described in Definition 9.14, which can be done by picking each edge $e \in E(R)$ and testing the parallel copies $\{e_1, e_2, \dots, e_{\ell(e)}\}$ one by one, which again takes $2^{\mathcal{O}(k)}n$ time. Hence the total time required to extract the stitching is $2^{\mathcal{O}(k)}n$. \square

9.3 Reconstruction of Weak Linkages from Templates

Now we describe the construction of a weak linkage from a valid stitching.

Definition 9.15 (Weak Linkage of a stitching.). *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths, and let R be a backbone Steiner tree. Let $\{f_v\}_{v \in V(R)}$ be a stitching and suppose that it is valid. Then the weak linkage \mathcal{W} constructed from f_v is obtained as follows.*

- For each $v \in S \cup T$, Let $e_v \in E_H(v)$ be the unique edge such that $f_v(e_v) = e_v$.
- Then the walk W_v is defined as the sequence of edges $e_0, e_1, e_2, \dots, e_{p_v}$, where $e_{v_0} = e_v$, and for each $i \in \{0, \dots, p_v - 1\}$, the edge $e_i = \{v_i, v_{i+1}\}$ satisfies (i) $f_{v_{i+1}}(e_i) = e_{i+1}$ where $v_0 = v$; and (ii) $f_{v_{p_v}}(e_{p_v}) = p_v$.

- We iteratively construct a sequence of walks W_{v_1}, W_{v_2}, \dots , where the walk W_{v_i} starts from a vertex v_i that is not the endpoint of any of the previous walks. Finally, we output \mathcal{W} as the collection of these walks.

It is clear that running time for the construction of a weak linkage from a stitching $\{f_v\}_{v \in V(R)}$ is upperbounded by the number of pairs of edges in $E_H(V(R))$ such that they are images of each other in the stitching. The following observation follows directly from Definition 9.15 and Definition 9.10. and the fact that

Observation 9.3. *Let (G, S, T, g, k) be a good instance of Planar Disjoint Paths, and let R be a backbone Steiner tree. Let \mathcal{W} be a weak linkage that is pushed onto R , and let $\{\text{stitch}_v\}_{v \in V(R)}$ be the stitching of \mathcal{W} . Then the weak linkage constructed from this stitching is equal to \mathcal{W} .*

Let \mathcal{W}_{ALL} denote the collection of weak linkages extracted from $\widehat{\text{ALL}}$. The following Lemma is the main result of this section.

Lemma 9.11. *Let (G, S, T, g, k) be a good Yes-instance of Planar Disjoint Paths, and let R be a backbone steiner tree. Then, there exists a collection of $2^{\mathcal{O}(k^2)}$ simplified linkages \mathcal{W}_{ALL} such that there is a weak linkage $\mathcal{W} \in \mathcal{W}_{\text{ALL}}$ that is discretely homotopic in H to some solution of (G, S, T, g, k) . Further, collection can be enumerated in $2^{\mathcal{O}(k^2)}n$ time.*

Proof. By Lemma 8.19, the given instance has a solution that is discretely homotopic to some simplified weak linkage \mathcal{W} , and by Corollary 9.3, the pairing and template of \mathcal{W} , denoted by \mathcal{A} lies in the collection $\widehat{\text{ALL}}$. Then, by Lemma 9.9, the stitching $\{f_v\}_{v \in V(R)}$ is equal to $\{\text{stitch}_v\}_{v \in V(R)}$, the stitching of \mathcal{W} , and it can be computed in $2^{\mathcal{O}(k)}n$ time by Lemma 9.10. Finally, we can construct a weak linkage \mathcal{W}' from the stitching $\{f_v\}_{v \in V(R)}$, and by Observation 9.3, $\mathcal{W}' = \mathcal{W}$. Note that, as \mathcal{W} is a simplified weak linkage, its multiplicity is upperbounded by $\alpha_{\text{mul}}(k) = 2^{\mathcal{O}(k)}$. Hence by Observation 9.1 and 9.2, the number of pairs in $(e, e') \in E_H(v) \times E_H(v)$ such that $f_v(e) = e'$ is upperbounded by $k \cdot \alpha_{\text{mul}}(k) = 2^{\mathcal{O}(k)}$. Hence, it is clear that we reconstruct \mathcal{W} from the stitching $\{f_v\}_{v \in V(R)}$ in time $2^{\mathcal{O}(k)}n$ time.

To enumerate the collection \mathcal{W}_{ALL} , we iterate over $\widehat{\text{ALL}}$. For each $\mathcal{A} \in \widehat{\text{ALL}}$, we attempt to construct a stitching and if it returns an invalid stitching, we move on to next iteration. Otherwise, we construct a weak linkage from this stitching and output it. Observe that for each $\mathcal{A} \in \widehat{\text{ALL}}$ we can compute the corresponding weak linkage \mathcal{W} , if it exists, in time $2^{\mathcal{O}(k)}n$ time. Since $|\widehat{\text{ALL}}| = 2^{\mathcal{O}(k^2)}$, clearly $|\mathcal{W}_{\text{ALL}}| = 2^{\mathcal{O}(k^2)}$ and it can be enumerated in $2^{\mathcal{O}(k^2)}n$ time, we can enumerate \mathcal{W}_{ALL} in $2^{\mathcal{O}(k^2)}n$ time. \square

10 The Algorithm

Having set up all required definitions and notions, we are ready to describe our algorithm. Afterwards, we will analyze its running time and prove its correctness.

10.1 Execution of the Algorithm

We refer to this algorithm as PDPAlg. It takes as input an instance $(\tilde{G}, \tilde{S}, \tilde{T}, \tilde{g}, k)$ of Planar Disjoint Paths, and its output is the decision whether this instance is a Yes-instance. The specification of the algorithm is as follows.

Step I: Preprocessing. First, PDPAlg invokes Corollary 4.1 to transform $(\tilde{G}, \tilde{S}, \tilde{T}, \tilde{g}, k)$ into an equivalent good instance (G, S, T, g, k) of Planar Disjoint Paths where $|V(G)| = \mathcal{O}(|V(\tilde{G})|)$.

Step II: Computing a Backbone Steiner Tree. Second, PDPAlg invokes Lemma 6.8 with respect to (G, S, T, g, k) to compute a backbone Steiner tree, denoted by R . Then, PDPAlg computes the embedding of H with respect to R (Section 6.6).

Step III: Looping on \mathcal{W}_{ALL} . Now, PDPAlg invokes Lemma 9.11 to enumerate \mathcal{W}_{ALL} . For each weak linkage $\mathcal{W} \in \mathcal{W}_{\text{ALL}}$, the algorithm applies the algorithm of Corollary 5.1 to (G, S, T, g, k) and \mathcal{W} ; if the algorithm finds a solution, then PDPAlg determines that $(\tilde{G}, \tilde{S}, \tilde{T}, \tilde{g}, k)$ is Yes-instance and terminates.

Step IV: Reject. PDPAlg determines that $(\tilde{G}, \tilde{S}, \tilde{T}, \tilde{g}, k)$ is a No-instance and terminates.

10.2 Running Time and Correctness

Let us first analyze the running time of PDPAlg.

Lemma 10.1. *PDPAlg runs in time $2^{\mathcal{O}(k^2)}n^{\mathcal{O}(1)}$.*

Proof. By Corollary 4.1, the computation of (G, S, T, g, k) in Step I is performed in time $2^{\mathcal{O}(k)}n^2$. By Lemma 6.8, the computation of R in Step II is performed in time $2^{\mathcal{O}(k)}n^{3/2} \log^3 n$. Moreover, the computation of R can clearly be done in time linear in n . Let $H = H_G$ be the radial completion of G enriched with $4|V(G)| + 1$ parallel copies of each edge, and note that $|V(H)| = \mathcal{O}(|V(G)|)$. By Observation 6.9, we can compute the embedding of H with respect to R in time $\mathcal{O}(n^2)$. By Lemma 9.11, $|\mathcal{W}_{\text{ALL}}| = 2^{\mathcal{O}(k^2)}$, and it can be enumerated in $2^{\mathcal{O}(k^2)}n$ time. Finally, for each $\mathcal{W} \in \mathcal{W}_{\text{ALL}}$, Corollary 5.1 takes $n^{\mathcal{O}(1)}$ time to test if there is a solution that is discretely homotopic to \mathcal{W} . Thus, PDPAlg runs in time $2^{\mathcal{O}(k^2)}n^{\mathcal{O}(1)}$. \square

The reverse direction of the correctness of PDPAlg is trivially true.

Lemma 10.2. *Let $(\tilde{G}, \tilde{S}, \tilde{T}, \tilde{g}, k)$ be an instance of Planar Disjoint Paths. If $(\tilde{G}, \tilde{S}, \tilde{T}, \tilde{g}, k)$ is accepted by PDPAlg, then (G, S, T, g, k) is a Yes-instance.*

Now, we handle the forward direction of the correctness of PDPAlg.

Lemma 10.3. *Let $(\tilde{G}, \tilde{S}, \tilde{T}, \tilde{g}, k)$ be a instance of Planar Disjoint Paths. If $(\tilde{G}, \tilde{S}, \tilde{T}, \tilde{g}, k)$ is a Yes-instance, then (G, S, T, g, k) is accepted by PDPAlg.*

Proof. Suppose that $(\tilde{G}, \tilde{S}, \tilde{T}, \tilde{g}, k)$ is a Yes-instance of Planar Disjoint Paths. Then, by Corollary 4.1, (G, S, T, g, k) is a Yes-instance of Planar Disjoint Paths. Then by Lemma 8.19 and Lemma 9.11, there is a collection \mathcal{W}_{ALL} of $2^{\mathcal{O}(k^2)}$ simplified weak linkages containing at least one simplified weak linkage \mathcal{W}^* , that is discretely homotopic in H to some solution of (G, S, T, g, k) . Here H is the radial completion of G enriched with $4|V(G)| + 1$ parallel copies of each edge. Then in Step III, Corollary 5.1 ensures that we obtain a solution to (G, S, T, g, k) in the iteration we consider \mathcal{W}^* . Hence PDPAlg accepts the instance $(\tilde{G}, \tilde{S}, \tilde{T}, \tilde{g}, k)$. \square

Lastly, we remark that PDPAlg can be easily modified not only to reject or accept the instance $(\tilde{G}, \tilde{S}, \tilde{T}, \tilde{g}, k)$, but to return a solution in case of acceptance, within time $2^{\mathcal{O}(k^2)}n^{\mathcal{O}(1)}$.

References

- [1] I. ADLER, *6th workshop on graph classes, optimization, and width parameters 2013, list of open problems*. http://www.cs.upc.edu/~sedthilk/grow/Open_Problems_GROW_2013.pdf, 2013. 1, 3
- [2] I. ADLER, S. G. KOLLIPOULOS, P. K. KRAUSE, D. LOKSHANOV, S. SAURABH, AND D. M. THILIKOS, *Irrelevant vertices for the planar disjoint paths problem*, J. Comb. Theory, Ser. B, 122 (2017), pp. 815–843. 1, 2, 3, 7, 9, 15, 16
- [3] I. ADLER AND P. K. KRAUSE, *A lower bound for the tree-width of planar graphs with vital linkages*, CoRR, abs/1011.2136 (2010). 3

- [4] J. BASTE AND I. SAU, *The role of planarity in connectivity problems parameterized by treewidth*, Theor. Comput. Sci., 570 (2015), pp. 1–14. [2](#)
- [5] H. L. BODLAENDER, F. V. FOMIN, D. LOKSHANOV, E. PENNINKX, S. SAURABH, AND D. M. THILIKOS, *(Meta) kernelization*, J. ACM, 63 (2016), pp. 44:1–44:69. [7](#), [13](#), [24](#)
- [6] C. CHEKURI AND J. CHUZHUY, *Polynomial bounds for the grid-minor theorem*, J. ACM, 63 (2016), pp. 40:1–40:65. [2](#)
- [7] J. CHUZHUY AND D. H. K. KIM, *On approximating node-disjoint paths in grids*, in Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015, August 24-26, 2015, Princeton, NJ, USA, vol. 40 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015, pp. 187–211. [1](#)
- [8] J. CHUZHUY, D. H. K. KIM, AND S. LI, *Improved approximation for node-disjoint paths in planar graphs*, in Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016, ACM, 2016, pp. 556–569. [1](#)
- [9] J. CHUZHUY, D. H. K. KIM, AND R. NIMAVAT, *New hardness results for routing on disjoint paths*, in Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017, ACM, 2017, pp. 86–99. [1](#)
- [10] ———, *Almost polynomial hardness of node-disjoint paths in grids*, in Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018, ACM, 2018, pp. 1220–1233. [1](#)
- [11] ———, *Improved approximation for node-disjoint paths in grids with sources on the boundary*, in 45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic, vol. 107 of LIPIcs, 2018, pp. 38:1–38:14. [1](#)
- [12] V. COHEN-ADDAD, É. C. DE VERDIÈRE, P. N. KLEIN, C. MATHIEU, AND D. MEIER-FRANKENFELD, *Approximating connectivity domination in weighted bounded-genus graphs*, in Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016, 2016, pp. 584–597. [13](#)
- [13] S. CORNELSEN AND A. KARRENBauer, *Accelerated bend minimization*, J. Graph Algorithms Appl., 16 (2012), pp. 635–650. [29](#)
- [14] M. CYGAN, D. MARX, M. PILIPCZUK, AND M. PILIPCZUK, *The planar directed k -vertex-disjoint paths problem is fixed-parameter tractable*, in 54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA, 2013, pp. 197–206. [1](#), [3](#), [8](#), [37](#), [38](#)
- [15] E. D. DEMAINE, F. V. FOMIN, M. T. HAJIAGHAYI, AND D. M. THILIKOS, *Subexponential parameterized algorithms on bounded-genus graphs and H -minor-free graphs*, J. ACM, 52 (2005), pp. 866–893. [1](#), [2](#)
- [16] E. D. DEMAINE AND M. HAJIAGHAYI, *Linearity of grid minors in treewidth with applications through bidimensionality*, Combinatorica, 28 (2008), pp. 19–36. [2](#)
- [17] G. DING, A. SCHRIJVER, AND P. D. SEYMOUR, *Disjoint paths in a planar graph—a general theorem*, SIAM Journal on Discrete Mathematics, 5 (1992), pp. 112–116. [1](#), [38](#)

- [18] F. V. FOMIN, D. LOKSHTANOV, D. MARX, M. PILIPCZUK, M. PILIPCZUK, AND S. SAURABH, *Subexponential parameterized algorithms for planar and apex-minor-free graphs via low treewidth pattern covering*, in IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, New Brunswick, New Jersey, USA, IEEE Computer Society, 2016, pp. 515–524. [1](#)
- [19] F. V. FOMIN, D. LOKSHTANOV, AND S. SAURABH, *Excluded grid minors and efficient polynomial-time approximation schemes*, J. ACM, 65 (2018), pp. 10:1–10:44. [1](#)
- [20] S. FORTUNE, J. E. HOPCROFT, AND J. WYLLIE, *The directed subgraph homeomorphism problem*, Theor. Comput. Sci., 10 (1980), pp. 111–121. [1](#)
- [21] A. FRANK, *Packing paths, cuts, and circuits—a survey*, Paths, Flows and VLSI-Layout, (1990), pp. 49–100. [1](#)
- [22] M. GROHE, K. KAWARABAYASHI, AND B. A. REED, *A simple algorithm for the graph minor decomposition - logic meets structural graph theory*, in Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013, SIAM, 2013, pp. 414–431. [2](#)
- [23] B. M. P. JANSEN, D. LOKSHTANOV, AND S. SAURABH, *A near-optimal planarization algorithm*, in Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014, 2014, pp. 1802–1811. [13](#), [20](#)
- [24] R. M. KARP, *Reducibility among combinatorial problems*, in Proceedings of a symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, USA, 1972, pp. 85–103. [1](#)
- [25] R. M. KARP, *On the computational complexity of combinatorial problems*, Networks, 5 (1975), pp. 45–68. [1](#)
- [26] K. KAWARABAYASHI, Y. KOBAYASHI, AND B. A. REED, *The disjoint paths problem in quadratic time*, J. Comb. Theory, Ser. B, 102 (2012), pp. 424–435. [1](#)
- [27] K. KAWARABAYASHI AND P. WOLLAN, *A shorter proof of the graph minor algorithm: the unique linkage theorem*, in Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010, ACM, 2010, pp. 687–694. [2](#)
- [28] ———, *A simpler algorithm and shorter proof for the graph minor decomposition*, in Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011, ACM, 2011, pp. 451–458. [2](#)
- [29] P. N. KLEIN AND D. MARX, *A subexponential parameterized algorithm for subset TSP on planar graphs*, in Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014, SIAM, 2014, pp. 1812–1830. [1](#)
- [30] M. R. KRAMER AND J. VAN LEEUWEN, *The complexity of wirerouting and finding minimum area layouts for arbitrary VLSI circuits*, Advances in computing research, 2 (1984), pp. 129–146. [1](#)
- [31] R. J. LIPTON AND K. W. REGAN, *David Johnson: Galactic Algorithms*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 109–112. [2](#)

- [32] D. LOKSHANOV, D. MARX, AND S. SAURABH, *Slightly superexponential parameterized problems*, SIAM J. Comput., 47 (2018), pp. 675–702. [2](#)
- [33] J. F. LYNCH, *The equivalence of theorem proving and the interconnection problem*, ACM SIGDA Newsletter, 5 (1975), pp. 31–36. [1](#)
- [34] B. MOHAR AND C. THOMASSEN, *Graphs on Surfaces*, Johns Hopkins series in the mathematical sciences, Johns Hopkins University Press, 2001. [12](#)
- [35] Y. PERL AND Y. SHILOACH, *Finding two disjoint paths between two pairs of vertices in a graph*, J. ACM, 25 (1978), pp. 1–9. [1](#)
- [36] M. PILIPCZUK, M. PILIPCZUK, P. SANKOWSKI, AND E. J. VAN LEEUWEN, *Network sparsification for steiner problems on planar and bounded-genus graphs*, ACM Trans. Algorithms, 14 (2018), pp. 53:1–53:73. [1](#)
- [37] N. ROBERTSON AND P. D. SEYMOUR, *Graph minors .xiii. the disjoint paths problem*, J. Comb. Theory, Ser. B, 63 (1995), pp. 65–110. [1](#), [2](#)
- [38] ———, *Graph minors. XVI. excluding a non-planar graph*, J. Comb. Theory, Ser. B, 89 (2003), pp. 43–76. [2](#)
- [39] ———, *Graph minors. XXII. irrelevant vertices in linkage problems*, J. Comb. Theory, Ser. B, 102 (2012), pp. 530–563. [2](#)
- [40] N. ROBERTSON, P. D. SEYMOUR, AND R. THOMAS, *Quickly excluding a planar graph*, J. Comb. Theory, Ser. B, 62 (1994), pp. 323–348. [2](#)
- [41] P. SCHEFFLER, *A practical linear time algorithm for disjoint paths in graphs with bounded tree-width*, TU, Fachbereich 3, 1994. [2](#)
- [42] A. SCHRIJVER, *Finding k disjoint paths in a directed planar graph*, SIAM J. Comput., 23 (1994), pp. 780–788. [1](#), [3](#), [4](#), [11](#), [14](#), [15](#)
- [43] A. SCHRIJVER, *Combinatorial optimization: polyhedra and efficiency*, vol. 24, Springer Science & Business Media, 2003. [1](#)

A Properties of Winding Number

In this section we sketch a proof of Proposition [7.1](#) using homotopy. Towards this, we introduce some notation that are extensions of the terms introduced Section [7](#) in the continuous setting. Recall that we have a plane graph $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ and we are interested in the winding number of paths in this graph, where I_{in} and I_{out} are two cycles such that I_{out} is the outer-face, and there are no vertices or edges in the interior of I_{in} . Let us denote the (closed) curves defined by these two curves by ρ_{out} and ρ_{in} , respectively. Then consider the collection of all the points in the plane that lie in the exterior of ρ_{in} and interior of ρ_{out} . The closure of this set of points defines a surface called a *ring*, which we denote by $\text{Ring}(\rho_{\text{in}}, \rho_{\text{out}})$ by abusing notation. Observe that the graph $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ is embedded in this ring, where the vertices of I_{in} and I_{out} lie on ρ_{in} and ρ_{out} respectively.

A curve α in the $\text{Ring}(\rho_{\text{in}}, \rho_{\text{out}})$ *traverses* it if it has one endpoint in ρ_{in} and the other in ρ_{out} . We then orient this curve from its endpoint in ρ_{in} to its endpoint ρ_{out} . A curve β *visits* $\text{Ring}(\rho_{\text{in}}, \rho_{\text{out}})$ if both its endpoints line on either ρ_{in} or ρ_{out} . In this case we orient this curve as follows. We first fix an arbitrary ordering of all points in the curve ρ_{in} and another one for all the points in the curve ρ_{out} . We then orient β from the smaller endpoint to the greater one.

Consider two curves α, α' that are either traversing or visiting $\text{Ring}(\rho_{\text{in}}, \rho_{\text{out}})$ are *homotopic* if there exists a homotopy of the ring that fixes ρ_{in} and ρ_{out} and transforms α to α' . Note that homotopic curves have same endpoints. Two curves β, β' are *transversally intersecting* if $\beta \cap \beta'$ is a finite collection of points. Let us remark that the above orientation is preserved under homotopy, since any two homotopic curves have the same endpoints. Furthermore, when we speak of oriented curves in a ring, it is implicit that such curves are either visitors or traversing the ring. Now we are ready to define the winding number of oriented curves in a ring.

Definition A.1 (Winding Number of Transversally Intersecting Curves). *Two curves α, β in $\text{Ring}(\rho_{\text{in}}, \rho_{\text{out}})$ intersect transversally if $\alpha \cap \beta$ is a finite set of points. For two curves α and β in $\text{Ring}(\rho_{\text{in}}, \rho_{\text{out}})$ that intersect transversally, we define the winding number $\overline{\text{WindNum}}(\alpha, \beta)$ as the signed number of traversings of β along α . That is, for every intersection point of α and β we record $+1$ if β crosses α from left to right, and -1 if it crosses from right to left (of course, with respect to the chosen direction of traversing β) and 0 if it does not cross at that point. The winding number $\overline{\text{WindNum}}(\alpha, \beta)$ is the sum of the recorded numbers.*

It can be easily observed that if α and α' are homotopic curves traversing $\text{Ring}(\rho_{\text{in}}, \rho_{\text{out}})$ and both intersect β transversally, then $\overline{\text{WindNum}}(\alpha, \beta) = \overline{\text{WindNum}}(\alpha', \beta)$. Here we rely on the fact that two homotopic curves have the same end-points. Observe that that every intersection point of the curves α' and β is a traversing point, i.e. the point assigned either $+1$ or -1 . Therefore, we can extend the notion of the winding number to pairs of curves not necessarily intersecting transversally as follows.

Definition A.2 (Winding Number). *If α, β are two curves in $\text{Ring}(\rho_{\text{in}}, \rho_{\text{out}})$, then we define $\overline{\text{WindNum}}(\alpha, \beta)$ to be the winding number $\overline{\text{WindNum}}(\alpha', \beta')$ for any α', β' such that α and α' are homotopic, β and β' are homotopic, and α', β' intersect transversally, and each common point of α' and β' is a traversing point.*

Note that such α', β' always exist and the definition of $\overline{\text{WindNum}}(\alpha, \beta)$ does not depend on the particular curves. Let us now proceed towards a proof of Proposition 7.1.

Lemma A.1. *Suppose α, β, γ are curves traversing a ring $\text{Ring}(\rho_{\text{in}}, \rho_{\text{out}})$ that pairwise intersect transversally. Further, letting a, b, c and a', b', c' be the endpoints of α, β, γ on ρ_{in} and ρ_{out} respectively, suppose that a, b, c are different and appear in the clockwise order on ρ_{in} , and that a', b', c' are different and appear in the clockwise order on ρ_{out} . Then*

$$\overline{\text{WindNum}}(\alpha, \beta) + \overline{\text{WindNum}}(\beta, \gamma) = \overline{\text{WindNum}}(\alpha, \gamma).$$

Proof. First, we argue that we may assume that $\overline{\text{WindNum}}(\alpha, \gamma) = 0$. This can be done as follows. Let $k = \overline{\text{WindNum}}(\alpha, \gamma)$. Glue a ring $\text{Ring}(\rho_{\text{out}}, \rho'_{\text{out}})$ to $\text{Ring}(\rho_{\text{in}}, \rho_{\text{out}})$ along ρ_{out} , for some non-self-traversing closed curve ρ'_{out} that encloses ρ_{out} , thus obtaining ring $\text{Ring}(\rho_{\text{in}}, \rho'_{\text{out}})$. Pick a'', b'', c'' in the clockwise order on ρ'_{out} . Extend α to a curve α' traversing $\text{Ring}(\rho_{\text{in}}, \rho'_{\text{out}})$ using any curve within $\text{Ring}(\rho_{\text{out}}, \rho'_{\text{out}})$ connecting a' with a'' . Next, extend β to β' in the same way, but choose the extending segment so that it does not cross α' . Finally, extend γ to γ' in the same way, but choose the extending segment so that it crosses α' (and thus also β') exactly $-k$ times (where we count signed traversings). Thus we have $\overline{\text{WindNum}}(\alpha', \gamma') = 0$ and if we replace α, β, γ with α', β', γ' , both sides of the postulated equality are decremented by k . Hence, it suffices to prove this equality for α', β', γ' , for which we know that $\overline{\text{WindNum}}(\alpha', \gamma') = 0$.

Having assumed that $\overline{\text{WindNum}}(\alpha, \gamma) = 0$, it remains to prove that $\overline{\text{WindNum}}(\alpha, \beta) + \overline{\text{WindNum}}(\beta, \gamma) = 0$, or equivalently

$$\overline{\text{WindNum}}(\alpha, \beta) + \overline{\text{WindNum}}(\gamma^{-1}, \beta) = 0. \tag{1}$$

Since $\overline{\text{WindNum}}(\alpha, \gamma) = 0$, we may further replace α and γ with homotopic curves that do not cross at all. Note that this does not change the winding numbers in the postulated equality. Hence, from now on we assume that α and γ are disjoint.

Let us connect c with a using an arbitrary curve ϵ through the interior of the disk enclosed by ρ_{in} , and let us connect a' with c' using an arbitrary curve ϵ' outside of the disk enclosed by ρ_{out} . Thus, the concatenation of $\alpha, \epsilon', \gamma^{-1}, \epsilon$ is a closed curve in the plane without self-traversings; call it δ . Then δ separates the plane into two regions R_1, R_2 . Since a, b, c appear in the same order on ρ_{in} as a', b', c' on ρ_{out} , it follows that b and b' are in the same region, say R_1 .

Now consider travelling along β from the endpoint b to the endpoint b' . Every traversing of α or γ along β is actually a traversing of δ that contributes to the left hand side of (1) with $+1$ if on the traversing β passes from R_1 to R_2 , and with -1 if it passes from R_2 to R_1 . Since β starts and ends in R_1 , the total sum of those contributions has to be equal to 0, which proves (1). \square

Lemma A.2. *For any curves α, β, γ traversing a ring $\text{Ring}(\rho_{\text{in}}, \rho_{\text{out}})$, it holds that*

$$|(\overline{\text{WindNum}}(\alpha, \beta) + \overline{\text{WindNum}}(\beta, \gamma)) - \overline{\text{WindNum}}(\alpha, \gamma)| \leq 1.$$

Proof. By slightly perturbing the curves using homotopies, we may assume that they pairwise intersect transversally. Further, we modify the curves in the close neighborhoods of ρ_{in} and ρ_{out} so that we may assume that the endpoints of α, β, γ on ρ_{in} and ρ_{out} are pairwise different and appear in the same clockwise order on both cycles; for the latter property, we may add one traversing between two of the curves, thus modifying one of the numbers $\overline{\text{WindNum}}(\alpha, \beta), \overline{\text{WindNum}}(\beta, \gamma), \overline{\text{WindNum}}(\alpha, \gamma)$ by one. It now remains to use Lemma A.1. \square

Proof of Proposition 7.1. Recall that the graph $\text{Ring}(I_{\text{in}}, I_{\text{out}})$ is embedded in the ring. The first property follows directly from the definition of winding numbers. For the second property, we apply Lemma A.2 to the curves defined by the paths α, β and γ in $\text{Ring}(\rho_{\text{in}}, \rho_{\text{out}})$. \square