

The planar directed k -Vertex-Disjoint Paths problem is fixed-parameter tractable

Marek Cygan ^{*} Dániel Marx [†] Marcin Pilipczuk [‡] Michał Pilipczuk [§]

Abstract

Given a graph G and k pairs of vertices $(s_1, t_1), \dots, (s_k, t_k)$, the k -Vertex-Disjoint Paths problem asks for pairwise vertex-disjoint paths P_1, \dots, P_k such that P_i goes from s_i to t_i . Schrijver [28] proved that the k -Vertex-Disjoint Paths problem on planar directed graphs can be solved in time $n^{O(k)}$. We give an algorithm with running time $2^{2^{O(k^2)}} \cdot n^{O(1)}$ for the problem, that is, we show the fixed-parameter tractability of the problem.

^{*}Institute of Informatics, University of Warsaw, Poland, cygan@mimuw.edu.pl.

[†]Computer and Automation Research Institute, Hungarian Academy of Sciences (MTA SZTAKI), Hungary, dmarx@cs.bme.hu.

[‡]Institute of Informatics, University of Warsaw, Poland, malcin@mimuw.edu.pl.

[§]Department of Informatics, University of Bergen, Norway, michal.pilipczuk@ii.uib.no.

Contents

1	Introduction	3
2	Overview of the algorithm	5
2.1	Irrelevant vertex rule	5
2.2	Decomposition and duality theorems	9
2.3	Bundles and bundle words	13
2.4	Guessing bundle words	16
2.5	Handling a ring component	19
2.6	Summary	21
3	Irrelevant vertex rule	22
3.1	Bends	23
3.2	Paths and segments in bends	27
3.3	Rerouting nested segments	30
3.4	Rerouting in a bend	34
4	Min-max theorems for paths, cycles, and cuts	43
4.1	Framework	43
4.2	Understanding homotopies in a ring	44
4.3	Min-max theorems for alternating cuts	50
4.3.1	Alternations	50
4.3.2	Many alternating cycles, or a short cut of a ring	52
4.3.3	Many alternating paths, or a short circular cut in a ring	52
5	Decomposition	57
6	Bundles and bundle words	63
6.1	Bundle arcs, bundles and component multigraph	64
6.2	Bundle words and spirals	66
6.3	Structure of bundle words and bundle profiles	70
6.4	Ring components: basics	73
6.5	Bundle words with holes	76
6.6	Minimal solution	77
6.7	Ring components: bound on the interaction	78
7	From bundle words to disjoint paths	82
8	Guessing bundle words and their winding numbers	86
8.1	Zooms	87
8.2	Deducing bundle words	90
8.3	Ring components: deducing winding numbers	99
9	Summary: proof of the main result	107

1 Introduction

A classical problem of combinatorial optimization is finding disjoint paths with specified endpoints:

k-Vertex-Disjoint Paths Problem (*k*-DPP)
Input: A graph G and k pairs of vertices $(s_1, t_1), \dots, (s_k, t_k)$.
Question: Do there exist k pairwise vertex-disjoint paths P_1, \dots, P_k such that P_i goes from s_i to t_i ?

We consider only the vertex-disjoint version of the problem in this paper; disjoint means vertex disjoint if we do not specify otherwise. If the number k of paths is part of the input, then the problem is NP-hard even on undirected planar graphs [20]. However, for every fixed k , Robertson and Seymour showed that there is a cubic-time algorithm for the problem in general undirected graphs [24]. Their proof uses the structure theory of graphs excluding a fixed minor and therefore extremely complicated. More recently, a significantly simpler, but still very complex algorithm was announced by Kawarabayashi and Wollan [17]. Obtaining polynomial running time for fixed k is significantly simpler in the special case of planar graphs [23]; see also the self-contained presentations of Reed et al. [21] or Adler et al. [1].

The problem becomes dramatically harder for directed graphs: it is NP-hard even for $k = 2$ in general directed graphs [8]. Therefore, we cannot expect an analogue of the undirected result of Robertson and Seymour [24] saying that the problem is polynomial-time solvable for fixed k . For directed planar graphs, however, Schrijver gave an algorithm with polynomial running time for fixed k :

Theorem 1.1 (Schrijver [28]). *The k -Vertex-Disjoint Paths Problem on directed planar graphs can be solved in time $n^{O(k)}$.*

The algorithm of Schrijver is based on enumerating all possible homology types of the solution and checking in polynomial time whether there is a solution for a fixed type. Therefore, the running time is mainly dominated by the number $n^{O(k)}$ of homology types. Our main result is improving the running time by removing k from the exponent of n :

Theorem 1.2. *The k -Vertex-Disjoint Paths Problem on directed planar graphs can be solved in time $2^{2^{O(k^2)}} \cdot n^{O(1)}$.*

In other words, we show that the k -Disjoint Paths Problem is fixed-parameter tractable on directed planar graphs. The fixed-parameter tractability of this problem was asked as an open question by Bodlaender, Fellows, and Hallett [3] already in 1994, in one of the earliest papers on parameterized complexity. The question was reiterated in the open problem list of the classical monograph of Downey and Fellows [7] in 1999. Note that, for undirected planar graphs, the algorithm with best dependence on k is due to Adler et al. [1] and has running time $2^{2^{O(k)}} \cdot n^{O(1)}$. Therefore, for the more general directed version of the problem, we cannot expect at this point a running time with better than double-exponential dependence on k .

For general undirected graphs, the algorithm of Robertson and Seymour [24] relies heavily on the structure theory of graphs excluding a fixed minor; in fact, this algorithm is one of the core achievements of the Graph Minors series. More recent results on finding subdivisions [11] or parity-constrained disjoint paths [16] also build on this framework. Even in the much simpler planar case, the algorithm presented by Adler et al. [1] uses the concepts and tools developed in the study of excluded minors. In a nutshell, their algorithm has three main components. First, if treewidth (a

measure that plays a crucial role in graph structure theory) is bounded, then standard algorithmic techniques can be used to solve the k -Vertex-Disjoint Paths Problem. Second, if treewidth is large, then (planar version of) the Excluded Grid Theorem [4, 12, 22, 25] implies that the graph contains a subdivision of a large wall, which further implies that there is a vertex enclosed by a large number of disjoint concentric cycles, none of them enclosing any terminals. Finally, Adler [1] et al. show that such a vertex is irrelevant, in the sense that it can be removed without changing the answer to the problem. Thus by iteratively removing such irrelevant vertices, one eventually arrives to a graph of bounded treewidth.

Can we apply a similar deep and powerful theory in the directed version of the problem? There is a notion of directed treewidth [14] and an excluded grid theorem holds at least for planar graphs [13] (and more generally, for directed graphs whose underlying undirected graph excludes a fixed minor [15]). However, the other two algorithmic components are missing: it is not known how to solve the k -Vertex Disjoint Paths problem in $f(k) \cdot n^{O(1)}$ time on directed graphs having bounded directed treewidth and the directed grids excluded by these theorems do not seem to be suitable for excluding irrelevant vertices. There are other notions that try to generalize treewidth to directed graphs, but the algorithmic applications are typically quite limited [2, 9, 10, 18, 19, 26]. In particular, the k -Vertex-Disjoint Paths Problem is known to be W[1]-hard on directed acyclic graphs [29], which is strong evidence that any directed graph measure that is small on acyclic graphs is not likely to be of help.

Our algorithm does not use any tool from the structure theory of undirected graphs, or any notion of treewidth for directed graphs. The only previous results that we use are the results of Ding, Schrijver, and Seymour [5, 6] on various special cases of the directed disjoint paths problem, the cohomology feasibility algorithm of Schrijver [28], and a self-contained combinatorial argument from Adler et al. [1]. Therefore, we have to develop our own tools and in particular a new type of decomposition suitable for the problem. A concept that appears over and over again in this paper is the notion of alternation: we are dealing with sequences of paths and cycles having alternating orientation (i.e., each one has an orientation that is the opposite of the next one), we measure the “width” of a sequence of arcs by the number of alternations in the sequence, and we measure “distance” between faces by the minimum alternation on any sequence of arcs between them. Section 2 gives a high-level overview of the algorithm. Let us highlight here the most important steps and the main ideas:

- **Irrelevant vertices.** Analogously to Adler et al. [1], we prove that a vertex enclosed by a large set of concentric cycles having *alternating orientation* and not enclosing any terminals is irrelevant. As expected, the proof is more complicated and technical than in the undirected case (Section 3).
- **Duality of alternation.** We show that alternation has properties that are similar to the classical properties of undirected planar graphs (Section 4). We prove an approximate duality between alternating paths and the minimum alternation size of a cut (reminiscent of max-flow min-cut duality), and between concentric cycles and alternation distance (reminiscent of the fact that two faces far away in a planar graph are separated by many disjoint cycles).
- **Decomposition.** We present a novel kind of decomposition into “disc” and “ring” components (Section 5). The crucial property of the decomposition is that the set of arcs leaving a component has bounded alternation. That is, the components are connected by a bounded number of *bundles*, each containing a set of “parallel” arcs with the same orientation.
- **Handling ring components.** Ring components pose a particular challenge: we have to understand how many turns a path of the solution does when connecting the inside and the

outside. We prove a rerouting argument showing that only a bounded number of possibilities has to be taken into account for the winding numbers of these paths.

- **Guessing bundle words.** Given a decomposition, a path of the solution can be described by a word consisting of a sequence of symbols representing the bundles visited by path, in the order they appear in the path. Note that a bundle can be used several times by a path of the solution, thus the word can be very long. Our goal is to enumerate a bounded number of possible bundle words for each path of the solution. These words, together with our understanding of what is going on inside the rings, allow us to guess the homology type of the solution, and then invoke Schrijver’s cohomology feasibility algorithm to check if there is a solution with this homology type.

In the next section we present an informal overview of the algorithm; all formal arguments follow in the rest of the paper.

The techniques introduced in this paper were developed specifically with the k -Vertex-Disjoint Paths Problem in mind. It is likely that some of the duality arguments or decomposition techniques can have applications for other problems involving planar directed graphs.

In general directed graphs, vertex-disjoint and edge-disjoint versions of the disjoint paths problems are equivalent: one can reduce the problems to each other by simple local transformations (e.g., splitting a vertex into an in-vertex and an out-vertex). However, such local transformations do not preserve planarity. Therefore, our result has no implications for the edge-disjoint version of the problem on planar directed graphs. Let us note that in planar graphs the edge-disjoint version seems very different from the vertex-disjoint version: as the paths can cross at vertices, the solution does not have a topological structure of the type that is exploited by both Schrijver’s algorithm [28] and our algorithm. The complexity of the planar edge-disjoint version for fixed k remains an open problem; it is possible that, similarly to general graphs [8], it is NP-hard even for $k = 2$.

One can define a variant of the planar edge-disjoint problem where crossings are not allowed. That is, in the noncrossing edge-disjoint version paths are allowed to share vertices, but if edge e_1 entering v is followed by e_2 , and edge f_1 entering v is followed by f_2 , then the cyclic order of these edges cannot be (e_1, f_1, e_2, f_2) or (e_1, f_2, e_2, f_1) around v . It is easy to see that this version can be reduced (in a planarity-preserving way) to the vertex disjoint version by replacing each vertex by a large bidirected grid. Therefore, our algorithm can solve the noncrossing edge-disjoint version of the k -Disjoint Paths Problem as well.

2 Overview of the algorithm

The goal of this section is to give an informal overview of our main result — the fixed-parameter algorithm for finding k disjoint paths in directed planar graphs.

2.1 Irrelevant vertex rule

Let us first recall how to solve the k -disjoint paths problem in the undirected (even non-planar) case. The algorithm of Robertson and Seymour [24] considers two cases. If the treewidth of the input graph G is bounded by a function of the parameter (k , the number of terminal pairs), then the problem can be solved by a standard dynamic programming techniques on a tree decomposition of small width of G . Otherwise, by the Excluded Grid Theorem [4, 12, 22, 25], G contains a large grid as a minor.

The idea now is to distinguish a vertex v of G , whose deletion does not change the answer of the problem; that is, there exist the required k disjoint paths in G if and only if they exist in




Figure 1: A situation where a shortcut can be made and how it can be made. There are more than 2^k horizontal segments of the paths, crossed by sufficiently many vertical chords, that, in the directed setting, are required to be of alternating orientation. Moreover, it is assumed that no other path of any path intersects the gray area, so that the paths remain pairwise disjoint after rerouting.

$G \setminus v$. Note that the disjoint paths problem can become only harder if we delete a vertex; thus, to pronounce v irrelevant, one needs to prove that any solution using the vertex v can be redirected to a similar one, omitting v .

In the case of planar graphs one may apply the following quite intuitive reasoning. Assume that G contains a large grid as a minor; as there are at most $2k$ terminals, a large part of this grid does not enclose any terminal. In such a part, a vertex v hidden deep inside the grid seems irrelevant: any solution using v needs to traverse a large part of the grid to actually contain v , and it should be possible to “shift” the paths a little bit to omit v . This reasoning can be made formal, and Adler et al. [1] proved that, in undirected planar graphs, the middle vertex of a grid of exponential (in k) size is irrelevant. In fact, they show a bit stronger statement: if we have sufficiently many (around 2^k) concentric cycles on the plane, such that the outermost cycle does not enclose any terminal, then any vertex on the innermost cycle is irrelevant.

One of the main argument in the proof of Adler et al. [1] is as follows. Assume that there are many pairwise disjoint segments of the solution that cross sufficiently many orthogonal paths (henceforth called *chords*) in the graph; see Figure 1. Assume moreover that the aforementioned segments are the only parts of the solution that appear in the area enclosed by the outermost segments and chords (i.e., in the part of the plane depicted on Figure 1). Then, if the number of segments is more than 2^k , one can redirect some of them, using the chords, and shortening the solution. Thus, in a minimal (in some carefully chosen sense) solution, a set of more than 2^k paths cannot go together for a longer period of time.

The argument of Adler et al. [1] described in the previous paragraph redirects the paths of the solution using the chords in an undirected way, and hence the direction in which a chord is used is unpredictable, depending on the order in the which the segments appear on the paths of the solution. Hence, if we want to transfer this argument to the directed setting, then we need to make some assumption on the direction of the chords. It turns out that what we need is that the chords are directed paths with alternating orientation. This ensures that we always have a chord going in the right direction at any place we would possibly need it.

If a set of paths intersect the innermost cycle, then they need to traverse all cycles. Adler et al. [1] show how to find a subset of these paths and how to cut out chords from the cycles in a way that satisfies the conditions of the rerouting argument. In the directed setting, in order




Figure 2: A d -bend B with chords C_0, C_1, \dots, C_d , and how it can be cut out from concentric cycles, using parts of the cycles as chords.

to obtain chords of alternating orientation, we need to assume that the cycles have alternating orientation too.

Definition 2.1. We say that cycles C_1, \dots, C_d form a *sequence of concentric cycles with alternating orientation* in a plane graph G if

1. they are pairwise vertex disjoint,
2. for every $1 \leq i < d$, cycle C_i encloses C_{i+1} , and
3. for every $1 \leq i < d$, exactly one of the cycles C_i and C_{i+1} is oriented clockwise.

Luckily, it turns out that such a sequence of cycles is sufficient for the irrelevant vertex rule. In Section 3 we prove the following.

Theorem 2.2 (Irrelevant vertex rule). *For any integer k , there exists $d = d(k) = 2^{O(k^2)}$ such that the following holds. Let G be an instance of k -DPP and let $C_1, C_2, C_3, \dots, C_d$ be a sequence of concentric cycles in G with alternating orientation, where C_1 is the outermost cycle. Assume moreover that C_1 does not enclose any terminal. Then any vertex of C_d is irrelevant.*

At the heart of the proof of Theorem 2.2 lies the rerouting argument described above, which states that a solution can be rerouted and shortened if a set of more than 2^k paths travel together through sufficiently many (exponential in k) chords cut out from the alternating cycles C_i . However, it is much harder to prove the existence of these paths and chords needed for the rerouting argument than in the undirected case, and we now sketch how it could be done.

Consider the situation assumed in Theorem 2.2 and assume we have a solution where one path, say P , intersects the innermost cycle. On one side of P we obtain a structure we call a *bend*, depicted on Figure 2. The parts of the cycles are called *chords*, a bend with d chords is a *d -bend*. Moreover, *the type of the bend* is the number of different paths from the solution that intersect the interior or the boundary of the bend; our initial bend is of type at most k . Our main technical claim in the proof of Theorem 2.2 is that in a (somehow defined) minimal solution there do not exist d -bends of type t , for $d > f(k, t)$ and some function $f(k, t) = 2^{O(kt)}$, that do not enclose any terminals.




Figure 3: Part of a path creates a bend inside another bend.

Assume we have a d -bend B of type t , for some large d , enclosed by a part of a path P_i of the solution. We analyze the *segments* of the solution: the maximal subpaths of the paths P_1, \dots, P_k in the interior of the bend. If the interval vertices of the last two chords are not intersected by any segment, then one of these two chords has the right orientation to serve as a shortcut for the path P_i , contradicting the minimality of the solution. Therefore, we can assume that all but the last two chords are intersected by segments. If any segment of $P_{i'}$ intersects the j -th chord of B , then it itself induces a j' -bend B' inside B , for some $j' = j - O(1)$ (see Figure 3). Hence, if the path P_i itself does not intersect the interior of the d -bend B , any bend inside B is of strictly smaller type, and the claim is proven by induction on t .

Otherwise, we can argue that several segments of P_i enter the interior of the d -bend B . Our goal is to prove that there is a large set of segments of P_i entering B that form a nested sequence and they travel together through a large number of chords deep inside the bend, with no other segment of P_i between them. Then we can argue that any other segment of some $P_{i'}$ with $i' \neq i$ intersecting these chords is also nested with these segments, otherwise they would create a large bend of strictly smaller type, and induction could be applied. Therefore, we get a large set of paths travelling together through a large number of chords, and the rerouting argument described above can be invoked.

To find this nested sequence of segments of P_i , we analyze how P_i intersects the chords of B . We construct the following auxiliary graph H : start with a subgraph of G consisting of P and the chords of B and suppress all vertices of degree 2. Let H^* be the dual of H and T^* be subgraph of H^* consisting only of chord arcs. It is not hard to see that T^* is a tree, with at least $d - O(1)$ vertices; see Figure 4 for an illustration.

Roughly speaking, if the segments of P_i reaching deep inside the bend are not nested, then we can find a face f , hidden deeply inside the d -bend B , such that at least three of the connected components of $T^* \setminus f$ cross many chords. If one of these components has the property that none of the faces appearing in this component contains a terminal, then the part of the path P_i that encloses this component is, by the definition of T^* , encloses a bend of type at most $t - 1$. Thus, by the induction hypothesis, it cannot cross more than $f(k, t - 1)$ chords of the d -bend B . However, if all such connected components of $T^* \setminus f$ contain faces with terminals inside, we cannot argue anything about f : the path P may need to do travel in such strange manner in order to go around




Figure 4: A d -bend with chords C_0, C_1, \dots, C_d appearing on a path P_b . The dotted lines show the edges of the tree T^* .

some terminals. The crucial observation is that there are $O(k)$ faces for which this situation can arise: as there are $2k$ terminals, there are only $O(k)$ vertices of the tree T^* such that at least three components of $T^* \setminus f$ contain faces with terminals. Therefore, if we avoid these $O(k)$ special faces, then we can find the required set of nested segments and we can find a place to apply the rerouting argument. This finishes the sketch of the proof of the irrelevant vertex rule (Theorem 2.2).

We would like to note that we can test in polynomial time if the irrelevant vertex rule applies: if we guess one faces enclosed by C_r and the orientation of C_r , we can construct the cycles in a greedy manner, packing the next cycle as close as possible to the previously constructed one. However, we do not use this property in our algorithm: the decomposition algorithm, described in the next subsection, returns an irrelevant vertex situation if it fails to produce a suitable decomposition.

We would also like to compare the assumptions of Theorem 2.2 with the conjectured canonical obstruction for small directed treewidth, depicted on Figure 5. It has been shown that a planar graph [13], or, more generally, a graph excluding a fixed undirected minor [15], has small directed treewidth unless it contains a large directed grid (as in Figure 5), in some minor-like fashion, and this statement is conjectured to be true for general graphs [14]. Although the assumption of bounded directed treewidth may be easier to use than the bounded-alternation decomposition presented in the next subsection, we do not know how to argue about irrelevancy of some vertex or arc in the directed grid. Thus, we need to stick with our irrelevant vertex rule with relatively strong assumptions (a large number of alternating cycles), and see in the rest of the proof what can be deduced if such a situation does not occur.

2.2 Decomposition and duality theorems

Once we have proven the irrelevant vertex rule (Theorem 2.2), we may see what can be deduced about the structure of the graph if the irrelevant vertex rule does not apply. Recall that in the




Figure 5: A directed grid — a conjectured canonical obstacle for small directed treewidth.




Figure 6: An illustration of the definition of $S(\gamma, p)$ for $p \in \gamma \cap G$.

undirected case the absence of an irrelevant vertex implied a bound on the treewidth of the graph, and hence the problem can be solved by a standard dynamic programming algorithm.

In our case the situation is significantly different. As we shall see, the assumptions in Theorem 2.2 are rather strong, and, if the irrelevant vertex rule is not applicable, the problem does not become as easy as in the bounded-treewidth case. Recall that Theorem 2.2 assumed a large number of cycles of alternating orientation, and these alternations were crucial for the rerouting argument. It turns out that, if such cycles cannot be found, we can decompose the graph into relatively simple pieces using cuts of bounded alternation.

Consider a directed curve γ on the plane that intersects the plane graph G only in a finite number of points (i.e., γ does not “slide” along any arc of G). For any point $p \in \gamma \cap G$ we define $S(\gamma, p) \subseteq \{-1, +1\}$ as follows: $-1 \in S(\gamma, p)$ if it is possible for a path in G to cross γ in p from left to right, and $+1 \in S(\gamma, p)$ if it is possible to cross γ from right to left (see Figure 6). The *alternation* of γ is the length of the longest sequence of alternating $+1$ and -1 s that is embeddable (in a natural way) into the sequence $S(\gamma, p)_{p \in \gamma \cap G}$.

Note that the existence of a curve γ with alternation \mathfrak{a} connecting faces f_1 and f_2 proves that f_1 and f_2 cannot be separated by a sequence of more than \mathfrak{a} concentric cycles of alternating orientation. Thus, a curve of bounded alternation is in some sense dual to the notion of concentric cycles of bounded alternation. It turns out that this duality is tight: such a curve of bounded alternation is the only obstacle that prevents the existence of these concentric cycles. One can also formulate a duality statement similar to the classical max-flow min-cut duality, with a set of paths of alternating orientation playing the role of the flow and a curve of bounded alternation playing the role of a cut. The following lemma states both types of duality in an informal way (see Figures




Figure 7: Two cases in Lemma 2.3(1): cycles of alternating orientation between f_{in} and f_{out} or a curve of bounded alternation connecting f_{in} and f_{out} .

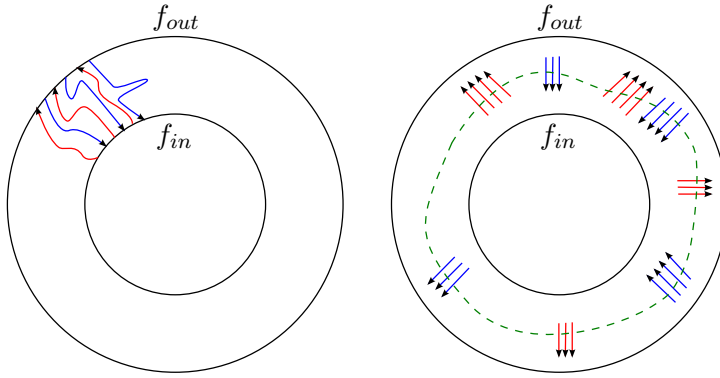


Figure 8: Two cases in Lemma 2.3(2): paths of alternating orientation connecting f_{in} and f_{out} or a curve of bounded alternation separating f_{in} and f_{out} .

7 and 8 for illustration).

Lemma 2.3 (Alternation dualities, informal statement.). *Let G be a graph embedded in a subset of a plane homeomorphic to a ring, and let f_{in} and f_{out} be the two faces of G that contain the inside and the outside of the ring, respectively. Let r be an even integer. Then, in polynomial time, one can in G :*

1. *either find a sequence of r cycles of alternating orientation, separating f_{in} from f_{out} , or find a curve connecting f_{in} with f_{out} with alternation at most r (Figure 7); and*
2. *either find a sequence of r paths, connecting f_{in} and f_{out} , with alternating orientation, or find a closed curve separating f_{in} from f_{out} with alternation at most $r + 4$ (Figure 8).*

Let us now give intuition on how to prove statements like Lemma 2.3. If we identify f_{in} and f_{out} , or more intuitively, extend the surface with a handle connecting f_{in} and f_{out} , we can perceive G as a graph on a torus. After some gadgeteering, we may use the following result of Ding, Schrijver, and Seymour [6]: if one wants (in a graph G on a torus) to route a set of vertex-disjoint cycles with prescribed homotopy class and directions, a canonical obstacle is a face-vertex curve γ (of some other homotopy class), where the sequence $S(\gamma, p)_{p \in \gamma \cap G}$ does not contain the expected subpattern of $+1$ and -1 s. Note that such a curve is not far from the curves promised by Lemma 2.3.

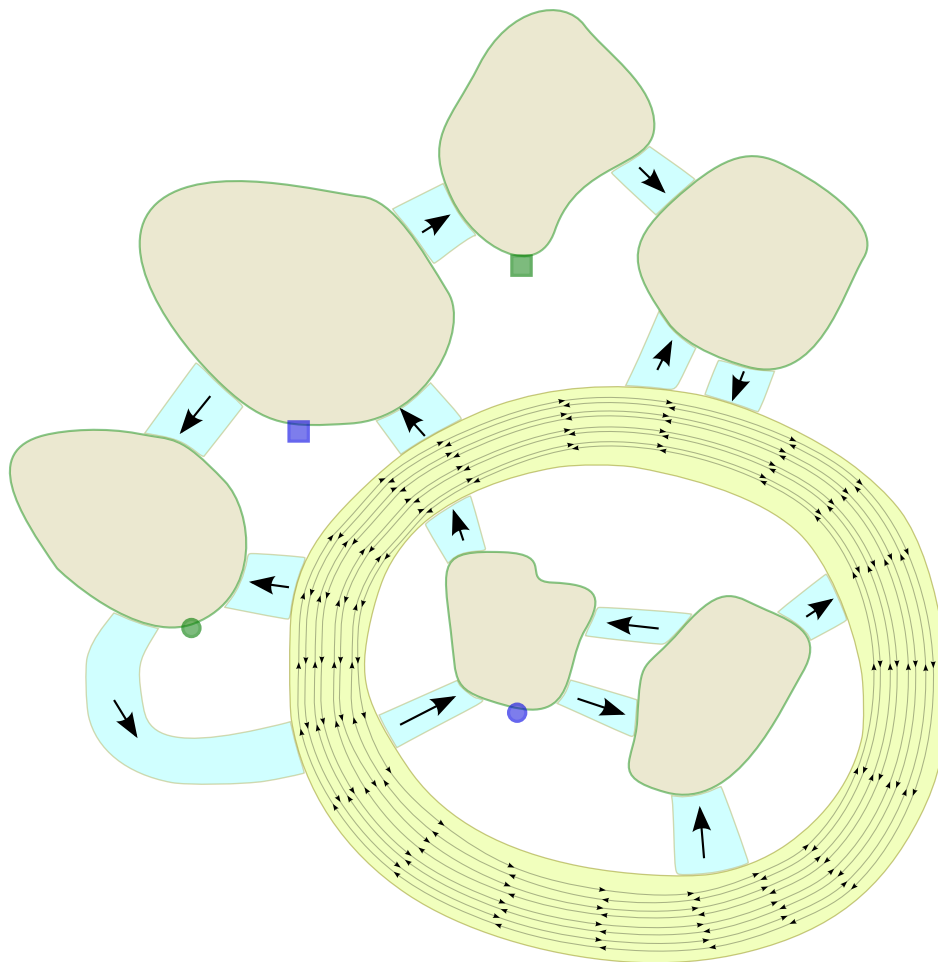


Figure 9: An example of a decomposition. The disc components are red, the ring component is yellow.

Equipped with this understanding of alternation, in Section 5 we prove a decomposition theorem that is crucial for our algorithm. We state this theorem here informally (see Figure 9 for an illustration); the precise statement appears in Theorem 5.8.

Theorem 2.4 (Decomposition theorem, informal statement). *Assume that G is a plane graph with k terminal pairs to which the irrelevant vertex rule is not applicable. Then one can partition the graph G into a bounded (in k) number of disc and ring components, using cuts of bounded total alternation; a disc (resp., ring) component occupies a subset of the plane that is isomorphic to a disc (resp., ring). Moreover, each terminal lives on the border of a disc component, and each ring component contains many concentric cycles of alternating orientation, separating the inside from the outside.*

The decomposition of Theorem 2.4 is obtained by iteratively refining a decomposition, moving a terminal to the boundary of a component in each step. If a disc component contains a terminal such that there is a curve of bounded alternation from the terminal to the boundary of the component, then the terminal can be moved to the boundary by removing the arcs intersected by the curve. This operation increases the alternation of the cut separating the component from the rest of graph only

by a bounded number, thus we can afford to perform one such step for each terminal. Otherwise, if there is no such curve, then Lemma 2.3(1) implies that there is a large set of concentric cycles of alternating orientations separating all the terminals in the component from the boundary. We again consider two cases. If there is large set of paths with alternating orientations crossing these cycles, then the paths and cycles together form some kind of grid, and we can easily identify a vertex that is separated from all the terminals by a large set of concentric cycles with alternating orientation. Such a vertex is irrelevant by Theorem 2.2, and hence can be removed. On the other hand, if there is no such set of paths, then Lemma 2.3(2) implies that there is a curve of bounded alternation separating the terminals of the component from the boundary of the component. We can use this curve to cut away a ring component and we can do this in such a way that after removing the ring component, one of the terminals is close to the boundary of the remaining part of the disc component (in the sense that there is a curve of bounded alternation connecting it to the boundary). Therefore, we can apply the argument described above to move this terminal to the boundary. Iteratively applying these steps until all the terminals are on the boundary of its component produces the required decomposition.

How can we use the decomposition of Theorem 2.4 to solve k -DPP? The disc components are promising to work with, as the k -DPP problem is polynomial if all terminals lie on the outer face of the graph [5]. In a topological sense, if the terminals are on the outer face, then the solutions are equivalent, whereas if the terminals are on the inside and outside boundaries of a ring, then the solutions can differ in how many “turns” they do along the ring. This possible difference in the number of turns create particular challenges when we are trying to apply the techniques of Schrijver [28] to find a solution based on a fixed homotopy class.

Theorem 2.4 would be nicer and more powerful if we could always obtain a decomposition using only disc components, but as Figure 10 indicates, this does not seem to be possible in general. Assume that one part of the graph (inside) is separated from the outside by a large number of concentric cycles of alternating orientation, but with cuts of bounded alternation between each consecutive cycles. Suppose that there are terminals inside the innermost cycle and outside the outermost cycle. Then the irrelevant vertex rule of Theorem 2.2 is not applicable, as we cannot find suitable set of cycles without enclosing some terminals inside the innermost cycle. Moreover, if we aim for bounded alternation cuts, we cannot cut through too many cycles. Thus, this set of concentric cycles need to be embedded in a separate ring component.

The formal statements of Lemma 2.3 are proven in Section 4, whereas the decomposition theorem is proven in Section 5.

2.3 Bundles and bundle words

From the previous subsection we know that, if the irrelevant vertex rule is not applicable, one may decompose the graph into a bounded number of disc and ring components, using cuts of bounded alternation. Let us reformulate this statement: we can decompose the graph into a bounded number of disc and ring components, connected by a bounded number of *bundles*; a *bundle* is a set of arcs of G that form a directed path in the dual of G , such that no other arc nor vertex of G is drawn between the consecutive arcs of the bundle. Thus, we obtain something we call *bundled instance* $(G, \mathcal{D}, \mathbb{B})$: a graph G with terminals, a family of components \mathcal{D} and a family of bundles \mathbb{B} . On Figure 9 one can see a partition of arcs between components into bundles. With any path P in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ we can associate its *bundle word*, denoted $\mathbf{bw}(P)$: we follow the path P from start to end and append a symbol $B \in \mathbb{B}$ whenever we traverse an arc belonging to a bundle B . That is, $\mathbf{bw}(P)$ is a word over alphabet \mathbb{B} ; see Figure 11 for an example.

Assume for a moment that there are no ring components in the decomposition; ring components




Figure 10: A situation, where a ring component is necessary: there are many concentric cycles of alternating orientation, but the arcs between the cycles have bounded alternation.

present their own challenges requiring an additional layer of technical work, but they do not alter the main line of reasoning. Assume moreover that we have computed somehow bundle words $\mathbf{bw}(P_i)$, $i = 1, 2, \dots, k$ for some solution $(P_i)_{i=1}^k$ for k -DPP on G . The important observation is that, given the bundle words, the cohomology feasibility algorithm of Schrijver [28] is able to extract (an approximation of) the paths P_i in polynomial time.

To show this, let us recall the algorithm of Schrijver [28] that solves k -DPP in polynomial time for every fixed k . The heart of the result of Schrijver lies in the proof that k -DPP is polynomial-time solvable if we are given a homotopy class of the solution. In simpler words, given a “pre-solution”, where many paths can traverse the same arc, even in wrong direction (but they cannot cross), one can in polynomial time check if the paths can be “shifted” (i.e., modified by a homotopy) so that they become a feasible solution. In such a “shift” (homotopy) one can move a path over a face, but not over a face that contains a terminal.¹ See Figure 12 for an illustration of different homotopy classes of a solution.

In our setting, we note that, in the absence of ring components, two solutions with the same set of bundle words of each paths are homotopical; thus, given bundle words of a solution, one can use the algorithm of Schrijver to verify if there is a solution consistent with the given set of bundle words. However, one should note that the homotopies are allowed to do much more than to only move paths within a bundle; formally, using the Schrijver’s algorithm we can either (i) correctly conclude that there is no solution with given set of bundle words $(p_i)_{i=1}^k$, or (ii) compute a solution $(P_i)_{i=1}^k$ such that the bundles of $\mathbf{bw}(P_i)$ is a subset (as a multiset) of the bundles of p_i .

Unfortunately, if the decomposition contains ring components as well, then the bundle words of a solution does not describe the homotopy class of the solution. What do we need to learn, apart from bundle words of the solution, to identify the homotopy class of the solution if ring components are present? The answer is not very hard to see: for any subpath of a path in a solution that crosses

¹Note that we can assume that each terminal is of degree one, and the notion of “face containing a terminal” is well-defined.




Figure 11: An example of a solution. The green path has bundle word $ANMDEFHDE$ and the blue path has bundle word $KLID$.




Figure 12: Different homotopy classes of a solution: in the first two figures, the solutions are of the same class, whereas on the third figure the homotopy class is different.




Figure 13: Winding numbers inside a ring component; formally defined as the signed number of crosses of some fixed chosen reference curve. The green line has winding number 0, the blue $+2$ and the red -1 .

some ring component (i.e., goes from the inside to the outside of vice versa) we need to know how many times it “turns” inside the ring component; we call it a *winding number* inside a component.²

Thus, our goal is to compute a small family of possible bundle words and winding numbers such that, if there exists a solution to k -DPP on G , there exists a solution consistent with one of the elements of the family. In fact, our main goal in the rest of the proof is to show that one can compute such family of size bounded in the parameter k .

2.4 Guessing bundle words

Assume again that there are no ring components; we are to guess the bundle words of one of the solutions. Recall that the number of bundles, $|\mathbb{B}|$, is bounded in k . Thus, if a bundle word of some path P_i from a solution $(P_i)_{i=1}^k$ is long, it needs to contain many repetitions of the same bundles.

Let us look at one such repetition: let uB be a subword of $\mathbf{bw}(P_i)$, where B is the first symbol of u and u contains each symbol of \mathbb{B} at most once. We call such place a *spiral*. Note that this spiral separates the graph into two parts, the inside and the outside, where any other path can cross the spiral only in a narrow place inside bundle B (see Figure 14). As the arcs of B go in one direction, any path P_j , $j \neq i$ can cross the spiral only once, in the same direction as P_i , and the path P_i cannot cross the spiral uB again. Note that we know exactly which paths cross the spiral uB : the paths that have terminals on different sides of the spiral uB .

There is also one important corollary of this observation on spirals. If $\mathbf{bw}(P_i) = xuBy$ for some words x, y and spiral uB , then, for any bundle B' that does not appear in u , only one of the words x or y may contain B' : the bundle B' is either contained inside the spiral uB or outside it. By some quite simple word combinatorics, we infer that $\mathbf{bw}(P_i)$ can be decomposed as $u_1^{r_1} u_2^{r_2} u_3^{r_3} \dots u_s^{r_s}$, where each word u_i contains each symbol of \mathbb{B} at most once, each r_i is an integer and $s \leq 2|\mathbb{B}|$.

²Formally, for a part of a path that starts and ends on the same side of the ring component we need to know also on which side it leaves the other side of the ring component; however, it turns out that this knowledge is quite easy to guess or deduce and we ignore this issue in the overview.

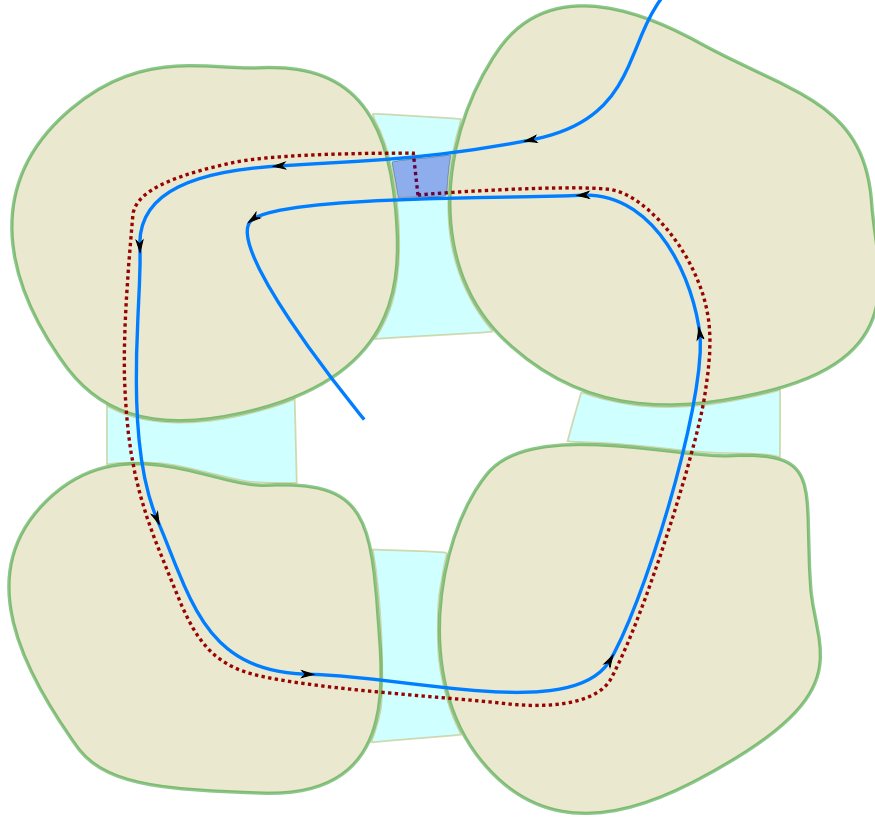


Figure 14: A spiral. Any other path may cross the dashed curve only in a narrow place in the top bundle (highlighted with darker blue colour).

Note that if we aim to guess bundle words of some solution $(P_i)_{i=1}^k$, there is only a bounded number of choices for the length s and the words u_i ; the difficult part is to guess the exponents r_i , if they turn out to be big (unbounded in k). That is, we can easily guess the global structure of the spirals (how they are nested, etc.), but we cannot easily guess the “width” of the spirals (how many turns of the same type they do). We need further analysis and insight in order to be able to guess these numbers as well.

Let us focus on a place in a graph where a path P_i in the solution $(P_i)_{i=1}^k$ contains a subword $u^r B$ of $\mathbf{bw}(P_i)$ for some large integer r , where B is the first symbol of u . The situation, depicted on Figure 15, looks like a large spiral; the *spiraling ring* is the area between the first and last spiral uB in the subword $u^r B$. Note that any path P_j that enters this area, actually needs to traverse all r spirals uB and $\mathbf{bw}(P_j)$ contains a subword $u^{r-1} B$; let $I \subseteq \{1, 2, \dots, k\}$ be the set of indices j such that P_j traverses uB . Moreover, note that, since B contains arcs going in only one direction, these parts of paths $(P_j)_{j \in I}$ are the only intersections of the solution $(P_i)_{i=1}^k$ with the spiraling ring.

Our main claim is that we can choose r to be as small as possible, just to be able to route the desired paths through the spiraling ring in G . More formally, we prove that if we can route $|I|$ directed paths through the spiraling ring, such that each path traverses B roughly r^* times (but they may start and end in different places than the parts of the solution $(P_i)_{i=1}^k$ traversing the spiraling ring), then we can modify the solution $(P_i)_{i=1}^k$ inside the spiraling ring such that $r \leq r^* + O(1)$. If we choose r^* to be minimal possible, we have $|r - r^*| = O(1)$ and we can guess r .

To prove that the paths can be rerouted, we show the following in Section 4.




Figure 15: A spiraling ring; the dashed lines are its borders.

Theorem 2.5 (rerouting in a ring, informal statement). *Let G be a plane graph embedded in a ring, with outer face f_{out} and inner face f_{in} . Assume that there exist two sequences of vertex-disjoint paths $(P_i)_{i=1}^s$ and $(Q_i)_{i=1}^s$ connecting f_{in} with f_{out} , such that P_i goes in the same direction (from f_{in} to f_{out} or vice versa) as Q_i , and the endpoints of $(P_i)_{i=1}^s$ lie in the same order on f_{in} as the endpoints of $(Q_i)_{i=1}^s$. Then one can reroute $(P_i)_{i=1}^s$ inside G , keeping the endpoints, such that the winding number of P_1 differs from the winding number of Q_1 by no more than 6.*

How to prove such a rerouting statement? We again make use of the results of Ding et al. [6] on a canonical obstacle for routing a set of directed cycles on a torus (as in the proof of Lemma 2.3). We connect f_{in} and f_{out} with a handle, and perceive the paths P_i and Q_i as a cycles on a torus. The winding number w_P of P_1 determines the homotopy class of the cycles $(P_i)_{i=1}^s$, and the winding number w_Q of Q_1 determines the homotopy class of the cycles $(Q_i)_{i=1}^s$. Now we observe that an obstacle for routing the same set of cycles with “homotopy” w for $w_Q + O(1) < w < w_P - O(1)$ (or, symmetrically, $w_P + O(1) < w < w_Q - O(1)$) would be an obstacle for “homotopy” either w_Q or w_P , a contradiction. Hence, almost all “homotopies” between w_Q and w_P are realizable. By some gadgeteering, we may force the cycles to use the same endpoints as the paths $(P_i)_{i=1}^s$, at the cost of $O(1)$ loss in the “homotopy” class.

We would like to note that Theorem 2.5 is a cornerstone of our result. The exponential time complexity of the algorithm of Schrijver [28] comes from the fact that the number of homotopy classes of a solution solution cannot be bounded by a function of k , because the number of possibilities for the number of turns of the solution in some ring-like parts of the graph cannot be bounded by a function of k . Theorem 2.5 overcomes this obstacle by showing that that for each

such ring, one can choose a canonical number of turns (that depends only on the ring, not how it is connected to the outside) and the solution can be assumed to spiral a similar number of turns than the canonical pass. In other words, if one would try to construct a $W[1]$ -hardness proof of k -DPP by a reduction from, say, k -CLIQUE, one cannot expect to obtain a gadget that encodes a choice of a vertex or edge of the clique by a number of turns a solution path makes in some ring-like part of the graph — and such an encoding seems natural, taking into account the source of the exponential-time complexity of the algorithm of Schrijver [28].

However, it still requires significant work to make use of Theorem 2.5. In the case of spiraling rings, the question that remains is how to get minimal exponent r^* such that $|I|$ paths can be routed through a spiraling ring with r^* turns. The idea is to isolate a part of the graph and parts of the bundle words of the solution where only one exponent is unknown, and then apply Schrijver’s algorithm for different choices of exponent; the desired value r^* is the smallest exponent for which Schrijver’s algorithm returns a solution.

However, in this approach we need to cope with two difficulties. First, the Schrijver’s algorithm may find a solution that follows our guidelines (bundle words) in a very relaxed manner. However, as at each step we choose the exponent r to be very close to the minimum possible number of turns in a spiraling ring, the bundle words of the solution found by the algorithm cannot differ much from the given ones, as they need to spiral at least the number of times we have asked them to (up to an additive constant).

Second, it is not always easy to isolate a part of the graph where one spiraling occurs. A natural thing to do is to cut the graph along bundles not used in the spiral and attach auxiliary terminals; however, in a situation on Figure 16 we cannot separate the middle spiraling ring from the two shorter ones outside and inside it. Luckily, it turns out that here the additional spirals use always *strictly smaller* number of bundles, and we can guess exponents r in terms u^r in the increasing order of $|u|$.

2.5 Handling a ring component

In the previous subsection we have sketched how to guess bundle words of the solution in absence of ring components. Recall that, for a ring component, and for any part of a path that traverses a ring component (henceforth called *ring passage*; note that ring passages are visible in bundle words of paths) we need to know its *winding number*: the number of times it turns inside the ring component.

As we have learnt already how to route paths in rings, it is tempting to use the aforementioned techniques to guess winding numbers: guess how many ring passages there are, and find one winding number w^* for which routing is possible (using Schrijver’s algorithm)³; the actual solution should be reroutable to a winding number w close to w^* .

However, there are two major problems with this approach. First, not only the ring passages of the solution use the arcs and vertices of a ring component, but parts of paths from the solution that start and end on the same side of the ring component (henceforth called *ring visitors*) may also be present. Luckily, we may assume that the ring components contain many concentric cycles of alternating orientation, as otherwise the decomposition algorithm would cut it though to obtain a disc component. If a ring visitor goes too deeply into the ring component, it creates a d -bend for too large d and we can reroute it. Thus, the ring visitors use only a thin layer of the ring component, and we can argue that we can still conduct the rerouting argument in the ring component without bigger loss on the bound on $|w - w^*|$; see Figure 17 for an illustration.

³Note that all winding numbers of passages in one ring component do not differ by more than one, and in this overview we assume they are equal.

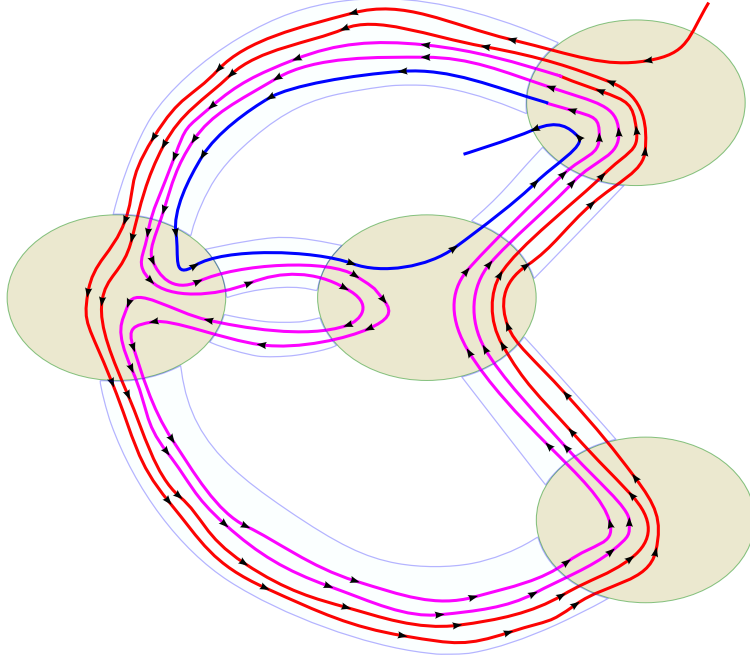


Figure 16: A nontrivial situation around a spiral on P_i . Path P_i first spirals two times using four bundles (red spiral), then makes the spiral we want to investigate by spiraling two times using six bundles (pink spiral), and finally makes one turn of a spiral using three bundles (blue spiral). Note that the set of bundles used in the red spiral and in the blue spiral is a subset of the set of bundles used in the pink spiral, which we aim to measure. The crucial observation is that these sets of bundles must be in fact **proper** subsets of the set used by the pink spiral.

Second, we do not have yet any means to control the number of ring passages, and the previous techniques of guessing bundle words have significant technical problems if we try to handle spiraling rings involving ring components. Hence, it is not trivial to guess the set of ring passages traversing a ring component. Here again we may use the concentric cycles hidden inside a ring component, as well as bounded alternation cuts that can be found repeatedly inside the ring component if the irrelevant vertex rule is not applicable. We argue that, if we have many ring passages, a large number of them need to travel together via a large number of concentric cycles of alternating orientation and we can use a rerouting argument in the spirit of the one used by Adler et al. [1].⁴ Overall, we obtain that there exists a solution with a bounded number of ring passages, and we are able to guess a good candidate for a winding number inside a ring component. To merge the techniques of the previous and this subsection, we need to handle ring visitors when guessing bundle words: these visitors may take part in some large spiraling ring. Luckily, as ring visitors are shallow in ring components, we can “peel” the ring components: separate a thin layer that may contain ring visitors, cut it through and treat it as disc component for the sake of bundle word guessing algorithm.

The analysis of bundle words, ring visitors and ring passages is done in Section 6. The Schrijver’s algorithm is recalled in Section 7. The final guessing arguments — both for bundle words and winding numbers — are described in Section 8.

⁴It is worth noticing that the bound on directions make it possible to use a simple flow argument instead of the techniques of Adler et al. [1].




Figure 17: Guessing winding number in a (flattened) ring component. The green pass is the actual solution we seek for, and the blue pass is the path we found. The red parts are used by ring visitors in the solution, thus we are not able to use the rerouting argument to the whole ring component. However, we are able to use it for the ring component with the red parts removed; we then argue, basing on the shallowness of ring visitors, that the winding numbers of the entire blue pass and the subpath of the blue pass that connects the red parts on the top with the red parts on the bottom do not differ much.

2.6 Summary

We conclude with a summary of the structure of the algorithm.

First, we invoke decomposition algorithm of Theorem 2.4. If it fails, it exhibits a place where the irrelevant vertex rule is applicable; apply the rule and restart the algorithm. Otherwise, compute bundled instance $(G, \mathcal{D}, \mathbb{B})$, with $|\mathcal{D}|$ and $|\mathbb{B}|$ bounded in k .

Given the bundled instance $(G, \mathcal{D}, \mathbb{B})$, we aim to branch into subcases whose number is bounded by a function of k , in each subcase guessing a set of bundle words for the solution, as well as winding numbers of each ring passage. Our branching will be exhaustive in the following sense: if G is a YES-instance to k -DPP, there will be a guess consistent with some solution (but not *all* solutions will have their consistent branches).

We branch in two phases. First, we guess the bundle words; the hard part is to guess exponents in spiraling rings u^r , where we argue that we can choose an exponent close to minimal possible number of turns in a spiraling ring, and detect this number using an application of Schrijver's algorithm for a carefully chosen subgraph of G . Second, we guess the winding numbers; here we argue that the winding numbers of the solution can be assumed to be close to a winding number of an arbitrarily chosen way to route ring passages through the ring component, ignoring the ring visitors.

Finally, given bundle words and winding numbers, we deduce the homotopy type of the solution and invoke Schrijver's algorithm on the entire graph to verify whether our guess is a correct one.

3 Irrelevant vertex rule

We prove the validity of the irrelevant vertex rule in this section.

Definition 3.1. Consider an instance of the k -path problem with graph G and terminals T . Let v be a nonterminal vertex. We say that v is *irrelevant* if the instance on G has a solution if and only if the instance on $G \setminus v$ has.

The main result of the section is the following:

Theorem 3.2 (Irrelevant Vertex Rule). *Consider an instance of the k -path problem with a graph G embedded in the plane and a set T of terminals. There is a $d := d(k) = 2^{O(k^2)}$ such that the following holds. Let C_0, \dots, C_d be an alternating sequence of concentric cycles (C_0 is outside) such that there is no terminal enclosed by C_0 . If v is a vertex of C_d , then v is irrelevant.*

We prove Theorem 3.2 by formulating a statement about unique solutions.

Definition 3.3. Let G be a graph with a set T of terminals. We say that a solution P_1, \dots, P_k is *unique* if for every solution P'_1, \dots, P'_k we have $P_i = P'_i$ for every i . Given a solution P_1, \dots, P_k , we say that an arc is *free* if it is not used by any P_i .

The main technical statement that we prove is that a unique solution cannot enter a large set of concentric free cycles:

Lemma 3.4. *There is a function $d(k) = 2^{O(k^2)}$ such that if an instance has a unique solution such that there is an alternating sequence of $d(k)$ free concentric cycles not enclosing any terminals, then the solution does not intersect the innermost cycle.*

The setting of Lemma 3.4 somewhat simplifies the proof, as we have to argue about a solution that is disjoint from the concentric cycles. Theorem 3.2 follows from Lemma 3.4 by the combination of a minimal choice argument and contracting arcs that are shared by the solution and the cycles.

Proof (or Theorem 3.2). Let E_C be the union of the arc sets of every cycle C_i . If the instance has no solution, then v is irrelevant by definition. Otherwise, let P_1, \dots, P_k be a solution using the minimum number of arcs *not* in E_C ; let E_P be union of the arc sets of every path P_i . We may assume that this solution uses vertex v , otherwise we are done. We create a new instance and a corresponding solution the following way. First, let us remove every arc not in $E_P \cup E_C$. Next, we contract every arc of $E_P \cap E_C$. Note that it is not possible that the solution uses every arc of a cycle C_j , thus each cycle C_j is contracted into a cycle C'_j . (The length of the cycle C'_j might be 1 or 2; in order to avoid dealing with loops and parallel arcs, we may subdivide such arcs without changing the problem.) Let G' be the graph obtained this way. Clearly, each path P_i becomes a (possibly shorter) path P'_i with the same endpoint and the paths remain disjoint, hence P'_1, \dots, P'_k is a solution of the new instance that is disjoint from the cycles C'_1, \dots, C'_d . The arcs of $E_C \setminus E_P$ are free, and hence the cycles C'_1, \dots, C'_d are free.

As one of the paths P_i used vertex v of C_d , there is a path P'_i intersecting C'_d . Therefore, by Lemma 3.4, P'_1, \dots, P'_k is not the unique solution. Let Q'_1, \dots, Q'_k be a solution different from P'_1, \dots, P'_k . As the arcs of $E_P \setminus E_C$ form disjoint paths connecting the terminals, there has to be an arc $e^* \in E_P \setminus E_C$ not used by any Q'_i . We construct a solution Q_1, \dots, Q_k of the original instance from Q'_1, \dots, Q'_k by uncontracting the arcs contracted during the construction of the new instance. We have to verify that each Q_i is a directed path. In general, if we contract an arc \vec{xy} into a single vertex v_{xy} in a directed graph, then a path Q going through v_{xy} in the contracted graph might not

correspond to any directed path in the original graph: it is possible that Q enters v_{xy} on an arc that corresponds to an arc entering y and it leaves v_{xy} on an arc that leaves x , hence replacing v_{xy} with the arc \overrightarrow{xy} in Q does not result in a directed path. However, in our case, when we contracted an arc \overrightarrow{xy} , then \overrightarrow{xy} is the only arc leaving x and it is the only arc entering y : this follows from the fact that we remove every arc not on any path P_i or on any cycle C_j . Therefore, the paths Q_1, \dots, Q_k obtained by reversing the contractions form a solution of the original instance. Let us observe that, as no Q'_i uses the arc $e^* \in E_P \setminus E_C$ in G' , no Q_i uses the corresponding arc in G , which is an arc not in E_C . Moreover, if some Q_i uses an arc that is not in E_C , then this arc is in E_P and hence used by some $P_{i'}$ (otherwise we would have removed it in the construction of G'). Thus Q_1, \dots, Q_k is a solution that uses strictly fewer arcs not in E_C than P_1, \dots, P_k , which contradicts the minimality of the choice of P_1, \dots, P_k . \square

3.1 Bends

Definition 3.5. Let G be an embedded planar graph, T a set of terminals, and P_1, \dots, P_k a solution. A d -bend $B = (P; C_0, \dots, C_d)$ consists of

- A subpath P of some P_i with endpoints x_0 and y_0 , and distinct vertices $x_0, \dots, x_d, y_d, \dots, y_0$ appearing on it in this order or in the reverse of this order.
- For every every $0 \leq i \leq d$, a free path C_i with endpoints $\{x_i, y_i\}$ (the *chords*) such that
 - The paths C_i are pairwise vertex disjoint.
 - The orientations of the paths are alternating, i.e., for every $0 \leq i < d$, vertex x_i is the start vertex of C_i if and only if x_{i+1} is the end vertex of C_{i+1} .
 - The internal vertices of C_i are not on P .
 - The undirected cycle formed by C_i and $P[x_i, y_i]$ (see footnote⁵) encloses C_j for every $j > i$.

We say that d -bend B (*strictly*) *encloses* a vertex/face/path if the undirected cycle formed by P and C_0 (*strictly*) encloses it. We say that B is *terminal free* if it does not enclose any terminals. We say that d -bend B *appears* on a path Q if P is a subpath of Q .

Note that we do not specify the orientation of the paths P and C_0 , (only that the orientations of C_0, \dots, C_d are alternating).

Observe that if $B = (P; C_0, \dots, C_d)$ is a d -bend of G , then it is a d -bend of every (embedded) supergraph of G : adding new arcs in an embedding-preserving way does not ruin any of the properties. It also follows that if B is a d -bend of a subgraph G , then it is a d -bend in G as well.

Definition 3.6. Given a solution P_1, \dots, P_k and d -bend $B = (P; C_0, \dots, C_d)$, the *type* of a d -bend is the number of paths P_i that have a vertex enclosed by B .

Note that if B appears on some P_i , then a subpath of P_i is enclosed by B and hence counted in the type of B . We prove the following statement by induction on type t .

Lemma 3.7. *There is a function $f(k, t)$ such that if P_1, \dots, P_k is a unique solution, then no $f(k, t)$ -bend of type t .*

⁵ If P is a directed path, then $P[x, y]$ is the subpath with endpoints x and y . We do not specify which of x and y is the start vertex, i.e., $P[x, y] = P[y, x]$.




Figure 18: A d -bend $B = (P; C_0, \dots, C_d)$.

The following simple parity argument will be used several times:

Lemma 3.8. *Let C be an undirected cycle, Q a subpath of C , and L a path with endpoints not enclosed by C such that L does not share any arcs with C , does not go through the endpoints of Q and crosses Q an odd number of times. Then there is a vertex w_1 of $Q \cap L$ and a vertex w_2 of $(C \setminus Q) \cap L$ such that the subpath $L[w_1, w_2]$ is enclosed by C and the internal vertices of $L[w_1, w_2]$ are disjoint from C .*

Proof. Clearly, L crosses C an even number of times; let p_1, \dots, p_{2f} be these crossing points in the order they appear on L (see Figure 19). As the endpoints of L are not enclosed by C , the subpath $L[p_{2j-1}, p_{2j}]$ is enclosed by C for every $1 \leq j \leq f$. Path L crosses Q an odd number of times, which means that there is a $1 \leq j \leq f$ such that exactly one of the crossing points p_{2j-1} and p_{2j} is on Q . If $L[p_{2j-1}, p_{2j}]$ has no internal vertex on Q , then we are done. Otherwise, as there are no crossing points between p_{2j-1} and p_{2j} , path L touches C at every such internal vertex. Suppose that there are s such touching points (possibly $s = 0$); let them be q_1, \dots, q_s in the order they appear on L from p_{2j-1} and p_{2j} . Let $q_0 = p_{2j-1}$ and $q_{s+1} = p_{2j}$. As exactly one of q_0 and q_{s+1} is on Q , there has to be an $0 \leq i \leq s$ such that exactly one of q_i and q_{i+1} is on Q . Let $w_1 \in \{q_s, q_{s+1}\}$ be the vertex on Q and let w_2 be the other one; the subpath $L[w_1, w_2]$ now has no internal vertex on C and is enclosed by C (as $L[q_0, q_{s+1}]$ is enclosed by C), thus $L[w_1, w_2]$ satisfies the requirements. \square

The following lemma shows that if a solution uses the innermost cycle of an alternating sequence of concentric cycles, then there is a d -bend. This allows us to use the bound on d -bends to show that a unique solution cannot use the innermost cycle. That is, the following lemmas show that Lemma 3.7 implies Lemma 3.4.

Lemma 3.9. *Let C_0, \dots, C_{d+1} be an alternating sequence of free concentric cycles. Assume that a path P_b of the solution contains a subpath $P := P_b[x_0, y_0]$ such that $x_0, y_0 \in C_0$, P does not contain any other vertex of C_0 except for x_0 and y_0 , and P intersects C_{d+1} . Moreover, assume that for one of the two subpaths C'_0 of C_0 between x_0 and y_0 , the cycle $P \cup C'_0$ does not enclose any terminals. Then there exists a terminal free d -bend $(P, C'_0, C'_1, \dots, C'_d)$, where C'_i is a subpath of C_i for $1 \leq i \leq d$.*

Proof. Let z be an arbitrary vertex of $C_{d+1} \cap P$. We may also assume that the graph has no other




Figure 19: Lemma 3.8.




Figure 20: A possible layout of the i -chords and the selection of the chords C'_i . Note that each chord is enclosed by the cycle C .

arc than the arcs of P and the arcs of the C_i 's: if we can find a d -bend after removing the additional arcs, then there is a d -bend in the original graph as well.

Consider the cycle $C := C'_0 \cup P$. As C intersects C_{d+1} , it cannot enclose any C_i with $i \leq d$. Therefore, every C_i with $0 \leq i \leq d$ has an arc e_i not enclosed by C . For every $0 \leq i \leq d$, let us subdivide e_i with two new vertices and let us remove the arc between these two new vertices. Now these two new vertices are the endpoints of a path L_i ; it is clear that the endpoints of L_i are not enclosed by C .

Let Q be the subpath $P[x_0, z]$. Clearly, Q crosses each C_i with $1 \leq i \leq d$ an odd number of times. As Q does not use the arc e_i , this implies that L_i crosses Q an odd number of times. Applying Lemma 3.8 on the path L_i and the subpath Q of cycle C , we get that for every $1 \leq i \leq d$, there is at least one subpath $L_i[w_1, w_2]$ enclosed by C such that $w_1 \in Q$, $w_2 \notin Q$, and the path has no internal vertex on P . We will call such a subpath of L_i an i -chord (see Figure 20). In particular, C'_0 is the unique 0-chord. Note that these chords are internally vertex disjoint (but two i -chords can share an endpoint), thus they have a natural ordering along Q , starting with C'_0 . Let us observe that if an i -chord and a j -chord are consecutive in this ordering, then $|i - j| \leq 1$: this is because then there is a face whose boundary contains both chords and no such face is possible if $|i - j| > 1$ (recall that we assumed that there is no other arc in the graph than the arcs of P and the arcs of the C_i 's).

Let $C'_i = L_i[x_i, y_i]$ be the first i -chord in this ordering. We claim that for every $0 \leq i \leq d$, the orientation of C'_i in the cycle $C'_i \cup P[x_i, y_i]$ is the same as the orientation of the cycle C_i . For $i = 0$,




Figure 21: Proof of Lemma 3.10 via Lemma 3.9. The dashed segments are the i -chords (note that there are three 2-chords, sharing some endpoints).

this follows from the fact that P is enclosed by C_0 . For some $1 \leq i \leq d$, let K be the last chord before C'_i . We claim that K is an $(i-1)$ -chord. By our observation in the previous paragraph, it is either an $(i-1)$ -chord, i -chord, or $(i+1)$ -chord. As C'_i is the first i -chord, K cannot be an i -chord. If K is an $(i+1)$ -chord, then there has to be another i -chord between C'_0 and K (as the difference is at most one between consecutive chords), again contradicting the choice of C'_i . Thus K is an $(i-1)$ -chord. Consider the face F whose boundary contains both K and C'_i : this face is between C_{i-1} and C_i . If C_i is oriented clockwise (the other case follows by symmetry), then C'_i goes counterclockwise on the the boundary of F . The undirected cycle $C'_0 \cup P[x_0, x_i] \cup C'_i \cup P[y_0, y_i]$ encloses C'_{i-1} , hence it encloses F as well, which means that the undirected cycle $C'_i \cup P[x_i, y_i]$ does not enclose F . It follows that if the orientation of C'_i is counterclockwise on the boundary of F , then it is clockwise on the undirected cycle $C'_i \cup P[x_i, y_i]$, i.e., same as the orientation of C_i . Therefore, we have shown that the orientations of the chords C'_0, \dots, C'_d are alternating, hence they form a d -bend. Moreover, as $P \cup C'_0$ does not enclose any terminals, the bend is terminal free. \square

Lemma 3.10. *Let C_0, \dots, C_{d+1} be an alternating sequence of free concentric cycles such that the outermost cycle C_0 does not enclose any terminals. If a path P_b of the solution intersects C_{d+1} , then there is a terminal-free d -bend appearing on P_b .*

Proof. Let z be an arbitrary vertex of $C_{d+1} \cap P_b$. As the endpoints of P_b are not enclosed by C_0 , both subpaths of P from z to its endpoints have to intersect C_0 . Let x_0 and y_0 be the intersections on these two subpaths closest to z . Let $P = P_b[x_0, y_0]$. Note that P is enclosed by C_0 and x_0, y_0 are the only vertices of P on C_0 . Let C'_0 be one of the two subpaths of C_0 between x_0 and y_0 chosen arbitrarily; by our assumption, no internal vertex of C'_0 is on P . Moreover, as C_0 does not enclose any terminals, neither does $C'_0 \cup P$. The claim follows from Lemma 3.9 on P, C'_0 and cycles C_0, C_1, \dots, C_{d+1} . \square




Figure 22: Proof of Lemma 3.11. The dotted segments are the i -chords, and the strong line is Q .

3.2 Paths and segments in bends

In this section, we discuss how new bends can be formed from paths appearing in a bend and how this can be used to obtain bounds on bends. From now on, we fix t and assume that Lemma 3.7 holds for $t - 1$, i.e., $f(k, t - 1)$ is defined.

The following lemma will be our main tool in constructing a new bend whenever there is a path that starts on some chord, visits many chords, and returns to the same chord. The proof is very similar to the proof of Lemma 3.10, but note that here we cannot assume that the path is enclosed by the bend B (see Figure 22).

Lemma 3.11. *Consider a solution P_1, \dots, P_k and let $B = (P, C_0, \dots, C_d)$ be a terminal-free d -bend appearing on some P_b . Let P' be a subpath of $P_{b'}$ (possibly $b = b'$) that is vertex-disjoint from P and going from $v_1 \in C_x$ to $v_2 \in C_x$ and intersecting C_y . Assume furthermore that the internal vertices of $C_x[v_1, v_2]$ are not on $P' \cup P_b$ and the cycle $C_x[v_1, v_2] \cup P'$ does not enclose any terminals. If $|x - y| \geq d' + 2$, then there is a terminal-free d' -bend $B' = (P'; C'_0, \dots, C'_{d'})$ with $C'_0 = C_x[v_1, v_2]$.*

Proof. We claim that the cycle $C := C_x[v_1, v_2] \cup P'$ does not enclose any vertex v of P . Since the endpoints of P_b (which are terminals) are not enclosed by C , the two subpaths of P_b from v to the endpoints of P_b intersect the cycle at two distinct points. This is only possible if $P_b = P_{b'}$ and v_1, v_2 are the two intersection points. However, then path P' and the subpath of P_b going from v_1 to v_2 via v forms a cycle, a contradiction. Therefore, no vertex of P is enclosed by C ; in particular, the endpoints of C_i are not enclosed by C for any $0 \leq i \leq d$.

In the rest of the proof, we will consider only the subgraph containing only the d -bend B and the path P' . It is clear that if we find the required d' -bend B' in this subgraph, then it is a d' -bend of the original graph as well.

Let $C_i^* = C_{x+i}$ if $y > x$ and let $C_i^* = C_{x-i}$ if $y < x$. Let z be an arbitrary vertex of $C_y \cap P'$ and let Q be the subpath $P'[v_1, z]$. Clearly, Q crosses each C_i^* with $1 \leq i \leq d' + 1$ an odd number




Figure 23: Lemma 3.11: An example where path P' starts on C_0 , intersects C_5 , and there is a 3-bend $(P; C'_0, C'_2, C'_3, C'_4)$. Note that X cannot be a chord as it intersects C'_0 and Y cannot be the chord after C'_0 as it has the wrong orientation.

of times. We have seen that C does not enclose the endpoints of C'_i . Applying Lemma 3.8 on the path C_i^* and the subpath Q of cycle C , we get that for every $1 \leq i \leq d' + 1$, there is at least one subpath $C_i^*[w_1, w_2]$ with $w_1 \in Q$, $w_2 \notin Q$, and no internal vertex on P . We will call such a subpath of C_i^* an i -chord. In particular, $C_0^*[v_1, v_2] = C_x[v_1, v_2]$ is a 0-chord, but there could be other 0-chords as well. Note that these chords are internally vertex disjoint (but two i -chords can share an endpoint), thus they have a natural ordering along Q . Let us observe that if an i -chord and a j -chord are consecutive in this ordering, then $|i - j| \leq 1$: this is because then there is a face whose boundary contains both chords and no such face is possible if $|i - j| > 1$.

Let $C'_i = C_i^*[x'_i, y'_i]$ be the first i -chord in this ordering. We claim that for every $1 \leq i \leq d'$, the orientation of C'_i (i.e., whether it goes from x'_i to y'_i or the other way around) and the orientation of C'_{i+1} are opposite. Let K be the last chord before C'_i . We claim that K is an $(i - 1)$ -chord. By our observation in the previous paragraph, it is either an $(i - 1)$ -chord, i -chord, or $(i + 1)$ -chord. As C'_i is the first i -chord, K cannot be an i -chord. If K is an $(i + 1)$ -chord, then there has to be another i -chord between C'_0 and K (as the difference is at most one between consecutive chords), again contradicting the choice of C'_i . Thus K is an $(i - 1)$ -chord. Recall that we are considering only the subgraph consisting of the d -bend B and the path P' . Therefore, there is a face F whose boundary contains both K and C'_i . Observe that C'_i and C'_{i+1} have the same orientation along the the boundary of F and this orientation depends on the parity of i . Note furthermore that face F is enclosed by the undirected cycle $C'_0 \cup P'[x'_0, x'_i] \cup C'_i \cup P'[y'_i, y'_0]$ and it is *not* enclosed by the cycle $C'_i \cup P'[x'_i, y'_i]$. Therefore, the orientation of C'_i along the cycle $C'_i \cup P'[x'_i, y'_i]$ is the opposite of its orientation along the boundary of the face F , that is, this orientation also depends on the parity of i . This proves that the orientation of C'_i and C'_{i+1} are opposite for every $1 \leq i \leq d'$. If the orientation of C'_0 and C'_1 are also opposite (see Figure 23 for an example), then we get a d' -bend $B' = (P; C'_0, C'_1, \dots, C'_{d'})$ and we are done. If the orientation of C'_0 is the same as the orientation of C'_1 (implying that it is the opposite of C'_2), then we get a d' -bend $B' = (P; C'_0, C'_2, C'_3, \dots, C'_{d'+1})$. \square

Note that the example on Figure 23 that the requirement $|x - y| \geq d' + 2$ in Lemma 3.11 is tight: path P' intersects C_0, \dots, C_5 , but there is only a 3-bend enclosed by P' .

Definition 3.12. Let $B = (P; C_0, \dots, C_d)$ be a d -bend in a solution P_1, \dots, P_k . A *segment* (with respect to B) is a subpath S of some P_b with endpoints on C_0 , enclosed by B , and no internal vertex on C_0 . A segment of path Q is a segment S with respect to B that is a subpath of Q . We say that a segment S with endpoints $q_1, q_2 \in C_0$ (strictly) encloses X if the undirected cycle $S \cup C_x[q_1, q_2]$ (strictly) encloses X . Two segments are *nested* if one encloses the other; a set of segments is nested if they are pairwise nested. If segment S_1 encloses segment S_2 , then we say that X is *between* S_1 and S_2 if X is enclosed by S_1 and no vertex of X is strictly enclosed by S_2 .

In particular, P is a segment of bend $B = (P; C_0, \dots, C_d)$. As the chords are free, every segment is arc disjoint from every chord. Note that two distinct segments of a path P_b are not necessarily disjoint: they can share endpoints. However, they cannot share both endpoints (otherwise they would form a cycle in the path P_b) and cannot share any internal vertices (since every internal vertex of the path P_b has exactly two arcs of P_b incident to it).

The parity argument of Lemma 3.8 can be stated with a chord C_i playing the role of L and a segment playing the role of the cycle C :

Lemma 3.13. *Let $B = (P; C_0, \dots, C_d)$ be a d -bend, let S be a segment, and let Q be a subpath of S with endpoints on C_x and C_y . For every $x < i < y$, there are vertices $w_1, w_2 \in S \cap C_i$ such that $w_1 \in Q$, $w_2 \notin Q$, $C_i[w_1, w_2]$ is enclosed by S and has no internal vertex on S , and $S[w_1, w_2]$ includes an endpoint of Q .*

Proof. Clearly, Q crosses C_i an odd number of times. Let v_1 and v_2 be the endpoints of S on C_0 . We apply Lemma 3.8 on the line $L = C_i$ and the cycle $C = S \cup C_0[v_1, v_2]$. As $w_1 \in Q$ and $w_2 \notin Q$, the path $S[w_1, w_2]$ includes an endpoint of Q . \square

In the following lemmas, we prove certain bounds on d -bends under the assumption that the solution is unique. Note that the following proof and the proof of Lemma 3.17 are the only places where we directly use the fact that a solution is unique; all the other arguments build on these two proofs.

Lemma 3.14. *Consider a unique solution P_1, \dots, P_k and let $(P; C_0, \dots, C_d)$ be a terminal-free d -bend of type t appearing on P_b . If no internal vertex of C_0 is on P_b , then $d \leq f(k, t - 1) + 2$.*

Proof. Observe that P is the only segment of P_b with respect to B : as C_0 has no internal vertex on P_b , the endpoints of a segment S of P_b have to be x_0 and y_0 , and therefore if S and P are different, they would form a cycle in P_b , a contradiction. This means that C_d and C_{d-1} have no internal vertex on P_b . Now for some $i \in \{d, d - 1\}$, the orientation of C_i is such that paths $P[x_0, x_i]$, C_i , $P[y_0, y_i]$ can be concatenated to obtain a directed path P' . If no internal vertex of C_i is used by the solution, then we can replace P by P' in P_b to obtain a new solution, contradicting the assumption that P_1, \dots, P_k is unique. Thus C_i is intersected by a segment S of some $P_{b'}$ with $b \neq b'$. Now Lemma 3.11 implies that there is a d' -bend $B' = (S; C'_0, \dots, C'_{d'})$ with $d' = i - 2 \geq d - 3$ and C'_0 being the subpath of C_0 connecting the endpoints of S . Observe that $S \cup C'_0$ is enclosed by B , thus the type of B' is at most the type of B . In fact, the type of B' is strictly smaller: as S and C'_0 are disjoint from P_b , no vertex of P_b is enclosed by B . By the induction hypothesis of Lemma 3.7, this implies $d' < f(k, t - 1)$ and therefore $d \leq f(k, t - 1) + 2$ follows from $d' \geq d - 3$. \square

Lemma 3.15. *Consider a unique solution P_1, \dots, P_k and let $B = (P, C_0, \dots, C_d)$ be a terminal-free d -bend of type t appearing on some P_b . Let Q be an subpath of P_b with endpoints $v_1, v_2 \in C_x$*

that intersects C_y . If $C_x[v_1, v_2]$ has no internal vertex on P_b and the cycle $C_x[v_1, v_2] \cup Q$ does not enclose any terminals, then $|x - y| \leq f(k, t - 1) + 4$ holds.

Proof. If $C_x[v_1, v_2]$ has no internal vertex on P_b , then Lemma 3.11 implies the existence of a terminal-free d' -bend $B' = (Q; C'_0, \dots, C'_{d'})$ with $d' \geq |x - y| - 2$ and $C'_0 = C_x[v_1, v_2]$. However, since C'_0 has no internal vertex on P_b , Lemma 3.14 implies $d' \leq f(k, t - 1) + 2$ and the claim follows. \square

Lemma 3.16. *Consider a unique solution P_1, \dots, P_k and let $(P; C_0, \dots, C_d)$ be a terminal-free d -bend of type t appearing on some P_b . Let S be a segment of P_b that intersects C_y for some $y \geq f(k, t - 1) + 6$. Then S encloses another segment of P_b that intersects $C_{y'}$ for $y' = y - f(k, t - 1) - 5 \geq 1$.*

Proof. Let us first verify the claim for the case when $P = S$. If no segment of P_b intersects the internal vertices of $C_{y'}$, then $B' = (P[x_{y'}, y_{y'}]; C_{y'}, \dots, C_d)$ is a d' -bend with $d' = d - y' \geq f(k, t - 1) + 5$ such that no internal vertex of $C_{y'}$ is on P_b , which would contradict Lemma 3.14. Thus in the following, we assume that $P \neq S$.

Let z be a vertex of $S \cap C_y$ and let s_1, s_2 be the endpoints of S on C_0 . Let Q be the subpath $S[s_1, z]$. By Lemma 3.13, there is a subpath $C_{y'}[w_1, w_2]$ enclosed by S and $w_1 \in Q$, $w_2 \in S \setminus Q$ and having no internal vertex on S (note that $y' \geq 1$ and the endpoints of $C_{y'}$ are not enclosed by S as $S \neq P$). This implies that $S[w_1, w_2]$ contains z and hence intersects $C_{y'}$. If the internal vertices of $C_{y'}[w_1, w_2]$ are intersected by P_b , then we are done, as every such vertex is on a segment S' (different from S) enclosed by S . Otherwise, we apply Lemma 3.15 on the segment $S[w_1, w_2]$ (note that $P \neq S$ implies that $S[w_1, w_2]$ is disjoint from P), and $y - y' \geq f(k, t - 1) + 5$ gives a contradiction. \square

3.3 Rerouting nested segments

Our goal is to show that if we have a large nested set of segments crossing many chords and having no other segment between them, then we can simplify the solution by rerouting. However, it seems hard to ensure the requirement that there are no other segments in between; therefore, we weaken the requirement by asking that there is a large area (intersecting many chords) such that there are no other segments between our nested segments in this area (see Figure 24).

Lemma 3.17. *Consider a unique solution P_1, \dots, P_k . Let $B = (P; C_0, \dots, C_d)$ be a d -bend and let S^1 and S^2 be two nested segments with respect to B . Let $x^1, y^1 \in S^1$ and $x^2, y^2 \in S^2$ be vertices with $x^1, x^2 \in C_x$ and $y^1, y^2 \in C_y$ for some $y \geq x + 2(2^k + 1)$ such that subpaths $C_x^* := C_x[x^1, x^2]$ and $C_y^* := C_y[y^1, y^2]$ are between S^1 and S^2 . If the set of segments intersecting C_x^* or C_y^* (including S^1 and S^2) is nested between S^1 and S^2 , then this set has size at most 2^k .*

The proof of Lemma 3.17 requires two tools. As an important step in the proof of the undirected irrelevant vertex rule of Adler et al. [1, Lemma 6], it is proved that if there is a disc in the plane such that a solution of the k -disjoint paths problem that contains nothing else than a set of $\ell > 2^k$ parallel paths, then one can “redraw” the solution by replacing the part of the solution in the disc by introducing a set of strictly less than ℓ new noncrossing edges inside the disc. Then if these new “virtual” edges can be actually realized by the edges of the graph inside the disc, then it follows that we can modify the solution, and hence it is not unique.

The following lemma is a directed analogue of the statement of Adler et al. [1, Lemma 6]. The same proof⁶ works for the directed case; in fact, the proof in [1] first assigns an arbitrary orientation to each path, thus it is clear that there is a correct directed solution in the modified graph.

⁶The proof appears in the full version, which can be accessed at <http://www.ii.uib.no/~daniello/papers/planarUniqueLinkage>.




Figure 24: Condition of Lemma 3.17: the segments intersecting $C_x[x_1, x_2]$ or $C_y[y_1, y_2]$ (i.e., segments entering the gray area) are nested.

Lemma 3.18. *Let G be a graph embedded in the plane with k pairs of terminals with a solution P_1, \dots, P_k . Let D be a closed disc not containing any terminals and suppose that $x_1, \dots, x_\ell, y_\ell, \dots, y_1$ appear on the boundary of D (in this order) for some $\ell > 2^k$. For every $1 \leq i \leq \ell$, let Q_i be a directed path with endpoints $\{x_i, y_i\}$ (going in arbitrary direction) such that these paths are pairwise disjoint and contained in D . Suppose that every Q_i appears as a subpath of the solution and the solution uses no other vertex in D . Let V_S and V_E be the starting points and end points of all the Q_i 's (clearly, $|V_S| = |V_E| = \ell$). Then there is a noncrossing matching of size strictly less than ℓ between V_S and V_E such that if G' is the planar graph obtained by removing every arc in D and adding every e_j oriented from V_S to V_E , then there is a solution in G' .*

To use Lemma 3.18 to modify a solution, we have to show that there are pairwise disjoint directed paths inside the disc D that correspond to the noncrossing matching. The existence of such directed paths connecting vertices on the boundary of the disc can be conveniently shown by the sufficient and necessary condition given by Ding, Schrijver, and Seymour [5], which will be our second tool in the proof of Lemma 3.17. We review this condition next.

Let G be a directed graph embedded in the plane; for simplicity we will treat only the case when the boundary of the infinite face is a simple cycle, as it is already sufficient for our purposes. Let $(r_1, s_1), \dots, (r_k, s_k)$ be pairs of terminals on the boundary such that all $2k$ terminals are different. We say that (r_i, s_i) and (r_j, s_j) *cross*, i.e., they appear in the order r_i, r_j, s_i, s_j or in the order r_i, s_j, s_i, r_j on the boundary of the infinite face. We say that the *noncrossing condition* holds if no two pairs cross. We say that the *cut condition* holds if the following is true for every curve C of the disc D going from a point on the boundary of D to another point on the boundary of D , intersecting G only in the vertices, and not intersecting any terminal:

If C separates $(r_{i_1}, s_{i_1}), \dots, (r_{i_n}, s_{i_n})$ in this order, then C contains vertices p_1, \dots, p_n , in this order so that for each $j = 1, \dots, n$:




Figure 25: Proof of Lemma 3.17. The figure shows two nested segments between S^1 and S^2 . The light gray shows the disc D_0 , which is extended to D by the dark gray area.

- if r_{i_j} is at the left-hand side of C , then at least one arc of G is entering C at p_j from the left and at least one arc of G is leaving C at p_j from the right.
- if r_{i_j} is at the right-hand side of C , then at least one arc of G is entering C at p_j from the right and at least one arc of G is leaving C at p_j from the left.

It is easy to see that both the noncrossing and the cut conditions are necessary for the existence of pairwise vertex-disjoint paths connecting the terminal pairs. The result of Ding, Schrijver, and Seymour states that these two conditions are sufficient:

Theorem 3.19 ([5]). *Let $(r_1, s_1), \dots, (r_k, s_k)$ be pairs of terminals on the infinite face of an embedded planar graph. There exist k pairwise vertex-disjoint directed paths P_1, \dots, P_k such that P_i goes from r_i to s_i if and only if the noncrossing and the cut conditions hold.*

We are now ready to prove Lemma 3.17. Let us remark that this proof is the second and last place (after Lemma 3.14) where the uniqueness of the solution is directly used.

Proof (of Lemma 3.17). Let S_1, \dots, S_ℓ be the segments nested between S^1 and S^2 intersecting C_x^* or C_y^* in the order they are nested ($S_1 = S^1, S_\ell = S^2$). If $\ell \leq 2^k$, then we are done. Otherwise, we argue that it can be assumed that $\ell = 2^k + 1$. Let $S' := S_{2^k+1}$. Let x' be the vertex of S' on C_x^* closest to x^1 one C_x^* (such a vertex exists, as S' is nested between S^1 and S^2). Similarly, let y' be the vertex of S' on C_y^* closest to y^1 . Now we can replace S^2, x^2, y^2 , with S', x', y' , respectively. Note that $C_x[x^1, x']$ is between S^1 and S' , hence only the segments $S_1, \dots, S_{2^k+1} = S'$ can intersect it, satisfying the conditions of the lemma being proved. Thus if we can arrive to a contradiction for $\ell = 2^k + 1$, the lemma follows.

Let D_0 be the disc enclosed by the cycle $S^1[x^1, y^1] \cup C_x^* \cup S^2[x^2, y^2] \cup C_y^*$. It is tempting to try to apply Lemma 3.18 on the disc D_0 and the subpaths of the S_i 's inside D_0 , but as Figure 25 shows, these subpaths are not necessarily parallel (and therefore Figure 24 gives only a simplified picture of what is happening inside the gray area). To avoid this difficulty, we extend D_0 to a disc D the following way: for every subpath Q of every S_i , if Q has both endpoints on C_x^* or both endpoints on C_y^* and no other vertex in D_0 , then we add to D the disc enclosed by Q and the subpath of C_x^*

or C_y^* between the endpoints of Q . This process creates a disc D whose boundary consists of the subpath $S^1[x^1, y^1]$, a path B_x from x^1 to x^2 , the subpath $S^2[x^2, y^2]$, and a path B_y from y^1 to y^2 .

Claim 3.20. *Only the segments S_1, \dots, S_ℓ intersect D .*

Proof. Suppose that a segment S contains a vertex x in D . If x is in D_0 , then it is clear that S is part of the nested sequence of segments. Otherwise, x is in D because it is enclosed by a cycle formed by a subpath Q of some S_i and a subpath of C_x^* or C_y^* . Since every segment intersecting C_x^* or C_y^* is in the nested sequence by assumption, we can conclude that S is also one of these nested segments. \square

Among all intersections of S_i with C_x^* or C_y^* consider those two that are closest to the endpoints of S_i ; one of these intersections (denote it by x_i) has to be on C_x^* , the other (denote it by y_i) has to be on C_y^* . Let $Q_i := S_i[x_i, y_i]$.

Claim 3.21. *For every $1 \leq i \leq \ell$,*

1. *vertex x_i is on B_x ,*
2. *vertex y_i is on B_y , and*
3. *Q_i is enclosed by D .*

Proof. If x_i is not on the boundary of D , then it is enclosed by a subpath Q of some S_j and a subpath of C_x^* , which contradicts the assumption that there is a subpath of S_i from x_i to an endpoint of S_i and not intersecting $C_x^* \cup C_y^*$. Similarly, every y_i is on the boundary segment B_y of D .

Suppose that a vertex z of Q_i is not enclosed by D (and hence by D_0). Let Q' be the subpath of Q_i containing z with both endpoints on the boundary of D_0 and no internal vertex enclosed by D_0 . Then either both endpoints are on C_x^* or both endpoints are on C_y^* . In either case, the definition of D would add the disc enclosed by Q' to D , contradiction that z is not enclosed by D . \square

It is clear that the solution uses no other vertex of D than the vertices of the Q_i 's: we have seen that every vertex in D belongs to some S_i and this vertex of S_i has to be part of Q_i . Therefore, the conditions of Lemma 3.18 hold for Q_1, \dots, Q_ℓ and we may assume the existence of the matching $M = \{e_1, \dots, e_{\ell'}\}$ of size $\ell' < 2^k + 1$. We consider the arcs $e_1, \dots, e_{\ell'}$ to be directed, with orientation as given by Lemma 3.18.

Claim 3.22. *There exist pairwise vertex-disjoint paths $R_1, \dots, R_{\ell'}$ enclosed by D such that R_j has the same start vertex r_j and end vertex s_j as the start and end vertex of e_j , respectively.*

Proof. We use Theorem 3.19 to prove the existence of these paths. The fact that the matching is noncrossing implies that the noncrossing condition holds. Suppose now that a curve C in D with endpoints on the boundary of D violates the cut condition. If C intersects both B_x and B_y , then it has a subpath with an endpoint on C_x^* and an endpoint on C_y^* . Therefore, C crosses all the chords between C_x and C_y , thus we can find the required vertices p_1, \dots, p_{2^k+1} (note that $y \geq x + 2(2^k + 1)$). Therefore, we may assume that C is disjoint from B_y (the case when C has no endpoint on B_x is similar).

Without loss of generality, we may assume that C separates the pairs $(r_1, s_1), \dots, (r_q, s_q)$ only and in this order. As C is disjoint from B_y , there is a subpath B' of the boundary of D that connects the endpoints of C and is disjoint from B_y ; let b_1, b_2 be the endpoints of B' . Now C separates the pair (r_j, s_j) if and only if exactly one of r_j and s_j is on B' ; denote this vertex by

$\beta_j \in \{r_j, s_j\}$. Moreover, the order in which C separates the pairs correspond to the order in which the vertices β_j corresponding to the separated pairs appear on B' , that is, vertices $b_1, \beta_1, \dots, \beta_z, b_2$ appear in this order on B' . Each vertex β_j is a vertex x_{i_j} for some $1 \leq i_j \leq 2^k + 1$ and there is a corresponding path Q_{i_j} that enters β_j if $\beta_j = r_j$ and leaves β_j if $\beta_j = s_j$. Let p_j to be the first intersection of C with Q_{i_j} (note that this cannot be x_{i_j} or y_{i_j} , as C does not intersect the terminals). The path Q_{i_j} shows that p_j has the required arcs entering and leaving, and therefore this sequence witnesses that C does not violate the cut condition. \square

As the matching M was given by Lemma 3.18, if we remove every arc enclosed by D and add the arcs $e_1, \dots, e_{\ell'}$, then there is a solution. Therefore, if we remove every arc enclosed by D except those that are on some R_i , then there is still a solution. This solution is different from P_1, \dots, P_k : the paths in the solution have at most $\ell' < 2^k + 1$ maximal subpaths enclosed by D . This contradicts the assumption that P_1, \dots, P_k is a unique solution. \square

3.4 Rerouting in a bend

Before beginning the main part of the inductive proof of Lemma 3.4, we need to introduce one more technical tool. We have to be careful to avoid certain subpaths of the solution, as they enclose terminals and hence the inductive argument cannot be applied on them. The following definitions will be helpful in analyzing this situation:

Definition 3.23. Let P_1, \dots, P_b be a solution, let $B = (P', C_0, \dots, C_d)$ be a d -bend appearing on a path P_b . Let H be the undirected P_b -graph formed by the paths P_b and C_0, \dots, C_d , with every degree-2 vertex suppressed. We call an edge of H a P_b -arc if it corresponds to a subpath of P_b or a chord arc if it corresponds to a subpath of some C_i . A subpath of C_i that corresponds to a chord arc (i.e., the endpoints of the subpath is on P_b and the internal vertices are disjoint from P_b) is called a P_b -bridge. The dual P_b -graph is the dual H^* of H . The subgraph T^* of H^* containing the chord arcs is the P_b -tree of B .

Note that the graphs defined in Definition 3.23 are all undirected. The following lemma justifies the name P_b -tree:

Lemma 3.24. Let $B = (P', C_0, \dots, C_d)$ be a d -bend appearing on an a path P_b of the solution. The P_b -tree T^* of B is a spanning tree of the dual P_b -graph.

Proof. Suppose that there is a cycle C^* in T^* . This cycle C^* corresponds to a cut C in the primal graph H , thus removing the edges of C disconnects the graph H . However, P_b is a connected spanning subgraph of H not containing any edge of the cut C , a contradiction. To see that T^* is connected and spanning, suppose that the dual graph H^* has a minimal cut consisting only of P_b -edges. Then the primal graph H contains a cycle consisting only of P_b edges, meaning that there is cycle in P_b , a contradiction. \square

A variant of the segment is the j -segment, which has its endpoints on C_j :

Definition 3.25. Let P_1, \dots, P_k be a solution and let $B = (P; C_0, \dots, C_d)$ be a d -bend. A j -segment is a subpath Q of some P_b with endpoints on C_j , no internal vertex on C_j , and enclosed by the cycle $C_j \cup P[x_j, y_j]$.

Note that every j -segment S_j is on a unique segment S , but segment S can contain multiple j -segments. Observe that if j -segment S_j is on segment S , then the subpath of C_j connecting the




Figure 26: A d -bend $(P; C_0, \dots, C_d)$ appearing on a path P_b . The dotted lines show the arcs of the P_b -tree.

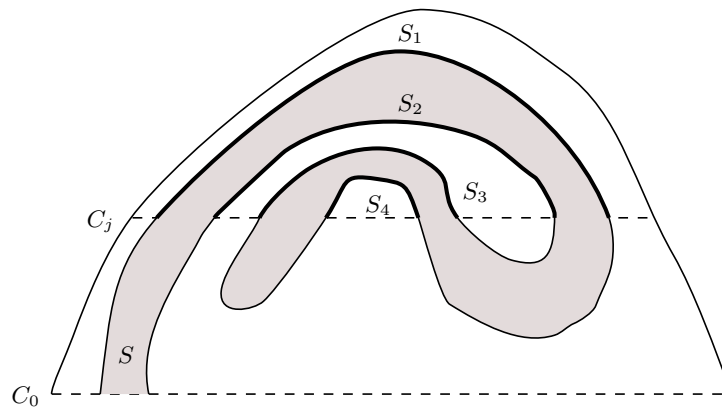


Figure 27: A segment S in a bend and four j -segments S_1, \dots, S_4 on S .




Figure 28: In T_b^* , there are two disjoint paths from $C_i[\alpha_1, \beta_1]$ to the faces between C_j and C_{j-1} , but there are no two disjoint paths from $C_i[\alpha_2, \beta_2]$, as $C_{j'}[\alpha', \beta']$ is a separator.

endpoints of S_j are not necessarily enclosed by S (see Figure 27). Therefore, to avoid confusion, we do not define the notion of “enclosing” for j -segments.

We need the following technical lemma in the proof of Claim 3.28. Intuitively, if S is a j -segment and e is a P_b -bridge having an endpoint on S , then the P_b -bridges having an endpoint on S give two paths from e to C_j in the P_b -tree T_b^* . For example, in Figure 28, one can see the two paths from $C_i[\alpha_1, \beta_1]$ to C_j . However, these two paths are not necessarily edge disjoint: for example, in Figure 28, there are no two disjoint paths from $C_i[\alpha_2, \beta_2]$ to C_j . The following lemma shows that the two disjoint paths always exist if the P_b -bridge e has *exactly* one endpoint on S (as it is the case with $C_i[\alpha_1, \beta_1]$, but not with $C_i[\alpha_2, \beta_2]$).

Lemma 3.26. *Let $B = (P; C_0, \dots, C_d)$ be a d -bend appearing on, let S be a j -segment of some P_b , and let edge e of the P_b -tree T_b^* correspond to P_b -bridge $C_i[\alpha, \beta]$ having exactly one endpoint on S . Let T_S^* be the subgraph of T_b^* containing edges corresponding to P_b -bridges having an endpoint on S . Then there are two edge-disjoint paths Q_1, Q_2 in T_S^* such that the first vertex of each Q_i is an endpoint of e and its last vertex corresponds to a face of the P_b -graph between C_j and C_{j-1} .*

Proof. Suppose without loss of generality that $\alpha \in S$. Note that the endpoints of e correspond to faces of the P_b -graph, one between C_i and C_{i+1} , the other between C_i and C_{i-1} . By Menger’s Theorem, if there are no such paths, then there is an edge e' of T_S^* separating the endpoints of e from the faces between C_j and C_{j-1} . Suppose that e' corresponds to P_b -bridge $C_{j'}[\alpha', \beta']$ for some $\alpha', \beta' \in P_b$. Observe that both α' and β' have to be on S , otherwise S has a subpath from α to one of its endpoints on C_j that is disjoint from $\{\alpha', \beta'\}$, and the P_b -bridges along this subpath of S give a path in T_S^* avoiding e' , contradicting the assumption that e' is a separator. By the same reason, α has to appear between α' and β' on S , that is, $S[\alpha', \beta']$ contains α .

Let C_S be the cycle formed by S and the subpath of C_j connecting the endpoints of S . Suppose first that C_S encloses β . Consider the cycle $C = C_{j'}[\alpha', \beta'] \cup S[\alpha', \beta']$. This cycle encloses α (as $\alpha \in S[\alpha', \beta']$) but cannot enclose β , which is on a j -segment of P_b different from S : $C_{j'}[\alpha', \beta']$ has no internal vertex on P_b . This is only possible if $C_{j'}[\alpha', \beta']$ is not enclosed by C_S . Thus the edges of T_S^* corresponding to P_b -bridges enclosed by C_S are disjoint from e' and connect e to the faces between C_j and C_{j-1} .

The argument is similar if C_S does not enclose β . Again, $C = C_{j'}[\alpha', \beta'] \cup S[\alpha', \beta']$ encloses α but does not enclose β . This is only possible if $C_{j'}[\alpha', \beta']$ is enclosed by C_S . Thus the edges of T_S^* corresponding to P_b -bridges not enclosed by C_S are disjoint from e' and connect e to the faces between C_j and C_{j-1} . \square

We are now ready to start the main part of the proof.

Lemma 3.27. *If $f(k, t - 1)$ is defined, then $f(k, t)$ is defined.*

Proof. We define the following constants:

$$s := 2^k + 1 \quad m := f(k, t - 1) + 5 \quad M := 40sm, \quad f(k, t) := M(2k + 4).$$

Note that this recursive definition of $f(k, t)$ implies that $f(k, t) = 2^{O(kt)}$. Suppose that a terminal-free d -bend $B = (P; C_0, \dots, C_d)$ of type t appears on path P_b in a unique solution for some $d \geq f(k, t)$. Let T_b^* be the P_b -tree as in Definition 3.23. We define a set F of special faces of the P_b -graph H containing

- the infinite face,
- faces strictly enclosing a terminal $s_{b'}$ or $t_{b'}$ for some $b' \neq b$,
- the at most two faces whose boundaries contain s_b and t_b (which are degree-1 vertices), and
- the two faces whose boundary contains the arc of P incident to x_0 .

Note that $|F| \leq 2k + 3$. Let T_0^* be the minimal subtree of T_b^* containing every vertex that corresponds to a face in F . Let X be the set of vertices of T_b^* that have degree at least 3 in T_0^* . As T_0^* has at most $2k + 3$ leaves, we have $|X| \leq 2k + 1$.

Consider the faces of H_b enclosed by the d -bend B . At most $2k + 1$ of these faces correspond to elements of X , thus $d \geq f(k, t)$ implies that there is an $h > M$ such that no face corresponding to X appears between C_h and C_{h-M} in the d -bend B . Let us fix such an h and let $h^* = h - 30sm$.

Let α and β be two vertices of P_b enclosed by B . We say that α sees β from the inside if α and β are both on the same C_i , vertex α is strictly enclosed by the segment of P_b containing β , and the subpath $C_i[\alpha, \beta]$ does not have any internal vertex on P_b . Note that this subpath $C_i[\alpha, \beta]$ may intersect some $P_{b'}$ with $b' \neq b$.

First we find a sequence of nested segments S_1, \dots, S_s of P_b (see Figure 29) such that, in a weak but precise technical sense, there are no further segments of P_b between them. What we show is that each segment has a subpath that is seen from the inside only by vertices of the next segment in the sequence.

Claim 3.28. *There are distinct nested segments S_1, \dots, S_s of P_b with respect to B and an $(h^* - im)$ -segment S'_i of each S_i such that*

1. S'_i does not intersect C_h .
2. S'_i intersects C_{h-im} ,




Figure 29: The segments in Claim 3.28.

3. every vertex of P_b seeing a vertex of S'_i from the inside is on S'_{i+1} .

Proof. There is at least one segment of P_b intersecting C_{h-m} : P itself is such a segment. Let S_1 be a segment of P_b intersecting C_{h-m} , but not enclosing any other segment of P_b intersecting C_{h-m} . It follows that S_1 does not intersect C_h : otherwise Lemma 3.16 implies that it encloses another segment intersecting $C_{h-f(k,t-1)-5} = C_{h-m}$. The segments S_2, \dots, S_s we construct in the rest of the proof are all enclosed by S_1 , thus they do not intersect C_h either. Let S'_1 be an arbitrary $(h^* - m)$ -segment of S_1 intersecting C_{h-m} .

Suppose that we have constructed such a sequence up to S_{i-1} and S'_{i-1} . We find S_i and S'_i the following way. Again by Lemma 3.16, there is a segment of P_b enclosed by S_{i-1} that intersects C_{h-im} , thus there is at least one S_i enclosed by S_{i-1} . Therefore, there is at least one $(h^* - im)$ -segment S'_i intersecting C_{h-im} . We show that there is at most one such $(h^* - im)$ -segment that contains vertices seeing S'_{i-1} from inside, thus we can define S'_i satisfying property (3).

Suppose $(h^* - im)$ -segments S' and S'' of P_b contain vertices α' and α'' that see $\beta', \beta'' \in S'_{i-1}$ from the inside, respectively. Let S^* be the $(h^* - im)$ -segment of S'_{i-1} (note that S'_{i-1} is a $(h^* - (i-1)m)$ -segment, thus its $(h^* - im)$ -segment is unique). Let L be the vertices of T_b^* corresponding to faces between C_{h^*-im} and C_{h^*-im-1} .

Let Z be the subtree of T_b^* that corresponds to P_b -bridges enclosed by S_{i-1} and having an endpoint on S'_{i-1} ; clearly, Z is connected. The edges corresponding to P_b -bridges with an endpoint on S^* contain a path Q from Z to L (see Figure 30). By Lemma 3.26 applied on the P_b -bridge $\alpha'\beta'$ and the segment S' (resp., $\alpha''\beta''$ and S''), we get two edge-disjoint paths Q'_1, Q'_2 (resp., Q''_1, Q''_2) that go between Z and L and each edge of the paths corresponds to a P_b -bridge having an endpoint on S' (resp., S''). We show that in T_b^* there are three edge-disjoint paths between Z and L . If there are no such paths, then Menger's Theorem implies that there are two edges e_1 and e_2 covering all such paths. Observe that each edge can be contained in at most two out of the five paths $Q, Q'_1, Q'_2, Q''_1, Q''_2$: otherwise the P_b -bridge corresponding to the edge would have an endpoint on all three of the $(h^* - im)$ -segments $S^*, S',$ and S'' , which is impossible. Therefore, one of these paths is

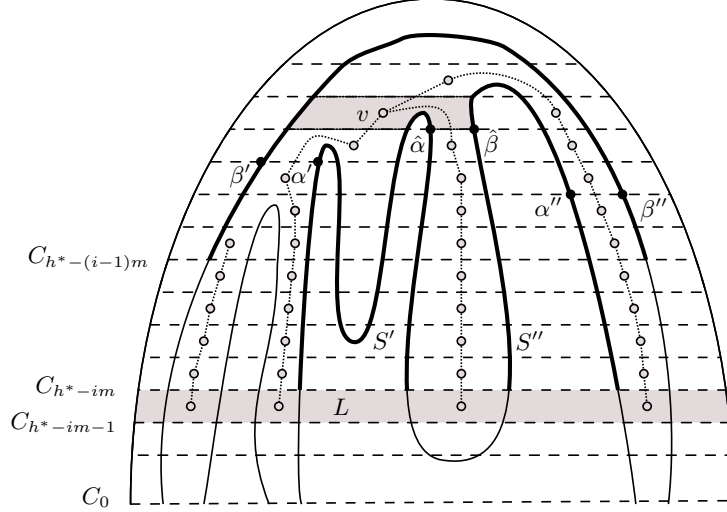


Figure 30: Proof of Claim 3.28. In T_b^* , there are two disjoint paths from $\alpha'\beta'$ to L , and two disjoint paths from $\alpha''\beta''$ to L using the edges shown on the figure. There is also a path from Z to L using edges that correspond to P_b -bridges having an endpoint on S_{i-1} . There are three edge-disjoint paths from vertex v to L in T_b^* , and we arrive to a contradiction using the path $P_b[\hat{\alpha}, \hat{\beta}]$ intersecting many chords but not enclosing any terminals.

disjoint from e_1 and e_2 , a contradiction.

Let Q_1, Q_2, Q_3 be three edge-disjoint paths from Z to L ; we can assume that they start at (possibly not distinct) vertices v_1, v_2, v_3 of Z and they contain no other vertices of Z . Let Z' be the minimal subtree of Z containing v_1, v_2 , and v_3 . This subtree Z' has a vertex v (possibly $v \in \{v_1, v_2, v_3\}$) and three edge-disjoint paths Z_1, Z_2, Z_3 (possibly of length 0) where Z_i goes from v to v_i . Now the concatenation of Z_i and Q_i for $i = 1, 2, 3$ gives three edge-disjoint paths from v to L . As $v \in Z$, the length of the paths is at least $m = f(k, t - 1) + 5$. Let $\hat{T}_1, \hat{T}_2, \hat{T}_3$ be the components of $T_b^* \setminus v$ that contain these paths (minus v). We have chosen h such that no face of X appears between C_h and C_{h-M} ; in particular, as $h^* - im > h - M$, vertex v is not in X . Therefore, it cannot happen that all three of $\hat{T}_1, \hat{T}_2, \hat{T}_3$ contain vertices from T_0^* : this would imply that $v \in T_0^*$ and has degree at least 3 in T_0^* , i.e., $v \in X$ by the definition of X . Suppose that $\hat{T} \in \{\hat{T}_1, \hat{T}_2, \hat{T}_3\}$ is disjoint from T_0^* and let \hat{e} be the edge connecting \hat{T} and v . Suppose that \hat{e} corresponds to P_b -bridge $C_j[\hat{\alpha}, \hat{\beta}]$. Note that $\hat{\alpha}, \hat{\beta} \notin P$, as they are enclosed by S_{i-1} (which is different from P).

We would like to invoke Lemma 3.15 with $Q = P_b[\hat{\alpha}, \hat{\beta}]$ to arrive to a contradiction, but we need to verify the conditions that no terminal is enclosed and this path is disjoint from P . This is the part of the proof where the definition of T_0^* and the fact that \hat{T} is disjoint from T_0^* comes into play. Consider the cycle C formed by $\hat{\alpha}\hat{\beta}$ and $P_b[\hat{\alpha}, \hat{\beta}]$. Removing the edge \hat{e} from T^* splits T^* into two parts, one of which is \hat{T} . The cycle C encloses the faces corresponding to one of these two parts; more precisely, it encloses the part that does not contain the infinite face. As \hat{T} is disjoint from T_0^* , it cannot contain the infinite face, thus C encloses exactly the faces of \hat{T} . This means that C does not enclose any face of T_0^* and hence does not enclose any terminals. Moreover, we claim that C is disjoint from P . As $\hat{\alpha}, \hat{\beta} \notin P$, if $P_b[\hat{\alpha}, \hat{\beta}]$ contains a vertex of P , then it fully contains P , including the arc of P incident to x_0 . Both faces incident to this arc is in T_0^* , thus if C contains this arc, then C encloses a face of T_0^* , a contradiction. Thus $P_b[\hat{\alpha}, \hat{\beta}]$ is disjoint from P . The tree \hat{T} contains a vertex of L (since it contains one of the three paths Q_1, Q_2, Q_3 , minus v), which means that C




Figure 31: A more complicated example of the segments in Claim 3.28. Segment S_i has an $(h^* - im)$ -segment different from S'_i , which is seen from inside by vertices on a segment S^* different from S_{i+1} , and also by vertices on a $(h^* - (i + 1)m)$ -segment of S_{i+1} different from S'_{i+1} .

encloses an arc of $C_{h^* - im}$. Therefore, $P_b[\hat{\alpha}, \hat{\beta}]$ intersects $C_{h^* - im}$ and $\hat{\alpha}, \hat{\beta}$ has no internal vertex on P_b . Thus we arrive to a contradiction by Lemma 3.15. \lrcorner

Note that the statement of Claim 3.28 is somewhat delicate. It does not claim that every vertex of S_i on C_j for some $j \geq h^* - im$ is only seen from inside only by S'_{i+1} ; it claims this only for a one specific $(h^* - im)$ -segment S'_i of S_i . Also, the $(h^* - (i + 1)m)$ -segment S'_{i+1} could contain more than one $(h^* - im)$ -segments (see Figure 31).

Not all vertices of S'_i are seen from inside by the vertices of S'_{i+1} : for example, if for some $v_1, v_2 \in S'_i \cap C_j$, the path $C_j[v_1, v_2]$ is enclosed by S_i and does not have any internal vertex on P_b , then v_1 and v_2 are not seen from inside by any vertex of S_{i+1} . Nevertheless, the following claim shows that if a subpath of S'_i intersects many chords, then it contains a vertex seen from inside by a vertex of S'_{i+1} .

Claim 3.29. *Let Q be a subpath of S'_i having an endpoint on C_x and an endpoint on C_y . If $x + f(k, t - 1) + 5 \leq j \leq y - f(k, t - 1) - 5$ holds for some j , then there exists at least one vertex of Q on C_j that is seen from inside by a vertex of S'_{i+1} .*

Proof. By Lemma 3.13, there are vertices $w_1, w_2 \in C_j \cap S_i$ such that subpath $C_j[w_1, w_2]$ is enclosed by S_i and $S_i[w_1, w_2]$ contains an endpoint of Q , which implies that $S_i[w_1, w_2]$ intersects C_x or C_y . Moreover, the internal vertices of $C_j[w_1, w_2]$ are not on S_i . Note that $S_i[w_1, w_2]$ is disjoint from P . As $|x - i|, |y - i| \geq f(k, t - 1) + 5$, Lemma 3.15 implies that a segment of P_b different from S_i intersects $C_j[w_1, w_2]$. Let z be the vertex of $C_j[w_1, w_2]$ closest to w_1 (but different from w_1) that is in P_b . As the internal vertices of $C_j[w_1, w_2]$ are not on S_i , vertex z is not on S_i . Now z sees w_1 from inside, hence z is on S'_{i+1} . \lrcorner

We would like to say that a subpath of some C_j between S'_1 and S'_s intersects no other segment of P_b than S_1, \dots, S_s . This is not completely trivial, as we can use the third property of Claim 3.28 for a vertex of S_i only if we show that it is on S'_i as well. Claims 3.30–3.32 provide such paths.

Claim 3.30. *Let x be a vertex of $S'_1 \cap C_j$ for some $j \geq h^* + 2sm$. Then for every $1 \leq i \leq s$, there is an undirected path W_i from x to a vertex $y \in S'_i$ such that*

- W_i is between S_1 and S_i ,
- W_i is between C_{j-2im} and C_{j+2im} ,
- W_i does not intersect any segment of P_b different from S_1, \dots, S_s .

Proof. The proof is by induction on i . A path W_1 consisting of only $x = y_1$ shows that the statement is true for $i = 1$. Suppose that y_i is on C_{j_i} . Let q be a vertex of S'_i either on C_{j_i-2m} or C_{j_i+2m} such that $S'_i[y_i, q]$ has no internal vertex on either C_{j_i-2m} or C_{j_i+2m} (as the endpoints of S'_i are on C_{h^*-im} and $j_i - 2m \geq j - 2im - 2m \geq h^* + 2sm - 2im - 2m \geq h^* - im$ holds, such a vertex q has to exist). Suppose therefore that q is on $C_{j_{i+1}}$, where j_{i+1} is either $j_i - 2m$ or $j_i + 2m$. Let $Q_i = S'_i[y_i, q]$. By Claim 3.29, there is a vertex z of Q_i on $C_{j_{i+1}}$ seen from inside by a vertex $y_{i+1} \in S'_{i+1}$. Appending $Q_i[y_i, z]$ and the subpath of $C_{j_{i+1}}$ between z and y_{i+1} to the path W_i gives the required undirected path W_{i+1} ending at y_{i+1} . If path W_i is between C_{j-2im} and C_{j+2im} , then W_{i+1} is between $C_{j-2(i+1)m}$ and $C_{j+2(i+1)m}$. \lrcorner

Claim 3.31. *For some $j \geq h^* + 4sm + 1$, let $C^* = C_j[v_1, v_2]$ be a subpath between S_1 and S_s with $v_1 \in S'_1$, $v_2 \in S_s$, and having no internal vertex on S_1 or S_s . Then S_1, \dots, S_s are the only segments of P_b intersecting C^* .*

Proof. Let S be a segment intersecting C^* , which means that S is enclosed by S_1 . The path C^* splits the area between S_1 and S_s into two regions. At least one of these two regions contains an endpoint of segment S ; let R^* be such a region and let S' be a subpath of S in this region between C_0 and C^* . Let R^* be enclosed by subpaths S_1^* of S_1 , subpath S_s^* of S_s , subpath C_0^* of C_0 , and C^* .

Let u be a vertex of S_1^* on $C_{j-2sm-1}$ closest to v_1 ; as $v_1 \in S'_1$ and $j - 2sm - 1 \geq h^* + 2sm$, we have $u \in S'_1$. Let W_s be the undirected path given by Claim 3.30. As path W_s connects u to S_s , contained between S_1 and S_s , and contained also between $C_{j-4sm-1}$ and C_{j-1} , it cannot intersect C^* and hence it is in the region R^* . Now path W_s separates C^* and C_0^* in R^* , thus W_s intersects S' . Since S_1, \dots, S_s are the only segments of P_s that W_s intersects, it follows that S is one of these segments. \lrcorner

Claim 3.32. *Let Q be a subpath of S'_1 from a vertex of C_x to a vertex of C_y . Suppose that $x + 2sm < j < y - 2sm$ and $j \geq h^* + 4sm + 1$ holds for some j . Then there is a vertex $\alpha \in Q \cap C_j$ and vertex $\beta \in S_s \cap C_j$ such that the subpath $C_j[\alpha, \beta]$ is between S_1 and S_s , has no internal vertex on S_1 and S_s , and intersects no segment of P_b other than S_1, \dots, S_s .*

Proof. Consider the subpath $C_j[w_1, w_2]$ given by Lemma 3.13 with $w_1, w_2 \in S'_1$. If it intersects S_s , then let β be the vertex of S_s closest to $\alpha := w_1$ on this subpath and we are done by Claim 3.31. Otherwise, we arrive to a contradiction as follows. The path $S_1[w_1, w_2]$ contains an endpoint of Q , hence it contains a vertex z that is either on C_x or on C_y . Applying Claim 3.30 on this vertex z gives a path W_s from z to S_s and enclosed by S_1 . Path W_s cannot intersect $C_j[w_1, w_2]$ as $|x - j|, |y - j| > 2sm$. Thus W_s is enclosed by the cycle $S_1[w_1, w_2] \cup C_j[w_1, w_2]$. As W has an endpoint on S_s , this contradicts the assumption that $C_j[w_1, w_2]$ does not intersect S_s . \lrcorner

We are now ready to find the area required by Lemma 3.17, where all the segments are nested. Let $j_i = (h^* + 4sm + 1) + 2smi$ for $i = 0, \dots, 8$ (note that $j_i \leq h^* + 20sm + 1 < h - m$ for every such j_i). As S'_1 is an $(h^* - m)$ -segment and intersects C_{h-m} , we can choose vertices v_0, \dots, v_8 appearing on S'_1 in this order such that v_j is on C_{j_i} . For $q = 1, 3, 5, 7$, Claim 3.32 gives a subpath




Figure 32: Claim 3.33: The segments C_1^* , C_3^* , C_5^* , C_7^* connecting S_1 and S_s , and two additional segments that intersect C_3^* or C_5^* .

$C_q^* := C_{j_q}[\alpha_q, \beta_q]$ between S_1 and S_s with α_q being an internal vertex of $S'_1[v_{q-1}, v_{q+1}]$. This means that C_1^* , C_3^* , C_5^* , C_7^* are distinct and connect S_1 and S_s in this order.

Claim 3.33. *Every segment S intersecting C_3^* or C_5^* is nested between S_1 and S_s .*

Proof. By Claim 3.32, if S is a segment of P_b , then it is one of S_1, \dots, S_s , hence the claim is certainly true. Consider now a segment S of $P_{b'}$ for some $b' \neq b$ and let Q be a subpath of S from one of its endpoints on C_0 to a vertex z of C_3^* or C_5^* . Clearly, Q has to intersect either C_1^* or C_7^* (see Figure 32). Let us assume that S intersects C_1^* (the case when S intersects C_7^* is similar).

Assume by contradiction that S is not nested between S_1 and S_s . Let w_1 and w_2 be the endpoints of C_1^* ; clearly, S does not enclose w_1 and w_2 (as they are on S_1 and S_s , respectively). Let us use Lemma 3.8, on the line $C_1^*[w_1, w_2]$, on the cycle formed by S and the subpath of C_0 connecting the endpoints of S , and on the subpath Q of S connecting C_0 and z while intersecting C_1^* . We get a subpath $S[w_1, w_2]$ containing z such that $w_1, w_2 \in C_1^*$ and $C_1^*[w_1, w_2]$ contains no vertex of S . Let us observe that the internal vertices of $C_1^*[w_1, w_2]$ are not on any segment of P_b : any such segment would be enclosed by S , thus it is in contradiction with Claim 3.32, which states that this segment has to be one of S_1, \dots, S_s . Since $S[w_1, w_2]$ contains z , which is on C_3^* or C_5^* , the path $S[w_1, w_2]$ intersects both C_1^* and C_3^* . By Lemma 3.11, there is an m -bend $(S[w_1, w_2], C'_0, \dots, C'_m)$, with $C'_0 = C_1^*[w_1, w_2]$. The type of this bend is at most $t - 1$ (as both $S[w_1, w_2]$ and $C^*[w_1, w_2]$ are disjoint from P_b), which contradicts $m > f(k, t - 1)$. \lrcorner

We have shown that the conditions of Lemma 3.17 hold for $S^1 := S_1$, $S^2 := S_s$, $C_x[x^1, x^2] := C_3^*$, and $C_y[y^1, y^2] := C_5^*$. Thus there are less than $2^k = s$ segments nested between S_1 and S_s , contradicting the existence of the sequence S_1, \dots, S_s . \square

4 Min-max theorems for paths, cycles, and cuts

4.1 Framework

In this section we consider graphs embedded on surfaces. By abusing the notation, we identify the graph with its image in the embedding.

Definition 4.1. Let G be a digraph embedded on a compact surface Σ . A directed curve N on Σ is called *non-degenerate*, if N intersects embedding of G in a finite number of points and for each intersection point x of N and G there exists a neighbourhood U_x of x such that $N \cap U_x$ separates some non-empty subsets of $(G \cap U_x) \setminus N$ in U_x . For a non-degenerate curve N , let

- $\vec{\nu}(N) = (x_1, x_2, \dots, x_p)$ be the sequence of vertices and edges through which N passes, in their order of appearance on N ;
- $\vec{\mu}(N) = (S_1, S_2, \dots, S_p)$ be the sequence of subsets of $\{-1, +1\}$, defined as follows:
 - if x_i is a vertex, then $-1 \in S_i$ if there exists an arc entering x_i from the left of N and an arc leaving x_i to the right of N , and $+1 \in S_i$ if there exists an arc entering x_i from the right of N and an arc leaving x_i to the left of N ;
 - if x_i is an edge, then $S_i = \{-1\}$ if the x_i traverses N from the left to right, and $S_i = \{+1\}$ if x_i traverses N from right to left.

Intuitively a directed curve is non-degenerate if it does not touch an edge (or vertex) and return to the same face. We point out that a non-degenerate curve is not necessarily non-self-crossing, it is just a smooth, regular image of an interval $[0, 1]$. A curve is called *simple* if it visits every vertex, arc, and face of G at most once; observe that every simple curve is also non-degenerate. We often consider closed curves, that is, smooth and regular images of a circle, and call such a curve a *noose*. The sequences $\vec{\nu}$ and $\vec{\mu}$ are defined in the same manner in this situation; note that they are unique modulo cyclic shifts.

From now on we assume that all the considered curves are non-degenerate (with a single exception of spiral cuts defined in Section 6); hence we ignore stating this attribute explicitly.

When we consider a curve or a noose in our algorithms, we may represent it as a sequence consisting of alternately vertices or edges and faces which the curve traverses. However, for some proofs it will be useful to imagine the curve as an actual topological object being an image of a circle or a closed interval.

Definition 4.2. For a sequence (S_1, S_2, \dots, S_p) where $S_i \subseteq \{-1, +1\}$ we say that a sequence (s_1, \dots, s_q) is *embeddable* into (S_1, S_2, \dots, S_p) if there exists an increasing function $\iota : [q] \rightarrow [p]$ such that $s_i \in S_{\iota(i)}$. In this case, function ι is called an *embedding*.

We are now ready to state the following Theorem of Ding, Schrijver, and Seymour [6].

Theorem 4.3 ([6]). *Let G be a digraph embedded on a torus Σ , and let C_1, C_2, \dots, C_k be closed, non-crossing directed curves on Σ of homotopies $\{(0, c_i)\}$, where $c_i = \pm 1$, located in this order on the torus. Then one can find vertex-disjoint directed cycles D_1, D_2, \dots, D_k in G homotopic to C_1, C_2, \dots, C_k if and only if there does not exist a noose N with the following property: if (p, q) is the homotopy of N , then $p \geq 0$ and no cyclic shift of $(c_1, c_2, \dots, c_k)^p$ is embeddable into $\vec{\mu}(N)$.*

Let us remark that by *homotopic* we mean that there exists a *continuous shift* of Σ that simultaneously shifts all the cycles C_1, C_2, \dots, C_k to D_1, D_2, \dots, D_k , i.e., a continuous map $h :$

$\Sigma \times [0, 1] \rightarrow \Sigma$ such that $\Sigma(\cdot, 0)$ is identity and $\sigma(\cdot, 1)$ maps every cycles C_i to the corresponding cycle D_i .

In the original paper of Ding, Schrijver, and Seymour [6], the noose N was required only to traverse vertices and faces, that is, it was a face-vertex curve. However, we may relax the requirement on the curve to just being non-degenerate, as after relaxation the assumed object still remains a counterexample, and this relaxation simplifies notation in the further parts of the article.

Let G be a digraph embedded into a ring R with outer face C_1 and inner face C_2 . We choose an arbitrary curve W in R that (i) avoids $V(G)$, (ii) crosses $E(G)$ in a finite number of points, (iii) connects C_1 and C_2 , and (iv) is directed from C_1 to C_2 , as the *reference curve*. For every path P connecting C_1 and C_2 in G , by $W(P)$, the *winding number* of P , we denote the number of signed crossings of W defined in the following manner: we traverse P in the direction from C_1 to C_2 and each time we intersect W , if W crosses from left to right (with respect to the directed curve W) then this contributes $+1$ to the number of crossings, and from right to left we count -1 crossing. Note that even if P is in fact a directed path directed from C_2 to C_1 , we count the crossings by traversing from C_1 to C_2 .

If a ring is equipped with a reference curve, we call it a *rooted ring*. We would like to stress here that the reference curve is required to be fully embedded into the ring, i.e., between cycles C_1 and C_2 . This property will be (implicitly) important to many claims, so we will always make sure that the constructed reference curves have this property.

Let $\mathcal{P} = (P_1, P_2, \dots, P_k)$ be a family of directed paths connecting C_1 and C_2 . For a sequence $\mathcal{C} = (c_1, c_2, \dots, c_k)$, where each c_i is ± 1 , we say that \mathcal{P} is *compatible* with \mathcal{C} if for each index i , path P_i is directed from C_1 to C_2 if $c_i = +1$, and from C_2 to C_1 otherwise. Let us note the following observation.

Observation 4.4. *Let G be a digraph embedded into a rooted ring with outer face C_1 and inner face C_2 . Let s_1, s_2, \dots, s_k be terminals lying on the outer face in clockwise order, and let t_1, t_2, \dots, t_k be terminals lying on the inner face in clockwise order. Let us fix a sequence $\mathcal{C} = (c_1, c_2, \dots, c_k)$, consisting of entries ± 1 . Let $\mathcal{P} = (P_1, P_2, \dots, P_k)$ be a family of directed vertex-disjoint paths connecting corresponding s_i with t_i , compatible with \mathcal{C} . Then the winding numbers of paths P_i differ by at most 1.*

For such family of paths $\mathcal{P} = (P_1, P_2, \dots, P_k)$, we denote $W(\mathcal{P}) := W(P_1)$. Note that for every i we have that $|W(P_i) - W(\mathcal{P})| \leq 1$.

4.2 Understanding homotopies in a ring

We use Theorem 4.3 to prove the following result. Intuitively, it says that the possible homotopies of families of paths crossing a ring behave roughly in a convex way.

Theorem 4.5. *Let G be a digraph embedded into a rooted ring with outer face C_1 and inner face C_2 . Let s_1, s_2, \dots, s_k be terminals lying on the outer face in clockwise order and t_1, t_2, \dots, t_k be terminals lying on the inner face in clockwise order; assume further that the reference curve W connects the interval between s_k and s_1 on the boundary of C_1 and the interval between t_k and t_1 on the boundary of C_2 . Let us fix a sequence $\mathcal{C} = (c_1, c_2, \dots, c_k)$, consisting of entries ± 1 . Let $\mathcal{P} = (P_1, P_2, \dots, P_k)$ be a family of directed vertex-disjoint paths connecting corresponding s_i with t_i , compatible with \mathcal{C} . Assume further that there are some families $\mathcal{Q} = (Q_1, Q_2, \dots, Q_k)$ and $\mathcal{R} = (R_1, R_2, \dots, R_k)$ of directed vertex-disjoint paths connecting C_1 and C_2 that are also compatible with \mathcal{C} . Then, for every number α such that $W(\mathcal{Q}) + 6 \leq \alpha \leq W(\mathcal{R}) - 6$, there exists a family $\mathcal{P}' = (P'_1, P'_2, \dots, P'_k)$ such that \mathcal{P}' is compatible with \mathcal{C} , P'_i connects s_i and t_i , and $W(\mathcal{P}') = \alpha$.*

Proof. By the assumed property of the reference curve W we have that for every sequence \mathcal{A} of vertex-disjoint paths connecting s_1, s_2, \dots, s_k with corresponding t_1, t_2, \dots, t_k , all the paths of \mathcal{A} have the same winding number, equal to $W(\mathcal{A})$. Let us change the curve W to any curve W' that has the same endpoints as W , but has winding number α with respect to W . Then, winding numbers with respect to W' are exactly the same as with respect to W , but with $-\alpha$ additive constant. Therefore, without loss of generality we may assume that $\alpha = 0$.

Compactify the plane using one point in the infinity. Take any disk D_1 outside C_1 , any disk D_2 inside C_2 , and cut them out of the compacted plane. Then identify boundaries of D_1 and D_2 , thus obtaining a torus Σ . For every index i , insert an arc (t_i, s_i) if $c_i = +1$, and an arc (s_i, t_i) if $c_i = -1$. Observe that these new arcs can be realized on Σ without introducing crossings by drawing them through the identified boundaries of D_1 and D_2 . Let us denote the new graph by \overline{G} ; we will also refer to it as to *extended G*. For $i = 1, 2, \dots, k$, let the face F_i be the face of \overline{G} enclosed by arcs between s_i, t_i , between s_{i+1}, t_{i+1} , and fragments of boundaries of C_1, C_2 between terminals s_i, s_{i+1} and t_i, t_{i+1} , respectively, where $s_{k+1} = s_1$ and $t_{k+1} = t_1$. Moreover, we can extend the curve W to a closed curve \overline{W} on Σ by closing it through the face F_k . We can set the homotopy group on Σ so that \overline{W} has homotopy $(0, 1)$ and $\overline{D} := \partial D_1 = \partial D_2$ has homotopy $(1, 0)$.




Figure 33: Initial situation in the proof of Theorem 4.5. Paths of the family $\overline{\mathcal{P}}$ are coloured blue, and paths of the family \mathcal{Q} are coloured red. Family \mathcal{R} is not depicted in order not to overcomplicate the picture. The cycles in the family $\overline{\mathcal{P}}$ already have homotopies $(0, c_i)$ (i.e., $\beta = 0$), hence they may serve as $\overline{\mathcal{P}'}$.

Similarly, using new arcs connecting s_i and t_i we can close each P_i into a directed cycle \overline{P}_i , thus obtaining a family of vertex-disjoint cycles $\overline{\mathcal{P}}$. Note that a cycle \overline{P}_i has homotopy $(c_i \cdot \beta, c_i)$

for $\beta = W(\mathcal{P})$. Also, the postulated family of paths \mathcal{P}' exists if and only if there exists the same family closed by the new arcs connecting s_i and t_i ; that is, in the theorem statement we postulate existence of a family $\overline{\mathcal{P}'} = (\overline{P}'_1, \overline{P}'_2, \dots, \overline{P}'_k)$ of vertex-disjoint cycles on Σ , located in this order, such that \overline{P}'_i has homotopy $(0, c_i)$.

For the sake of contradiction assume that the assumed family $\overline{\mathcal{P}'}$ does not exist. By Theorem 4.3 there exists a noose N of some homotopy (p, q) , where $p \geq 0$, such that $(c_1, c_2, \dots, c_k)^p$ is not embeddable in $\vec{\mu}(N)$.

We now construct an infinite graph \widehat{G} , called also *unravalled G*, by taking the torus Σ , cutting it along \overline{W} and \overline{D} thus obtaining a rectangle, and filling the plane with copies of this rectangle in a grid-like manner. The halves of arcs cut when cutting \overline{W} and \overline{D} are joined with the corresponding halves from a neighbouring rectangle. The rectangles are called *cells*; we index them naturally by $\mathbb{Z} \times \mathbb{Z}$, where the first coordinate corresponds to the direction of \overline{D} , and the second to the direction of \overline{W} . Let us locate the cells on the plane in such a manner that a cell (a, b) is the square with coordinates of corners $(a - 1, b - 1), (a - 1, b), (a, b), (a, b - 1)$. By $F_i^{(a,b)}$ we denote the face of \widehat{G} originating in the face F_i of \overline{G} that lies between cells (a, b) and $(a, b + 1)$; the copy of the face F_k that lies on the meet of cells $(a, b), (a + 1, b), (a + 1, b + 1)$ and $(a, b + 1)$ has index (a, b) . Let \mathcal{F} be the set of all the faces of the form $F_i^{(a,b)}$. Note that then a line $y = c$ for $c \in \mathbb{Z}$ traverses all the faces of \mathcal{F} with the second coordinate of the upper index equal to c ; moreover, no vertex of \widehat{G} has integer second coordinate.

Let us define a family $\widehat{\mathcal{P}}$ as an infinite family of parallel vertex-disjoint infinite paths that are results of unravelling the cycles of $\overline{\mathcal{P}}$ on the plane. Also, we can define unravalled families $\widehat{\mathcal{Q}}$ and $\widehat{\mathcal{R}}$ as infinite families of vertex-disjoint paths originating in unravelling of \mathcal{Q} and \mathcal{R} , respectively. Note that each path from $\widehat{\mathcal{Q}}$ or $\widehat{\mathcal{R}}$ is finite and contained in one row of the grid.

Let us elaborate a bit on the family $\widehat{\mathcal{P}}$. Paths of $\widehat{\mathcal{P}}$ divide the plane into infinite number of strips, each with two ends and bounded by two paths of $\widehat{\mathcal{P}}$. The strips can be labeled by integer numbers so that each strip is neighbouring only with the strips with numbers differing by one. Observe that every path of $\widehat{\mathcal{P}}$ traverses every line $y = c$ for $c \in \mathbb{Z}$ exactly once, using one of the arcs between neighbouring faces of \mathcal{F} . Thus, $\widehat{\mathcal{P}}$ divides each line $y = c$ for $c \in \mathbb{Z}$ into intervals, corresponding to intersections of this line with strips. We call these intervals *strip intervals*.

Moreover, observe that the family $\widehat{\mathcal{P}}$ is invariant with respect to integer shifts, i.e., for any $(a, b) \in \mathbb{Z} \times \mathbb{Z}$ we have that $\widehat{\mathcal{P}} + (a, b) = \widehat{\mathcal{P}}$. Observe also that a shift $(a, 0)$ acts on a line $y = c$ for $c \in \mathbb{Z}$ by shifting the strip intervals $k \cdot a$ to the right. As strip intervals are intersections of strips with this line, shift $(a, 0)$ acts on the plane divided into strips by mapping each strip into the strip $k \cdot a$ to the right, in the order of strips.

We now move to families $\widehat{\mathcal{Q}}$ and $\widehat{\mathcal{R}}$. Take each path Q from $\widehat{\mathcal{Q}}$ contained in $\{y \in [c, c + 1]\}$ and extend it by short curves contained in single faces of \mathcal{F} at which Q starts and ends, so that Q connects the line $\{y = c\}$ with the line $\{y = c + 1\}$ (recall that each path in $\widehat{\mathcal{Q}}$ connects C_1 with C_2 , not D_1 with D_2). The original part of Q , being its intersection with \widehat{G} , is called the *core* of Q , and the two small curves at the ends are called *extensions* of Q . Clearly, we may add the extensions in such a manner that (i) curves originating in the same path from \mathcal{Q} are extended in identical manner, and (ii) all the curves from $\widehat{\mathcal{Q}}$ are still pairwise disjoint. Perform the same construction for the family $\widehat{\mathcal{R}}$.

Observe that families $\widehat{\mathcal{Q}}$ and $\widehat{\mathcal{R}}$ have a similar periodic behaviour inside parts of the plane $\{y \in [c, c + 1]\}$ for $c \in \mathbb{Z}$. Curves from $\widehat{\mathcal{Q}}$ ($\widehat{\mathcal{R}}$) divide each $\{y \in [c, c + 1]\}$ into small strips, naturally linearly ordered along $\{y \in [c, c + 1]\}$. The partition is invariant with respect to shifts by vectors $(a, 0)$ for $a \in \mathbb{Z}$: each such shift maps every strip to the strip that is $a \cdot k$ to the right of it, counting in the natural order of the strips.

Let us now examine what happens with the noose N . Clearly, N also unravels into an infinite family of infinite-length parallel curves (see Fig. 34). We take one of them and declare it \widehat{N} , the unravelled noose N . By slightly perturbing N we may assume that intersections of \widehat{N} with lines $\{y = c\}$ for $c \in \mathbb{Z}$ are not contained in \widehat{G} .

Note that \widehat{N} is periodic in the following sense: $\widehat{N} + (p, q) = \widehat{N}$. If we take a fragment of \widehat{N} between any $x \in \widehat{N} \setminus \widehat{G}$ and $x + (p, q) \in \widehat{N} \setminus \widehat{G}$, denote it by $\widehat{N}[x]$, then this corresponds to the closed curve N on Σ cut in the vertex x . In particular, no cyclic shift of \mathcal{C}^p can be embedded into $\vec{\mu}(\widehat{N}[x])$.

Assume now that $q = 0$, that is, N does not wind in the direction of \overline{W} . Take any $x \in \widehat{N} \setminus \widehat{G}$ and consider the curve $\widehat{N}[x]$. By the previous observations, $x + (p, 0)$ is located in the strip $k \cdot p$ to the right with respect to the strip containing x . It follows that any noose travelling from x to $x + (p, 0)$ must cross the $k \cdot p$ consecutive paths from \widehat{P} between these strips; this applies in particular to $\widehat{N}[x]$. We may now define an embedding of a cyclic shift of \mathcal{C}^p into $\vec{\mu}(\widehat{N}[x])$, by embedding a ± 1 into each crossing of these paths with $\widehat{N}[x]$, where we choose the sign depending on the direction in which a corresponding path is directed. This is a contradiction with the assumed properties of N .

Hence, we infer that $q \neq 0$. For the rest of the proof we assume that $q > 0$ and we will use only the family \widehat{Q} . The proof for $q < 0$ is symmetric and uses only family \widehat{R} .

Let us perform another slight modification to \widehat{N} . By slightly perturbing \widehat{N} within faces it travels through, we may assume that whenever \widehat{N} travels through a face $F \in \mathcal{F}$ (say, lying on a line $\{y = c\}$), it firstly touches the line $\{y = c\}$, then crosses all the extensions of paths from \widehat{Q} , then again touches the line $\{y = c\}$, and then finally leaves the face F . In other words, we may assume that before the first intersection and after the last intersection with the line $\{y = c\}$, no extension of a path from \widehat{Q} is crossed. This purely technical property will be used later to avoid counting crossings of extensions of paths from \widehat{Q} as true crossings contributing to $\vec{\nu}(N)$.

We now need the following two claims.

Claim 4.6. *Let M be a curve with an origin in a point $(x_1, c) \notin \widehat{G}$ and end in a point $(x_2, c) \notin \widehat{G}$, where $c \in \mathbb{Z}$. If $r = \max(0, \lfloor x_2 \rfloor - \lfloor x_1 \rfloor)$, then \mathcal{C}^r is embeddable in $\vec{\mu}(M)$.*

Proof. If $r = 0$ then the claim is trivial, so assume that $r > 0$. As (x_1, c) is located in the interval $(\lfloor x_1 \rfloor - 1, \lfloor x_1 \rfloor]$ on the line $y = c$ and (x_2, c) is located in the interval $[\lfloor x_2 \rfloor, \lfloor x_2 \rfloor + 1)$ on the same line, it follows that one can find r consecutive bundles of paths from \widehat{P} corresponding to cycles $\overline{P}_1, \overline{P}_2, \dots, \overline{P}_k$, separating (x_1, c) from (x_2, c) . Each such bundle is formed by paths of \widehat{P} that intersect an interval $[\lambda, \lambda + 1)$ on the line $\{y = c\}$ for $\lfloor x_1 \rfloor \leq \lambda < \lfloor x_2 \rfloor$, $\lambda \in \mathbb{Z}$. Similarly to the proof for the case $q = 0$, crossings of M with these paths define an embedding of \mathcal{C}^r into $\vec{\mu}(M)$. \square

Claim 4.7. *Let M be a curve contained in the part of the plane $\{y \in [c, c + 1]\}$, where the origin is of the form (x_1, c) and the end is of the form $(x_2, c + 1)$. Moreover, assume that M does not cross any extension of any path from \widehat{Q} . If $r = \max(-4, \lfloor x_2 \rfloor - \lfloor x_1 \rfloor)$, then \mathcal{C}^{r+4} is embeddable in $\vec{\mu}(M)$.*

Proof. If $r = -4$ then the claim is trivial, so assume that $r > -4$. We examine the mutual location of (x_1, c) and $(x_2, c + 1)$ with respect to the strips into which \widehat{Q} divides $\{y \in [c, c + 1]\}$. We claim that the strip containing (x_2, c) is at least $(r + 5) \cdot k$ strips to the right with respect to the one containing (x_1, c) . By the properties of \widehat{Q} with respect to shifts, we have that $(x_1 + r, c)$ is $r \cdot k$ strips to the right when compared with (x_1, c) ; note that possibly $r < 0$ — then it means $-r \cdot k$ strips to the left. Now note that $\lfloor x_2 \rfloor \geq x_1 + r$, so $(\lfloor x_2 \rfloor, c)$ is at least $r \cdot k$ strips to the right when compared with (x_1, c) . Observe that $(\lfloor x_2 \rfloor, c + 1)$ is at least $5 \cdot k$ strips to the right when compared to $(\lfloor x_2 \rfloor, c)$. This follows from the fact that the the segment between points $(\lfloor x_2 \rfloor, c)$ and

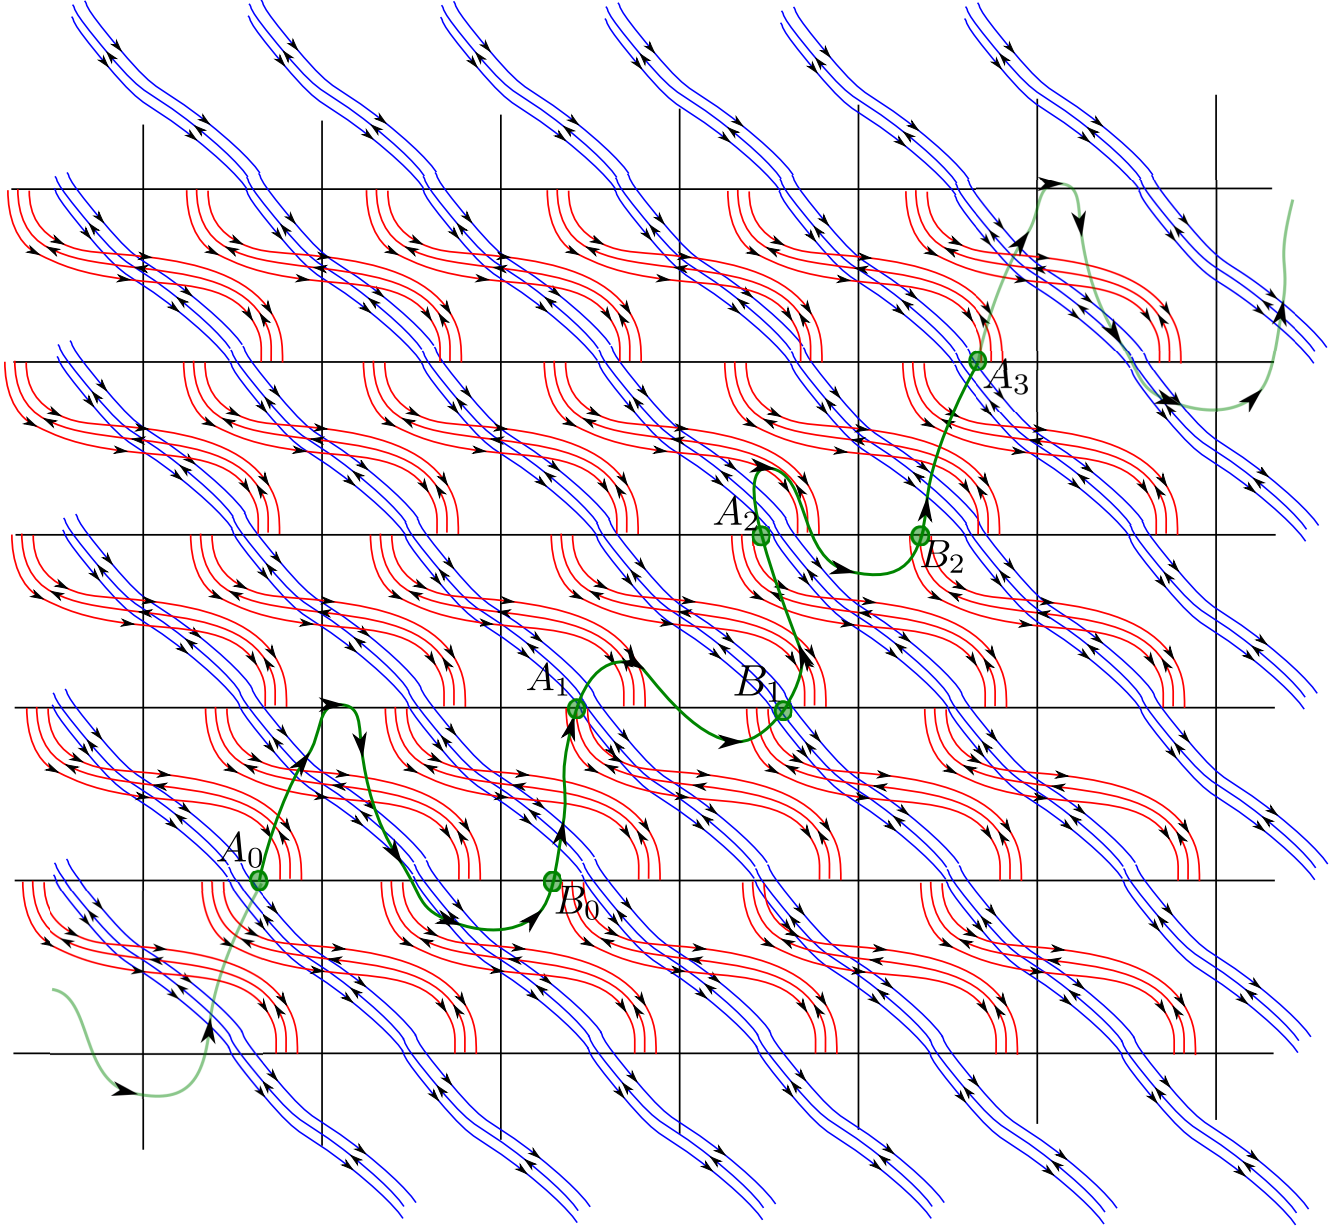


Figure 34: The unraveled graph \widehat{G} with families $\widehat{\mathcal{P}}$ and $\widehat{\mathcal{Q}}$ depicted (blue and red, respectively), where $\beta = -1$ and $(p, q) = (4, 3)$. On the curve \widehat{N} , the interval $\widehat{N}[x]$ between x and $x + (4, 3)$ is highlighted, and partitioned into parts K_i, L_i by points $x = A_0, B_0, A_1, B_1, A_2, B_2, A_3 = x + (4, 3)$. For the sake of simplicity, paths from the family \mathcal{Q} have winding numbers -1 instead of required at most -6 .

$(\lfloor x_2 \rfloor, c + 1)$ is exactly the line W , and as $W(\mathcal{Q}) \leq -6$, then by Observation 4.4 each curve of \mathcal{Q} has signed crossing number at least 5 with respect to W . This means that during its travel from $(\lfloor x_2 \rfloor, c)$ to $(\lfloor x_2 \rfloor, c + 1)$ we travel at least $5 \cdot k$ strips to the right in the partition of $\{y \in [c, c + 1]\}$ into the strips by $\widehat{\mathcal{Q}}$. As $\lfloor x_2 \rfloor \leq x_2$, point $(x_2, c + 1)$ can be only even further to the right when compared with $(\lfloor x_2 \rfloor, c + 1)$

Hence, in total $(x_2, c+1)$ lies at least $(r+5) \cdot k$ strips to the right with respect to (x_1, c) . As in the previous proofs, it follows that every curve M travelling from (x_1, c) to $(x_2, c+1)$, contained in $\{y \in [c, c+1]\}$ and not crossing any extensions of paths from $\widehat{\mathcal{Q}}$, must admit an embedding of a cyclic shift of $\mathcal{C}^{(r+5) \cdot k}$ into $\vec{\nu}(M)$. As every cyclic shift of $\mathcal{C}^{(r+5) \cdot k}$ has $\mathcal{C}^{(r+4) \cdot k}$ as a subword, we have that $\mathcal{C}^{(r+4) \cdot k}$ is embeddable into $\vec{\mu}(M)$. \square

Now observe that

$$\begin{aligned} \max(0, \lfloor x_2 \rfloor - \lceil x_1 \rceil) &\geq x_2 - x_1 - 2, \\ \max(-4, \lfloor x_2 \rfloor - \lceil x_1 \rceil) + 4 &\geq x_2 - x_1 + 2. \end{aligned}$$

Let us now examine the curve \widehat{N} . As \widehat{N} is invariant under the shift (p, q) and $q \neq 0$, it follows that the set of intersections of \widehat{N} with $\{y = 0\}$ is nonempty and contained in a compact interval on \widehat{N} . Let then x be the first intersection of \widehat{N} with $\{y = 0\}$. Consider the curve $\widehat{N}[x]$; similarly as in the case $q = 0$, we have that $\vec{\mu}(\widehat{N}[x])$ does not admit embedding of any cyclic shift of \mathcal{C}^p . By the choice of x and invariance of \widehat{N} under the shift (p, q) we have that $x + (p, q)$ is the first intersection of \widehat{N} with the line $\{y = q\}$.

We would like now to divide $\widehat{N}[x]$ into smaller curves. Let $A_0 = x$. We inductively define points $B_0, A_1, B_1, \dots, B_{q-1}, A_q$ as follows: B_i is the last intersection of the part of $\widehat{N}[x]$ after A_i with line $\{y = i\}$, while A_{i+1} is the first intersection of the part of $\widehat{N}[x]$ after B_i with the line $\{y = i+1\}$. As $x + (p, q)$ is the first intersection of \widehat{N} with $\{y = q\}$, it follows that $A_q = x + (p, q)$. Let $A_i = (a_i, i)$ and $B_i = (b_i, i)$. For $i = 0, 1, \dots, q-1$ let K_i be the part of $\widehat{N}[x]$ between A_i and B_i , and let L_i be the part of $\widehat{N}[x]$ between A_i and B_{i+1} .

Therefore, $\widehat{N}[x]$ consists of curves $K_0, L_0, K_1, L_1, \dots, K_{q-1}, L_{q-1}$ concatenated in this order. Each curve K_i satisfies the conditions of Claim 4.6, while each curve L_i satisfies the conditions of Claim 4.7; note that the claim for curves L_i follows from the technical property about visits of \widehat{N} in faces of \mathcal{F} that we ensured before introducing Claims 4.6 and 4.7. By applying these two claims we have that $\vec{\mu}(\widehat{N}[x])$ admits an embedding of \mathcal{C}^t , where

$$t \geq \sum_{i=0}^{q-1} (b_i - a_i - 2) + \sum_{i=0}^{q-1} (a_{i+1} - b_i + 2) = a_q - a_0 = q$$

This is a contradiction with the assumed properties of N . \square

We now provide two corollaries of Theorem 4.5, which will be used in our algorithm.

Lemma 4.8 (Rerouting in a ring). *Let G be a digraph in a rooted ring with outer face C_1 and inner face C_2 , where W is the reference curve.*

- Let p_1^1, \dots, p_s^1 be vertices on C_1 (clockwise order) and let p_1^2, \dots, p_s^2 be vertices on C_2 (clockwise order).
- Let q_1^1, \dots, q_s^1 be vertices on C_1 (clockwise order) and let q_1^2, \dots, q_s^2 be vertices on C_2 (clockwise order).
- Let P_1, \dots, P_s be a set of pairwise vertex-disjoint paths contained between C_1 and C_2 such that the endpoints of P_i are p_i^1 and p_i^2 .
- Let Q_1, \dots, Q_s be a set of pairwise vertex-disjoint paths contained between C_1 and C_2 such that the endpoints of Q_i are q_i^1 and q_i^2 .

- P_i and Q_i go in the same direction, i.e., P_i goes from p_i^1 to p_i^2 if and only if Q_i goes from q_i^1 to q_i^2 .

Then there is set P'_1, \dots, P'_s of pairwise vertex-disjoint paths between C_1 and C_2 such that P'_i and P_i have the same start/end vertices and $|W(P'_i) - W(Q_i)| \leq 6$.

Proof. The Lemma follows from the application of Theorem 4.5 to families

$$\mathcal{P}, \mathcal{Q}, \mathcal{R} = \{P_1, P_2, \dots, P_s\}, \{P_1, P_2, \dots, P_s\}, \{Q_1, Q_2, \dots, Q_s\},$$

or

$$\mathcal{P}, \mathcal{Q}, \mathcal{R} = \{P_1, P_2, \dots, P_s\}, \{Q_1, Q_2, \dots, Q_s\}, \{P_1, P_2, \dots, P_s\},$$

depending on the inequality between $W(\{P_1, P_2, \dots, P_s\})$ and $W(\{Q_1, Q_2, \dots, Q_s\})$. \square

Lemma 4.9 (One-way Spiral Lemma). *Let G be a digraph in a rooted ring with outer face C_1 and inner face C_2 , where W is the reference curve. Assume moreover that W induces a directed path W^* in the dual of G , directed from C_1 to C_2 .*

- Let p_1^1, \dots, p_s^1 be vertices on C_1 (clockwise order) and let p_1^2, \dots, p_s^2 be vertices on C_2 (clockwise order).
- Let q_1^1, \dots, q_s^1 be vertices on C_1 (clockwise order) and let q_1^2, \dots, q_s^2 be vertices on C_2 (clockwise order).
- Let P_1, \dots, P_s be a set of pairwise vertex-disjoint paths contained between C_1 and C_2 such that P_i goes from p_i^1 to p_i^2 .
- Let Q_1, \dots, Q_s be a set of pairwise vertex-disjoint paths contained between C_1 and C_2 such that Q_i goes from q_i^1 to q_i^2 .

Then there is set P'_1, \dots, P'_s of pairwise vertex-disjoint paths between C_1 and C_2 such that P'_i and P_i have the same start/end vertices and $-6 \leq |E(P'_i) \cap E(W^*)| - |E(Q_i) \cap E(W^*)| \leq 6$.

Proof. The Lemma follows from Lemma 4.8 and the assumption that W crosses only arcs directed in one direction, so for any path P the winding number with respect to W is the number of arcs of $E(W^*)$ used by P . \square

4.3 Min-max theorems for alternating cuts

4.3.1 Alternations

We now introduce the language and basic facts about measuring the number of alternations along a curve.

Definition 4.10. An *alternating sequence* is a sequence of entries ± 1 , where $+1$ and -1 appear alternately.

Note that for every p there are two alternating sequences of length p : one starting with $+1$ and one starting with -1 .

Definition 4.11. *Alternation* of sequence A of subsets of $\{-1, +1\}$, denoted $\mathfrak{a}(A)$, is the largest possible length of an alternating sequence embeddable into A . Alternation of a non-closed curve N , denoted $\mathfrak{a}(N)$, is defined as $\mathfrak{a}(\vec{\mu}(N))$, while alternation of a noose N , also denoted $\mathfrak{a}(N)$, is defined as the maximum of $\mathfrak{a}(\vec{\mu}(N))$ over all cyclic shifts of $\vec{\mu}(N)$.

We now define a notion of a *pretty curve* that will be useful for cutting the digraph using small number of alternations. Intuitively, a pretty curve uses as little crossings of vertices as possible.

Definition 4.12. A non-degenerate curve N is called *pretty* if all S_i -s in $\vec{\mu}(N)$ corresponding to crossings of vertices are equal to $\{-1, +1\}$ or \emptyset .

We state now the following observation that shows that pretty curves behave robustly with respect to alternations. We stress that in the following two observations we are considering non-closed curves, not nooses.

Observation 4.13. *Let G be a digraph embedded on a compact orientable surface Σ and let N be a non-degenerate curve on Σ . Then there exists a pretty curve N' that has the same endpoints as N , is homotopic to N and $\mathbf{a}(N') = \mathbf{a}(N)$.*

Proof. We just need to make some simple local modifications to N . Whenever N traverses a vertex v_i with $S_i = \{-1\}$, we have that all arcs adjacent to v_i from the right of N are directed from v_i , or all arcs adjacent to v_i from the left of N are directed from v_i . Then in N' we may either circumvent v_i by traversing it from the left or from the right side, thus traversing all the arcs adjacent to v_i from the left or from the right of N , respectively. Note that the corresponding change in $\vec{\mu}(N')$ is substitution of one term $\{-1\}$ with an arbitrary number of terms $\{-1\}$, which does not change the alternation of the sequence. We perform a symmetric operation for every traversed v_i with $S_i = \{+1\}$. Observe that the modifications performed do not spoil non-degeneracy. \square

We now show that one can conveniently reduce any pretty curve to a simple curve.

Observation 4.14. *Let N be a pretty curve in a graph G embedded on a compact orientable surface Σ . Then there exists a simple curve N' with the same endpoints as N , such that $\mathbf{a}(N') \leq \mathbf{a}(N)$, and N' is constructed by taking disjoint subcurves of N , and appending them in the order of their appearance on N , and possibly some consecutive two using small shortcutting curves, each entirely contained in the interior of one face of G .*

Proof. We build N' by travelling along N , and whenever we encounter a vertex or face visited more than once, we continue further from the last visit, thus omitting the interval on N between the first and the last visit (in case of double visit of a face, we may have to make a shortcut inside the face). N' built in such a manner is simple. Hence, we need to argue that $\mathbf{a}(N') \leq \mathbf{a}(N)$.

Observe that cutting out a segment between two visits of the same face corresponds to taking a curve with new $\vec{\mu}$ being a subsequence of the old one; hence, the alternation cannot increase. The non-trivial part is that when we cut out a segment between the first and the last visit of some vertex v , then the alternation also cannot increase. Let N_0 be the original curve and N'_0 be the curve after cutting out. Let S_f be the set in $\vec{\mu}(N_0)$ corresponding to the first visit of v and S_l be the set in $\vec{\mu}(N_0)$ corresponding to the last visit of v ; by the fact that N is pretty we have that S_f, S_l are equal to $\{+1, -1\}$ or \emptyset . Let S be the set in $\vec{\mu}(N'_0)$ corresponding to the only visit of v after the shortening.

If S_f or S_l is equal to $\{+1, -1\}$, then we have that $S \subseteq S_f \cup S_l$. This suffices for our purposes, as then every embedding of an alternating sequence into a cyclic shift of $\vec{\mu}(N'_0)$ can be lifted to an embedding into an appropriate cyclic shift of $\vec{\mu}(N_0)$ by mapping the term mapped to S in $\vec{\mu}(N'_0)$ either to S_f or to S_l .

If $S_f = S_l = \emptyset$ then v is a sink or source in G , which means that also $S = \emptyset$. Hence $\vec{\mu}(N'_0)$ is a subsequence of $\vec{\mu}(N)$ and the claim follows. \square

We remark that in Observation 4.14 the assumption of N being pretty is necessary. Moreover, the simple curve given by Observation 4.14 is not necessarily pretty; however, it is non-degenerate as it is simple.

4.3.2 Many alternating cycles, or a short cut of a ring

The first min-max theorem follows from the torus theorem of Ding, Schrijver, and Seymour, that is, Theorem 4.3.

Lemma 4.15 (Alternating cycles/cut duality). *Let G be a graph embedded in the plane and let γ_1, γ_2 be two disjoint closed Jordan curves not containing any point of G . Let G_γ be the subgraph of G that consists of the vertices and edges that lie between γ_1 and γ_2 , and let F_1, F_2 be the faces of G_γ that contain γ_1, γ_2 , respectively (it is possible that $F_1 = F_2$). Then, for any positive even integer r , in time polynomial in r and G , one can find either:*

1. a sequence C_1, C_2, \dots, C_r of alternating concentric cycles in G_γ that separate f_1 from f_2 , and C_i being separated from f_2 by C_j for $i < j$; or
2. a simple curve M that starts in F_1 , ends in F_2 , and such that M has alternation at most r .

Proof. Observe that finding a sequence C_1, C_2, \dots, C_r can be performed in polynomial time using the algorithm of Schrijver [28], while finding curve N can be easily done using breadth-first search in polynomial time. Hence, we are left with proving that one of these objects always exists.

Let us perform a similar operation to that from the proof of Theorem 4.5: we compactify the plane using a point in the infinity, cut the plane along γ_1 and γ_2 constraining ourselves to the part between them, and identify γ_1 and γ_2 thus creating a torus Σ . Thus we may imagine that that G is a graph embedded on torus Σ . We choose the homotopy group on Σ so that the identified γ_1, γ_2 have homotopy $(0, 1)$ while an arbitrarily chosen simple curve connecting γ_1 and γ_2 has homotopy $(1, 0)$.

By Theorem 4.3, if there is no feasible sequence C_1, C_2, \dots, C_r , this means that there exists a noose N of homotopy (p, q) for $p > 0$, such that no cyclic shift of $(-1, +1)^{pr/2}$ is embeddable in $\vec{\mu}(N)$ (recall that r is even). As N winds p times in the direction from γ_1 to γ_2 , it follows that one can find p disjoint segments N_1, N_2, \dots, N_p on N such that every N_i travels from γ_1 to γ_2 and does not touch γ_1 or γ_2 apart from the endpoints (is entirely contained in the part of the plane between γ_1 and γ_2). Assume that every N_i has alternation at least $r + 1$, hence some alternating sequence of length $r + 1$ can be embedded into $\vec{\mu}(N)$. It follows that one can embed $(-1, +1)^{r/2}$ into every $\vec{\mu}(N_i)$. Therefore, one can embed $(-1, +1)^{pr/2}$ into $\vec{\mu}(N)$, which is a contradiction with the assumed properties of N . We infer that there exists some $M_0 = N_i$ such that $\vec{\mu}(M_0)$ has alternation at most r . By applying Observation 4.13 and Observation 4.14 to the curve M_0 we obtain the desired curve M . \square

4.3.3 Many alternating paths, or a short circular cut in a ring

The second min-max theorem is considerably more difficult to prove. It is possible to prove it also by the means of Theorem 4.3, but in order to avoid unnecessary technical details, we choose to use the algorithm for cohomology feasibility problem of Schrijver [28].

Let Λ be a free group on r generators g_1, g_2, \dots, g_r . In the cohomology feasibility problem we are given a directed graph G together with some labeling ϕ of arcs with elements of Λ , and a downward-closed set $H(a) \subseteq \Lambda$ for each arc a (i.e., if a word w belongs to $H(a)$, then so does every subsequence of w). We say that a labeling $\psi : E(G) \rightarrow \Lambda$ is *cohomologous* to ϕ if there exists $F : V(G) \rightarrow \Lambda$, such that $\psi((u, v)) = F(u)^{-1} \cdot \phi((u, v)) \cdot F(v)$ for every $(u, v) \in E(G)$. The question asked in the problem is whether there is some ψ cohomologous to ϕ such that $\psi(a) \in H(a)$ for every $a \in E(G)$. The following theorem is the fundamental result standing behind the algorithm of Schrijver [28].

Theorem 4.16 ([28]). *The cohomology feasibility problem is polynomial time solvable, even if sets $H(a)$ are given as polynomial-time oracles that check belonging.*

We describe the instance of the cohomology feasibility problem as a quadruple (G, Λ, ϕ, H) . We remark that the algorithm works also in a more general setting than just for free groups, i.e., for free partially commutative groups, but we do not need this in this work. However, we will use the fact that in the algorithm one can require for an arbitrary subset of vertices $S \subseteq V(G)$ that $F(s) = 1_\Lambda$ for every $s \in S$. In this case, we add set the S as the fifth coordinate of the problem description.

In [27], Schrijver explicitly describes the obstacles for existence of the solution.

Theorem 4.17 ([27], Theorem 3 with adjusted terminology). *Let $I = (G, \Lambda, \phi, H)$ be an instance of the cohomology feasibility problem. Then there is a solution ψ to the instance I if and only if for every vertex u and every two undirected walks P, Q from u to u there exists $x \in \Lambda$ such that $x^{-1} \cdot \phi(P) \cdot x \in H(P)$ and $x^{-1} \cdot \phi(Q) \cdot x \in H(Q)$, where for an undirected walk S , $\phi(S)$ denotes the product of group elements along S , while $H(S)$ denotes the set of all possible such products where the factors are taken from $H(a)$ for consecutive arcs a of S . Moreover, the algorithm of Theorem 4.16 can provide walks P, Q contradicting this assumption, in case no solution was found.*

We remark that in the sense of the above theorem, a walk is undirected, that is, it does not necessarily respect the direction of arcs; it can go via an arc in the reverse direction, and if this is the case, the contribution to $\phi(S)$ and $H(S)$ is the normal contribution reversed.

Let G be a directed graph embedded into a ring R , where C_1 and C_2 are cycles being boundaries of the faces outside and inside the ring, respectively. Contrary to Lemma 4.15, in the following we assume that these two faces are different, and in fact C_1 and C_2 are disjoint. Let r be an even number. We say that a family of vertex-disjoint paths $\mathcal{P} = (P_1, P_2, \dots, P_r)$ is an *alternating join* of C_1 and C_2 of size r , if every path P_i connects C_1 and C_2 and is directed from C_1 to C_2 if i is odd, and from C_2 to C_1 if i is even, and paths P_i are located in clockwise order in the ring.

We now show how to, given a digraph G , construct an instance of the cohomology feasibility problem that encodes existence of an alternating join; the construction closely follows the lines of constructions of Schrijver for various other problems, but we include it for the sake of completeness.

We define an *extended dual* G^+ of G as follows: we construct the classical dual G^* , and for every vertex v and every two faces sharing v and not sharing an edge adjacent to v , we add an additional arc between the corresponding pair of vertices in the dual, in an arbitrary direction; we call these arcs *added*. We delete the vertices in the dual corresponding to faces with boundaries C_1, C_2 , i.e., we delete the inner and outer of the ring. Note that G^+ is not necessarily planar.

We take Λ to be a free group on r generators g_1, g_2, \dots, g_r . For every original arc a of the dual we define $H(a) = \{1, g_1, g_2, \dots, g_r\}$, and for every added arc a^+ we define $H(a^+) = \{1, g_1, g_2, \dots, g_r, g_1^{-1}, g_2^{-1}, \dots, g_r^{-1}\}$. Take any path P in G connecting C_1 and C_2 and put $\phi(a^*) = g_1 \cdot g_2^{-1} \cdot g_3 \cdot \dots \cdot g_{r-1} \cdot g_r^{-1}$ for all arcs a of P , where a^* is the arc in the dual corresponding to a . Moreover, if a^+ is an added arc connecting two faces F_1, F_2 sharing a vertex v , take the boundary of the face in the dual corresponding to the vertex v , define R to be any path on this boundary connecting F_1 and F_2 , and put $\phi(a^+) = \phi(R)$. Note that after deleting the outer and inner face, at least one such path exists. Put $\phi(a) = 1$ for all the other arcs.

Having defined the instance, we prove the following claim.

Lemma 4.18. *Let $I = (G^+, \Lambda, \phi, H)$ be the defined instance of the cohomology feasibility problem. Then I has a solution ψ if and only if there exists an alternating join of C_1 and C_2 of size r . Moreover, given a solution ψ one can construct the alternating join in polynomial time.*




Figure 35: The marked edges belong to the path P . On the right we have a subgraph of the extended dual depicting the values $\phi(a)$, where $g = g_1 \cdot g_2^{-1} \cdot g_3 \cdot \dots \cdot g_{r-1} \cdot g_r^{-1}$.

Proof. Assume first that an alternating join $\mathcal{P} = (P_1, P_2, \dots, P_r)$ exists. Put $\psi(a^*) = g_i$ if $a \in P_i$, $\psi(a^*) = 1$ for every other arc of the dual, and $\psi(a^+) = \psi(R)$ for every added arc, where R is defined as in the definition of the instance. It is easy to verify that ψ is indeed a solution: satisfaction of constraints imposed by the function H follows from vertex-disjointness, while being cohomologous can be easily seen by drawing paths aside one by one, modifying each path by a sequence of operations of jumping over a single face at a time.

Assume then that we are given a solution ψ to the instance I . Take any vertex $v \in V(G)$. Let a_1, a_2, \dots, a_t be arcs adjacent to v that do not lie on C_1 or C_2 , in this order on the plane. As constraints imposed by the function H are satisfied, each of the corresponding arcs in the dual can accommodate only a symbol g_i or 1 . The added arcs ensure that there are no two arcs accommodating two different generators and that every two consecutive arcs accommodating a generator have different directions. Hence, if by G_i we denote the subgraph of G consisting of arcs a such that $\psi(a^*) = g_i$ and vertices adjacent to them, then subgraphs G_i are vertex-disjoint.

Consider the graph G_i . Partition arcs of G_i into cycles and paths as follows: if we have an arc $(u, v) \in E(G_i)$, then we say that the next arc on the path or cycle is the arc $(v, w) \in E(G_i)$ that is the next arc incident to v accommodating a nontrivial symbol in clockwise order. Thus, G_i is partitioned into paths and cycles, where each path begins and ends on C_1 or C_2 . These paths and cycles are edge-disjoint and noncrossing by their construction; moreover, the paths can begin and finish only on C_1 or C_2 . We say that a path P is a *connector* if it connects from C_1 to C_2 ; the symbol associated with a connector from G_i is g_i if it goes from C_1 to C_2 , and g_i^{-1} if it goes from C_2 to C_1 . Let \mathcal{P} be the family of all connectors for all G_i -s. As connectors from \mathcal{P} are non-crossing, they can be naturally ordered along the cycle. Let \mathcal{C} be the sequence of symbols associated with connectors in this order; note that \mathcal{C} is defined uniquely up to a cyclic shift. As the other paths and cycles from partitionings of G_i -s are edge-non-crossing, one can find a cycle C in G^+ winding one time around the ring such that $\psi(C) = \mathcal{C}$.

As ψ is cohomological to ϕ , we have that $\psi(C)$ is conjugate to $\phi(C)$. It follows that \mathcal{C} must admit a cyclic shift of the sequence $(g_1, g_2^{-1}, g_3, \dots, g_{r-1}, g_r^{-1})$ as an embedding. Therefore, we can find r connectors, that is, paths connecting C_1 and C_2 , such that every two clockwise consecutive traverse the ring in different directions. Such a family is an alternating join of size r . Careful inspection of the proof shows that all the steps of the construction of this family can be performed in polynomial time, given a solution ψ . \square

Now, using obstacle characterization of Theorem 4.17 and cohomological formulation of Lemma 4.18, we are able to prove the following min-max theorem.

Lemma 4.19 (Alternating paths/circular cut duality). *Let G be a graph embedded into a ring with C_1 being the boundary of the outer face and C_2 the boundary of the inner face. Assume moreover that C_1 and C_2 are disjoint. Let r be an even integer. Then there exists a polynomial-time algorithm that returns either*

- *an alternating join of C_1 and C_2 of size r , or*
- *a simple noose inside the ring, separating C_1 and C_2 and having at most $r + 4$ alternations.*

Proof. Before we start the proof, without loss of generality we assume that C_1 and C_2 are in fact directed cycles going clockwise. Indeed, we can redirect the arcs of C_1 and C_2 in any manner, as we can safely assume that the alternating join will not use any of these arcs, and also redirecting these arcs do not influence alternation of any noose contained inside the ring.

We construct the instance of the cohomology feasibility problem $I = (G^+, \Lambda, \phi, H)$, as in Lemma 4.18, and apply the algorithm of Theorem 4.16. If the algorithm returns a solution, by Lemma 4.18 we can extract a sufficiently large alternating join and return it. Otherwise, by Theorem 4.17 we are left with two closed undirected walks P, Q in G^+ rooted in some vertex $u \in V(G^+)$, such that for no element $x \in \Lambda$ we have that $x^{-1}\phi(P)x \in H(P)$ and $x^{-1}\phi(Q)x \in H(Q)$. Fix an arbitrary reference curve and let w_P and w_Q be the winding numbers of P and Q in the ring; by reversing each cycle if necessary, we may assume that $w_P, w_Q \geq 0$. Note that then $\phi(R) = (g_1 \cdot g_2^{-1} \cdot g_3 \cdot \dots \cdot g_{r-1} \cdot g_r^{-1})^{w_R}$ for $R \in \{P, Q\}$.

For clarity, in the following, whenever considering alternation, we think of P, Q as non-closed undirected walks that only by coincidence begin and end in the same point. That is, we consider embeddings of alternating sequences into sequences $\vec{\mu}(P), \vec{\mu}(Q)$, where these sequences begin and end in v .

We now claim that if $(w_R \cdot r \leq \mathfrak{a}(R)$ or $w_R = 0)$ for both $R = P, Q$, then for $x = 1$ we have that $x^{-1}\phi(P)x \in H(P)$ and $x^{-1}\phi(Q)x \in H(Q)$, which is a contradiction. Indeed, if $w_P = 0$ then $\phi(P) = 1 \in H(P)$, and if $\mathfrak{a}(P) \geq w_P \cdot r$, then we can embed a sequence $(+1, -1)^{w_P \cdot r/2}$ in $\vec{\mu}(P)$; by taking consecutive $g_i^{\pm 1}$ from the sets $H(a)$ for images of the embedding, and 1 from all sets $H(a)$ for all the other arcs of P , we see that $\phi(P) \in H(P)$. The same argument holds for Q .

Without loss of generality assume then that $w_P \cdot r > \mathfrak{a}(P)$ and $w_P > 0$. By somehow abusing notation, from now on we identify the undirected walk P in G^+ with a naturally corresponding non-degenerated noose in G ; by Observation 4.13 we may assume that P is pretty. We consecutively extract nooses from P keeping the invariant that $w_P \cdot r > \mathfrak{a}(P)$ and that P is pretty. At each step we either find a simple noose with winding number 1 and alternation at most $r + 4$, which can be output by the algorithm, or shorten (with respect to some measure to be defined) noose P keeping the invariant. Moreover, if the extraction step cannot be applied, then P is already simple, has winding number 1 and (by the invariant) at most r alternations when treated as a noose, hence it can be output by the algorithm.

We proceed with a similar cutting scheme as in the proof of Observation 4.14. Assume then that the noose P is not simple, hence some vertex or face is visited more than once. Let us take a shortest interval on P between two consecutive visits of the same face or vertex; by minimality it follows that we may partition P into a simple noose N traversing this interval and the resulting noose P' that is P with the interval cut out (in case of visiting the same face twice, we may need to add small connections within this face). Note that the winding number of N is of absolute value at most 1, as it is simple.

If the cutting was performed due to visiting the same face twice, we have a simple situation: $\vec{\mu}(P')$ is $\vec{\mu}(P)$ with $\vec{\mu}(N)$ carved out, so $\mathfrak{a}(P') \leq \mathfrak{a}(P)$ and $\mathfrak{a}(P') \leq \mathfrak{a}(P) - \mathfrak{a}(N) + 3$ (we may lose at most two alternations on the cut points and potentially one alternation on a cyclic shift of $\vec{\mu}(N)$).

Hence if the winding number of N equals zero or $\mathfrak{a}(N) \geq r + 3$, we are keeping the invariant, and otherwise we may output N (or N reversed if its winding number is -1). Note that thus P' is still pretty.

Assume now that cutting was performed due to visiting the same vertex w twice. Note that $\vec{\mu}(P')$ is basically $\vec{\mu}(P)$ with $\vec{\mu}(N)$ carved out, only with possible manipulations on the term corresponding to passing through the vertex w . Again, it follows that $\mathfrak{a}(P') \leq \mathfrak{a}(P) - \mathfrak{a}(N) + 5$: if we can embed some alternating sequence into a cyclic shift of $\vec{\mu}(P')$ and some other alternating sequence into a cyclic shift of $\vec{\mu}(N)$, then after removing terms corresponding to passing through w from both embeddings, shifting the alternating sequence embedded into a cyclic shift of $\vec{\mu}(N)$ so that it corresponds to the shift rooted at w , and gluing the embedding together, we obtain an embedding into $\vec{\mu}(P)$ of an alternating sequence of length at least $\mathfrak{a}(N) + \mathfrak{a}(P') - 5$ (we may lose additional 2 alternations on gluing and one on shifting $\vec{\mu}(N)$). Hence, if N has nonzero winding number, then either we may reduce P by shortening it to P' (in case $\mathfrak{a}(N) \geq r + 5$) or output N (in case $\mathfrak{a}(N) \leq r + 4$). In order to ensure that P' is still pretty, we may need to apply Observation 4.13 to it. Note that this application can only make a small circumvent of the vertex w , and by the assumptions that C_1 and C_2 are directed cycles, this circumvent will not make P' go outside the ring.

Now consider the case when N has zero winding number; we need to argue that $\mathfrak{a}(P') \leq \mathfrak{a}(P)$, as then we can shortcut P to P' , possibly again applying Observation 4.13 to it. This, however, follows from the same argumentation as in the proof of Observation 4.14: from the fact that P is pretty we can argue that the term corresponding to passing w in $\vec{\mu}(P')$ is contained in the union of terms corresponding to passing w in $\vec{\mu}(P)$, which means that every sequence embeddable into $\vec{\mu}(P')$ is also embeddable into $\vec{\mu}(P)$.

We are left with arguing that the presented procedure will terminate in polynomial number of steps. Note that in every cutting step we either decrease the number of vertices visited by P , or do not increase the number of vertices visited by P but decrease the number of faces and arcs visited by P . Hence, the maximum number of steps performed is at most the number of vertices visited by P times the size of the graph. \square

The following corollary will be used in the algorithm.

Lemma 4.20 (Handling a large ring). *Let G be a directed graph on the plane with some terminals. Let r be an integer and let C_1, \dots, C_{2r+3} be an alternating sequence of concentric cycles in G with no terminal between C_1 and C_{2r+3} . Then there exists a polynomial-time algorithm that outputs either:*

- *a simple noose separating C_1 and C_{2r+3} having at most $2r + 8$ alternations, or*
- *a vertex v surrounded by a sequence of r alternating concentric cycles with no terminals inside them.*

Proof. Let R be a ring with C_1 being the outer cycle and C_{2r+3} being the inner cycle. We use the algorithm of Lemma 4.19 to either find a simple noose having at most $2r + 8$ alternations that separate C_1 from C_{2r+3} , which can be returned by the algorithm, or an alternating join of size $2r + 4$. Assume then that the join $P_1, P_2, \dots, P_{2r+4}$ has been found, where the indices reflect clockwise order of the paths on the ring.

Let v be any vertex that is on the intersection of cycle C_{r+2} and P_{r+2} . We modify the ring R using with the following operations. We delete the part of the ring between P_{2r+3} and P_1 , thus cutting through the ring and creating a graph with outer face boundary consisting of paths P_1, P_{2r+3} and subpaths of cycles C_1 and C_{2r+3} (see Fig. 36). Moreover, we delete the vertex v from




Figure 36: The outer face of the ring R' is marked with dashed gray.

the graph and declare the face in which it was embedded the inner face. Thus, the resulting graph may be viewed as embedded into a ring R' with the aforementioned inner and outer face.

We now apply Lemma 4.15 to the ring R' . We either find r alternating cycles around the inner face, which constitute r alternating cycles around v with no terminals embedded that can be returned by the algorithm, or a simple curve M starting in the inner face and ending in the outer face, having at most r alternations. Assume then that curve M was found.

Assume first that M reaches the part of the boundary of the outer face of R' that is on C_1 . Then M must have passed through cycles C_1, C_2, \dots, C_{r+1} ; these passages define an embedding of an alternating sequence of size $r + 1$ into $\bar{\mu}(M)$, which contradicts the fact that $\alpha(M) \leq r$. Similarly, if M reaches C_{2r+3} , P_1 , or P_{2r+3} , then it must have passed through sequences of paths $(C_{r+3}, C_{r+4}, \dots, C_{2r+3})$, $(P_1, P_2, \dots, P_{r+1})$, or $(P_{r+3}, P_{r+4}, \dots, P_{2r+3})$, respectively, and in each case we get a contradiction. Hence, the curve M could not be found and we are done. \square

5 Decomposition

In this section we show how to decompose the graph G into a bounded number of weakly connected subgraphs (called henceforth *components*), such that the interaction between the components is somehow limited, and the terminals are kept outside of the components. We start with defining a notion of an alternation suitable for weakly connected subgraphs of G ; the main property of our components is that we control their alternation.

Definition 5.1 (incident arcs). Let G be a graph and let H be its subgraph. By $\hat{E}(H)$ we denote the set of arcs of G incident to at least one vertex of H , but not belonging to H ; that is, $\hat{E}(H) = E(G) \setminus (E(H) \cup E(G[V(G) \setminus V(H)]))$.

Definition 5.2 (alternation of a face of a subgraph). Let G be a plane graph, let H be a weakly connected subgraph of G and let f be a face of H . Consider an undirected walk P in H that goes around the face f in such a direction that it leaves the face f to the right, and the subgraph H

to the left; that is, P goes counter-clockwise if f is the outer face of H , and clockwise otherwise. Consider a cyclic sequence $P(\hat{E}(H))$ of elements of $\{+1, -1\}$ constructed as follows: we go along $\hat{E}(H)$ and insert into $P(\hat{E}(H))$ an element $+1$ or -1 for each endpoint of an arc of $\hat{E}(H)$ we encounter along the walk P , depending on whether this is a starting or ending point of the arc. If multiple arcs are encountered at one vertex $v \in V(H)$, consider them in the counter-clockwise order as they appear on the face f . The *alternation* of the face f in H is the alternation of the sequence $P(\hat{E}(H))$.

Note that each arc $e \in \hat{E}(H)$ corresponds to exactly one entry in $P(\hat{E}(H))$ if it has exactly one endpoint in $V(H)$, and to two entries if it has both endpoints in $P(\hat{E}(H))$.

We also note that for any weakly connected subgraph H and its face f , there exists a noose $N(f, H)$ that goes parallelly and very closely to the walk P , is contained in the face f , does not visit any vertex of G and $\bar{\mu}(N(f, H)) = P(\hat{E}(H))$ (up to a cyclic shift). We call such a noose a *border noose* of H . Moreover, we may assume that a border noose of a H is sufficiently close to H , so that for any two disjoint weakly connected subgraphs of G , their border nooses and the areas enclosed by them are disjoint.

We are now ready to define components in our decomposition.

Definition 5.3 (disc component). Let G be a plane graph and let H be a weakly connected subgraph of G . We call H a *disc component* of G if every arc in $\hat{E}(H)$ is contained in the outer face of H . The *alternation* of a disc component H is the alternation of its outer face.

Definition 5.4 (ring component). Let G be a plane graph and let H be a weakly connected subgraph of G . We call H a *ring component* of G if there exists a face f_{IN} of H different than its outer face (called the *inner face* of the component) such that every arc in $\hat{E}(H)$ is contained either in the outer face of H or in f_{IN} . The *alternation* of a ring component H is the maximum of alternations of its outer face and the inner face f_{IN} .

By $\mathfrak{a}(H)$ we denote an alternation of a component H .

Note in both component definitions, we do not require that the graph H is an *induced* subgraph of G . In other words, we allow arcs in $\hat{E}(H)$ that have both endpoints in H .

Definition 5.5 (decomposition). Let G be a plane graph having a set $T \subseteq V(G)$ of terminals. Then a set \mathcal{D} of (disc and ring) components of G is a *decomposition* of G iff

- every vertex of G is in exactly one component of \mathcal{D} ;
- for each terminal $t \in T$ there exists a disc component $H_t^{\text{disc}} \in \mathcal{D}$ that consists of the vertex t only.

The *disc (ring) alternation* of the decomposition \mathcal{D} is the maximum alternation of a disc (ring) component in \mathcal{D} .

We will control two natural measures of a quality of a decomposition: its alternation and the number of its components. Moreover, we will require that ring components are embedded into a decomposition in a special way.

Definition 5.6 (isolating component). Consider a decomposition \mathcal{D} of a plane graph G with terminals T . We say that a disc component $H^{\text{disc}} \in \mathcal{D}$ is a *d-isolating* component if the subgraph of G induced by the vertices of H^{disc} (i.e., the subgraph H^{disc} together with all the arcs of $\hat{E}(H^{\text{disc}})$ that have both endpoints in H^{disc}) contains a sequence of d alternating concentric cycles and each edge of $\hat{E}(H^{\text{disc}})$ that has exactly one endpoint in H^{disc} either lies inside the innermost of these cycles (an *inner edge*) or lies outside of the outermost of these cycles (an *outer edge*).




Figure 37: A ring component with two levels of isolation.

We note that, if the graph G is weakly connected, in the space enclosed between the innermost and the outermost cycle from the definition of the isolating component H^{disc} there are only arcs and vertices of H^{disc} and arcs of $\hat{E}(H^{\text{disc}})$ that have both endpoints in H^{disc} .

Definition 5.7 (ring isolation). Consider a decomposition \mathcal{D} of a plane graph G with terminals T . We say that the decomposition has *ring isolation* (Λ, d) if for every ring component $H^{\text{ring}} \in \mathcal{D}$ with outer face f_{OUT} inner face f_{IN} there exist 2Λ disc components $H_{IN,\lambda}^{\text{disc}}, H_{OUT,\lambda}^{\text{disc}} \in \mathcal{D}$, $1 \leq \lambda \leq \Lambda$, such that each of these components is d -isolating and:

- the components $H_{IN,\lambda}^{\text{disc}}$, $1 \leq \lambda \leq \Lambda$, are contained inside f_{IN} and the components $H_{OUT,\lambda}^{\text{disc}}$, $1 \leq \lambda \leq \Lambda$ are contained in f_{OUT} ;
- for each $1 \leq \lambda < \Lambda$, each inner arc of $H_{IN,\lambda}^{\text{disc}}$ has the second endpoint in $H_{IN,\lambda+1}^{\text{disc}}$, and each outer arc of $H_{OUT,\lambda}^{\text{disc}}$ has the second endpoint in $H_{OUT,\lambda+1}^{\text{disc}}$;
- for each $1 < \lambda \leq \Lambda$, each outer arc of $H_{IN,\lambda}^{\text{disc}}$ has the second endpoint in $H_{IN,\lambda-1}^{\text{disc}}$, and each inner arc of $H_{OUT,\lambda}^{\text{disc}}$ has the second endpoint in $H_{OUT,\lambda-1}^{\text{disc}}$;
- each outer arc of $H_{IN,1}^{\text{disc}}$ and each inner arc of $H_{OUT,1}^{\text{disc}}$ has the second endpoint in H^{ring} .

Moreover, we require that no disc component serves as an isolation to two ring components, that is, the components $H_{IN,\lambda}^{\text{disc}}$ and $H_{OUT,\lambda}^{\text{disc}}$, $1 \leq \lambda \leq \Lambda$ are pairwise distinct for different ring components of \mathcal{D} .

Note that each of the d alternating concentric cycles of $H_{IN,\lambda}^{\text{disc}}$ (forming a d -isolation) is contained in the inner face f_{IN} , while each of the d alternating concentric cycles of $H_{OUT,\lambda}^{\text{disc}}$ encloses H^{ring} .

We say that a decomposition has *positive isolation* if it has isolation (Λ, d) for some $\Lambda, d > 0$.

We are now ready to state our decomposition theorem.

Theorem 5.8 (Main Decomposition Theorem). *Let G be a plane graph with k terminal pairs, each of them having degree 1. Let Λ , d and r be positive integers. Then in $O^*(2^{O(\Lambda(d+r)k^2)})$ time we can either find a set of r alternating concentric cycles with no terminal inside the outermost cycle, or compute a set of at most $2^{O(\Lambda(d+r)k^2)}$ pairs (G_i, \mathcal{D}_i) where:*

1. *each G_i is a plane graph of size polynomially bounded in the size of G , with k terminal pairs, where each terminal is of degree one in G_i ;*
2. *\mathcal{D}_i is a decomposition of the graph G_i with $O(\Lambda k^2)$ components, ring isolation d and alternation $O(\Lambda(d+r)k^2)$;*
3. *G is a YES-instance to k -DPP, if and only if there exists i such that G_i is a YES-instance to k -DPP.*

Proof. We decompose the graph in an iterative manner. At i -th iteration, we are given a graph G_0 with k terminal pairs and family \mathcal{C} of pairwise disjoint disc components of G_0 that are to be partitioned further. Moreover, we assume that each terminal in G_0 has degree one. Each iteration decreases the total number of terminals in the components of \mathcal{C} , thus the iteration ends after at most $2k$ steps.

In each iteration, we first filter out any component $H \in \mathcal{C}$ that does not contain any terminals. Such a component may be output as a disc component in the final decomposition. If \mathcal{C} becomes empty, we finish the algorithm.

Otherwise, we pick any component $H_0 \in \mathcal{C}$ to decompose it further (and remove it from \mathcal{C}). By κ denote the number of terminals contained in H_0 . If there is a terminal on the outer face of H , we create a new disc component that contains it, delete the terminal and its incident arc from H_0 and add back H_0 to \mathcal{C} . Each such step produces a new disc component, decreases κ and increases $\mathfrak{a}(H_0)$ by a constant. Moreover, it maintains the connectivity of H_0 .

Thus, from this point we may assume that H_0 contains $\kappa > 0$ terminals, but none of these terminals lies in the outer face of H_0 . Let us pick one terminal $t \in T \cap V(H_0)$ and let F_t be the face that contains t . Let γ_t be an arbitrary circle with centre t and with sufficiently small radius such that $\gamma_t \cap G_0$ consists of a single point which is an intersection of γ_t with the arc incident to t . Invoke Lemma 4.15 on the graph $G_0 \setminus \{t\}$, curves γ_t and a border noose $N = N(f_0, H_0)$, where f_0 is the outer face of H_0 . Find the largest even a for which we obtain a sequence C_1, C_2, \dots, C_a of alternating concentric cycles in H_0 . Moreover, invoke Lemma 4.15 for $a + 2$ and obtain a simple curve M connecting γ_t and N with at most $a + 2$ alternations.

Assume C_1 is the innermost and C_a is the outermost of the constructed cycles. For $s = 1, 2, \dots, \kappa$ let $i(s)$ be the largest integer $0 \leq i(s) \leq a$ such that there are at most s terminals contained in the area enclosed by the cycle $C_{i(s)}$ (note that no terminal lies on any of the cycles C_1, C_2, \dots, C_a , as terminals are of degree one in G_0). Note that $i(\kappa) = a$; we denote $i(0) = 0$.

Let $q = 2r + 1$ and let $S^{\text{big}} \subseteq \{1, 2, \dots, \kappa\}$ be the set of such integers that $s \in S^{\text{big}}$ whenever $i(s) - i(s-1) \geq 2\Lambda d + 2(\Lambda + 1)q + 2$. For each $s \in S^{\text{big}}$, apply Lemma 4.20 to the following $2(\Lambda + 1)$ sets of $q + 2 = 2r + 3$ concentric cycles:

- $C_{i(s-1)+(\Lambda-\lambda)(d+q)+1}, C_{i(s-1)+(\Lambda-\lambda)(d+q)+2}, \dots, C_{i(s-1)+(\Lambda-\lambda)(d+q)+q+2}$ for $0 \leq \lambda \leq \Lambda$;
- $C_{i(s)-(\Lambda-\lambda)(d+q)-q-1}, C_{i(s)-(\Lambda-\lambda)(d+q)-q}, \dots, C_{i(s)-(\Lambda-\lambda)(d+q)}$ for $0 \leq \lambda \leq \Lambda$.

If any of the applications of Lemma 4.20 returns a sequence of r concentric cycles, we return it as an outcome of the algorithm. Otherwise, we obtain $2\Lambda + 2$ circular cuts $N_s^{IN, \lambda}$ and $N_s^{OUT, \lambda}$, $0 \leq \lambda \leq \Lambda$: the cut $N_s^{IN, \lambda}$ separates $C_{i(s-1)+(\Lambda-\lambda)(d+q)+1}$ from $C_{i(s-1)+(\Lambda-\lambda)(d+q)+q+2}$ and the cut

$N_s^{OUT,\lambda}$ separates $C_{i(s)-(\Lambda-\lambda)(d+q)-q-1}$ from $C_{i(s)-(\Lambda-\lambda)(d+q)}$. As $i(s)-i(s-1) \geq 2\Lambda d+2(\Lambda+1)q+2$, we have that

$$i(s-1) + \Lambda(d+q) + q + 1 \leq i(s) - \Lambda(d+q) - q - 1$$

and the cycle $C_{i(s-1)+\Lambda(d+q)+q+2}$ lies between $N_s^{IN,0}$ and $N_s^{OUT,0}$.

Note that, for any $1 \leq \lambda \leq \Lambda$:

- the d cycles $C_{i(s-1)+(\Lambda-\lambda)(d+q)+q+2}, C_{i(s-1)+(\Lambda-\lambda)(d+q)+q+3}, \dots, C_{i(s-1)+(\Lambda-\lambda+1)(d+q)+1}$ are contained between $N_s^{IN,\lambda}$ and $N_s^{IN,\lambda-1}$,
- the d cycles $C_{i(s)-(\Lambda-\lambda-1)(d+q)}, C_{i(s)-(\Lambda-\lambda-1)(d+q)+1}, \dots, C_{i(s)-(\Lambda-\lambda)(d+q)-q-1}$ are contained between $N_s^{OUT,\lambda-1}$ and $N_s^{OUT,\lambda}$.

As the simple curve M connects γ_t with N , it crosses all cycles C_i as well as all circular cuts $N_s^{IN,\lambda}, N_s^{OUT,\lambda}$ for $s \in S^{\text{big}}, 1 \leq \lambda \leq \Lambda$. By slightly perturbing M , we may assume that M crosses these circular cuts either in a vertex of G or inside a face of G , never on an edge of G . For each $s \in S^{\text{big}}$, remove from M the subcurve between the intersection with $N_s^{IN,0}$ that is closest to γ_t on M and the intersection with $N_s^{OUT,0}$ that is closest to N on M ; if the intersection happens in the vertex of G , we remove a slightly shorter part of M , so that the remaining curve still traverses the intersection vertex and ends in a face of G . Denote by \mathcal{M} the set of remaining subcurves of M (note that there are at most $|S^{\text{big}}| + 1 \leq \kappa + 1$ of them).

We now note that all curves in \mathcal{M} intersect $O(\Lambda(d+r)\kappa)$ alternating cycles from the sequence C_1, C_2, \dots, C_a , as they intersect $O(\Lambda(d+r))$ cycles between $C_{i(s-1)+1}$ and $C_{i(s)}$ for each $1 \leq s \leq \kappa$. Recall that the alternation of M is at most $a + 2$, while M intersects all cycles C_1, C_2, \dots, C_a . Therefore the sum of alternations of all curves in \mathcal{M} is $O(\Lambda(d+r)\kappa)$.

We are now going to cut the graph along the circular cuts $N_s^{IN,\lambda}$ and $N_s^{OUT,\lambda}$ for $s \in S^{\text{big}}, 1 \leq \lambda \leq \Lambda$, as well as along the curves of \mathcal{M} . However, both circular cuts and the curves in \mathcal{M} may traverse a vertex, whereas we want to cut arcs only. To cope with this, we introduce the following bounded search tree strategy.

Recall that any curve of \mathcal{M} as well as any cut $N_s^{IN,\lambda}, N_s^{OUT,\lambda}$ is a simple curve. Therefore we can apply Observation 4.13 to each of those curves, making them pretty. Let N_0 be any of these curves, and let v be a vertex on N_0 . By the definition of a pretty curve (Definition 4.12), the set $S_v \in \bar{\mu}(N_0)$ that corresponds to v equals $\{+1, -1\}$ or \emptyset .

For a given curve N_0 , we first resolve all vertices with $S_v = \{+1, -1\}$. For each such vertex, we branch into four cases. We choose one side of the curve N_0 (left or right) around the vertex v and one type of arcs incident to v (incoming or outgoing arcs). In each branch, we delete from the graph G_0 the arcs of the chosen type from the chosen side of the curve and perturb N_0 slightly to omit the vertex v from the chosen side. In this way, we do not increase $\mathfrak{a}(N_0)$, as in $\bar{\mu}(N_0)$ we exchange the $S_v = \{+1, -1\}$ for a sequence (of arbitrary length) of equal one-element sets, corresponding to the remaining arcs from the chosen side of N_0 . Moreover, for any solution to the k -DPP problem, there exists a choice where we do not delete any arcs of the solution, as the vertex v may lie only on one path of the solution, and this path cannot enter and leave the vertex v from both sides of N_0 .

Note that the number of these vertices is at most $\mathfrak{a}(N_0)$, thus we create at most $4^{\mathfrak{a}(N_0)}$ branches.

Once we have resolved all vertices with $S_v = \{+1, -1\}$, we move to the second case. For a vertex v we may have $S_v = \emptyset$ in two situations. First, N_0 may visit the vertex v , but leaves all incident edges on one of its sides. However this would contradict the assumption that N_0 is non-degenerate.

Second, v may be a source or a sink. We consider here two options: either we modify the curve N_0 to omit the vertex v from the left or from the right. Let us investigate how such a change

will influence $\vec{\mu}(N_0)$ and $\mathfrak{a}(N_0)$. We replace $S_v = \emptyset$ with a sequence (of arbitrary length) of equal one-element sets corresponding to the traversed arcs of G ; such a situation may increase $\mathfrak{a}(N_0)$ by one. However, we have a choice of whether we insert a sequence of sets $\{+1\}$ or $\{-1\}$: if we omit the vertex v from the left, we cross the arcs incident to it in a different orientation than if we omit the vertex v from the right.

If in $\vec{\mu}(N_0)$ there exists a vertex v with $S_v = \emptyset$, but one neighbouring set S with $|S| = 1$, we may modify N_0 around v so that we replace S_v with a sequence of sets equal to S . It may be easily seen that if we apply this operation to all the sets S_v equal to \emptyset , then the alternation of N_0 does not increase in case N_0 is a non-closed curve (from \mathcal{M}) or increases by at most 1 in case N_0 is a noose (of form $N_s^{IN,\lambda}$ or $N_s^{OUT,\lambda}$). As we have removed all two-element sets from $\vec{\mu}(N_0)$, we may not perform the above operation only if $\vec{\mu}(N_0)$ consists only of empty sets. In this case, we modify N_0 around each vertex, so that $\vec{\mu}(N_0)$ is a sequence of sets $\{+1\}$; note that the alternation of N_0 changes from 0 to 1 in this case.

Let us now summarize. Recall that:

- $|\mathcal{M}| \leq \kappa + 1$;
- there are at most $2(\Lambda + 1)\kappa$ curves $N_s^{IN,\lambda}$, $N_s^{OUT,\lambda}$, $s \in S^{\text{big}}$, $0 \leq \lambda \leq \Lambda$;
- the sum of alternations of all curves of \mathcal{M} is $O(\Lambda(d+r)\kappa)$;
- each curve $N_s^{IN,\lambda}$ and $N_s^{OUT,\lambda}$ has alternation at most $2r + 8$ (by Lemma 4.20).

Therefore the aforementioned procedure generates $4^{O(\Lambda(d+r)\kappa)}$ subcases, and increases alternation of each curve by at most one.

Before we describe the outcome of the partitioning of H_0 , let us define the notion of *connectifying* a component. Assume we are given an open connected subset A of the plane, isomorphic to a disc or to a ring, such that no vertex of G_0 lies on the border of A (i.e., not in A , but in the closure of A). For each of the (one or two) borders of A that are homeomorphic to a circle, travel along the border and, for each arc that it intersects, subdivide it, inserting a new vertex inside A . For each two consecutive newly added vertices, connect them with a length-two path inside A , where the middle vertex of the path is a sink (i.e., both arcs from the path point from the subdivided arcs of G_0 towards the vertex in the middle). As G_0 is weakly connected, after this operation, the subgraph of G_0 consisting of all arcs and vertices completely contained in A is weakly connected, whereas the answer to k -DPP on G_0 does not change, as the added arcs are useless from the point of view of constructing directed paths.

We are now ready to partition the graph into the following components.

- We create a disc component containing the terminal t only.
- For each $s \in S^{\text{big}}$, we create the following $2\Lambda + 1$ components.
 - A ring component H_s^{ring} that is a connectification of a subgraph consisting all vertices and edges of the graph G_0 contained between the noose $N_s^{IN,0}$ and the noose $N_s^{OUT,0}$. The face of H_s^{ring} that contains $N_s^{IN,0}$ is the inner face of H_s^{ring} , and the face that contains $N_s^{OUT,0}$ is the outer face. Note that, as H_s^{ring} contains the cycle $C_{i(s-1)+\Lambda(d+q)+q+2}$, these faces are distinct.
 - For any $1 \leq \lambda \leq \Lambda$, a disc component $H_{s,IN,\lambda}^{\text{disc}}$ that is a connectification of a subgraph of G_0 enclosed in the area with its border being a concatenation of the following four curves: a minimal segment of a curve of \mathcal{M} connecting $N_s^{IN,\lambda-1}$ with $N_s^{IN,\lambda}$, the curve

$N_s^{IN,\lambda}$, the same minimal segment of a curve of \mathcal{M} , but now traversed backwards, and the curve $N_s^{IN,\lambda-1}$. Note that, as the chosen subset of the plane is homeomorphic to a disc, the $H_{s,IN,\lambda}^{\text{disc}}$ is in fact a disc component.

- For any $1 \leq \lambda \leq \Lambda$, a disc component $H_{s,OUT,\lambda}^{\text{disc}}$ created in the same manner as $H_{s,IN,\lambda}^{\text{disc}}$, but between curves $N_s^{OUT,\lambda-1}$ and $N_s^{OUT,\lambda}$.

We note that the d alternating cycles $C_{i(s-1)+(\Lambda-\lambda)(d+q)+q+2}, \dots, C_{i(s-1)+(\Lambda-\lambda+1)(d+q)+1}$ are contained in the subgraph of G induced by the vertices of $H_{s,IN,\lambda}^{\text{disc}}$, fulfilling all the requirements to make $H_{s,IN,\lambda}^{\text{disc}}$ a d -isolating component. Similarly, the components $H_{s,OUT,\lambda}^{\text{disc}}$ are d -isolating as well. Therefore H_s^{ring} has (Λ, d) -isolation, as required.

- For each $s \in S^{\text{big}} \cup \{\kappa+1\}$, we insert into \mathcal{C} the connectification of a subgraph enclosed by the following noose. We denote $N_0^{OUT,\Lambda} = \gamma_t$ and $N_{\kappa+1}^{IN,\Lambda} = N$. We concatenate a minimal segment of the curve of \mathcal{M} connecting $N_{\text{pred}(s)}^{OUT,\Lambda}$ and $N_s^{IN,\Lambda}$, the curve $N_s^{IN,\Lambda}$, again the same minimal segment but traversed backwards, and the curve $N_{\text{pred}(s)}^{OUT,\Lambda}$, where $\text{pred}(s)$ is the maximum element of S^{big} smaller than s , or 0 if it does not exist. Note that, as in the case of previous two components, the aforementioned area in the plane is homeomorphic to a disc, thus we insert into \mathcal{C} disc components only.

We note that at each step, we invoke the connectification algorithm a few times, but in pairwise disjoint subsets of the plane. Therefore, the enhanced graph G_0 remains plane.

Note that, after this step, the total number of terminals in \mathcal{C} decreased by one, as the terminal t is put in its own disc component and each other terminal is put into exactly one new recursive call. Therefore, the number of iterations is at most $2k$. Moreover, as the step creates at most κ ring components and $2\Lambda\kappa$ disc components that are not components with a terminal, we obtain at most $O(\Lambda k^2)$ components in total and at most $2^{O(\Lambda(d+r)k^2)}$ subcases.

Let us now bound the alternations of the constructed components. Each constructed ring component H_s^{ring} has two borders $N_s^{IN,0}$ and $N_s^{OUT,0}$ of alternation at most $2r+8$, with possible slight modifications due to the branching procedure that can add 1 to the alternation, thus its alternation is at most $2r+9$. Similarly, each isolating component $H_{s,IN,\lambda}^{\text{disc}}$ and $H_{s,OUT,\lambda}^{\text{disc}}$ is surrounded by two circular cuts of alternation $2r+9$ each and two subcurves of a curve of \mathcal{M} , again with possible slight modifications due to the branching procedure that can add 1 to the alternation. Recall that the sum of alternations of all curves in \mathcal{M} is $O(\Lambda(d+r)\kappa)$, thus each isolating disc component has alternation $O(\Lambda(d+r)k)$. Moreover, γ_t has alternation 1.

At each iteration, we put into \mathcal{C} components surrounded by two copies of a subcurve of a curve from \mathcal{M} and two nooses being either circular cuts $N_s^{IN,\Lambda}$, $N_s^{OUT,\Lambda}$, the curve γ_t or N . An operation of cutting away a disc component with a terminal may increase the alternation of a component by at most one per terminal, thus $2k$ in total. We infer that at iteration i , any component in \mathcal{C} has alternation bounded by $O(i(\Lambda(d+r)k))$. Therefore any computed disc component has alternation $O(\Lambda(d+r)k^2)$.

This concludes the proof of the decomposition theorem. □

6 Bundles and bundle words

Our decomposition theorem, Theorem 5.8, provides us either with a situation where an irrelevant vertex can be found, or with a bounded number of subcases, each of the subcase being a modified graph G with a decomposition \mathcal{D} of bounded number of components and with bounded alternation.

In this section we focus on solving one fixed subcase, that is, we focus on a single graph G with a decomposition \mathcal{D} .

6.1 Bundle arcs, bundles and component multigraph

Consider a graph G with its decomposition \mathcal{D} . For any $v \in V(G)$, let $H(v)$ be the component of \mathcal{D} that contains v . Let $\hat{E}(G, \mathcal{D})$ be the set of arcs of G that are not contained in any component of \mathcal{D} , that is, $\hat{E}(G, \mathcal{D}) = E(G) \setminus \bigcup_{H \in \mathcal{D}} E(H)$. Any element of $\hat{E}(G, \mathcal{D})$ is called a *bundle arc*.

The set of bundle arcs form a structure of a multigraph on the set of components \mathcal{D} ; we call this multigraph a *component multigraph*. If the number of components and the alternation of \mathcal{D} is bounded, the set of arcs of the component multigraph can be decomposed into a bounded number of *bundles* of arcs that go parallelly and in the same direction. We now formalize these notions.

Definition 6.1 (component multigraph). Let G be a plane graph and let \mathcal{D} be its decomposition. The *component multigraph* of G and \mathcal{D} , denoted $G_{\text{comp}}(G, \mathcal{D})$, is a multigraph with vertex set \mathcal{D} and arcs set $\{(H(u), H(v)) : (u, v) \in \hat{E}(G, \mathcal{D})\}$.

Note that the component multigraph is planar, and the embedding of G naturally imposes a non-standard embedding of $G_{\text{comp}}(G, \mathcal{D})$, where each disc component is contracted into a single point and each ring component is contracted into a closed curve, separating the part of the graph inside the ring component from the part outside it.

We sometimes abuse the notation and identify the bundle arc being an arc in the multigraph with the vertex set \mathcal{D} with the corresponding arc in G .

We now define a notion of a *bundle*, that gathers together bundle arcs that “serve the same role”.

Definition 6.2 (bundle). Let G be a graph and let \mathcal{D} be its decomposition. A sequence $B = (b_1, b_2, \dots, b_s)$ of bundle arcs is called a *bundle* if

- there exist two components H_1, H_2 , such that each bundle arc b_i leads from H_1 to H_2 (possibly $H_1 = H_2$);
- for any $1 < i \leq s$, in the graph $H_1 \cup H_2 \cup \{b_{i-1}, b_i\}$ the unique face to the left of b_{i-1} and to the right of b_i is empty, that is, does not contain any point of the embedding of the graph G .

Note that there may exist other arcs from H_1 to H_2 in G , that do not belong to B .

Consider faces of G . We say that a face is a *component face* if it belongs to some component H ; otherwise it is called a *mortar face*. A mortar face F can be either a *bundle face* if all the bundle arcs on the boundary of F belong to the same bundle (in which case there is at most 2 of them), or an *end-face* otherwise.

Lemma 6.3 (bundle in the dual). *Let G be a graph, \mathcal{D} be its decomposition and $B = (b_1, b_2, \dots, b_s)$ be a bundle. Then B is a directed path or a directed cycle in the dual of G .*

Proof. Let B connect H_1 with H_2 in \mathcal{D} . By the definition of a bundle, in the graph $H_1 \cup H_2 \cup B$ for any $1 < i \leq s$ there exists a face f_i whose border consists of b_{i-1} , b_i and parts of H_1 and H_2 ; moreover, f_i lies to the left of b_{i-1} and to the right of b_i . Therefore, in the dual of G , for $1 < i \leq s$, b_i is an arc between f_{i-1} and f_i . Moreover, no face f_i is an endpoint of b_s nor a starting point of b_1 . Therefore B is indeed a directed path or a cycle in the dual of G . \square

We now show that, given a graph G with a decomposition \mathcal{D} with bounded alternation, we can efficiently partition the bundle arcs into a bounded number of bundles.

Lemma 6.4 (bundle recognition). *Given a graph G embedded in a plane, together with its decomposition \mathcal{D} of alternation $\mathfrak{a}(\mathcal{D})$, one can in polynomial time partition the bundle arcs into a set \mathcal{B} of bundles, such that each component $H \in \mathcal{D}$ with is incident to $O(\mathfrak{a}(\mathcal{D})|\mathcal{D}| + |\mathcal{D}|^2)$ bundles.*

Proof. Consider first any two components H_1 and H_2 , $H_1 \neq H_2$. We are to partition the bundle arcs with one endpoint in H_1 and second endpoint in H_2 into $O(\mathfrak{a}(\mathcal{D}) + |\mathcal{D}|)$ bundles. If there is a constant number of them, the task is trivial, so let us assume there are at least three bundle arcs between these two components.

Note that, if any of the components H_1, H_2 is a ring component, the bundle arcs between H_1 and H_2 lie either in the inner face of the ring component, or in the outer face. Moreover, due the assumption of isolation it may not happen that both H_1 and H_2 are ring components. Let b_1, b_2, \dots, b_m be the bundle arcs with one endpoint in H_1 and the second endpoint in H_2 , in the order of their appearance on a walk around the face of H_1 that contains H_2 (i.e., the outer face if H_1 is a disc component, and the outer or the inner face if H_1 is a ring component). By planarity, this is also a reversed order of they appearance in a walk around the face of H_2 that contains H_1 .

Denote $b_0 = b_m$ and let A_i , $1 \leq i \leq m$, be the unique face of the graph $H_1 \cup H_2 \cup \{b_{i-1}, b_i\}$ that does not contain any arc b_j for $j \notin \{i-1, i\}$. Note that b_{i-1} and b_i may not be grouped together in the same bundle for the following reasons.

1. b_{i-1} and b_i go in different directions, that is, one of this bundle arcs leads from H_1 to H_2 and the second one from H_2 to H_1 . However, this may happen at most $O(\mathfrak{a}(\mathcal{D}))$ times.
2. There is a component inside A_i ; as the areas of the components are disjoint and disjoint with bundle arcs, this may happen at most $|\mathcal{D}|$ times.
3. There is a bundle arc that is a self-loop of H_1 or H_2 in the component multigraph contained in A_i . However, as both its endpoints are contained in A_i , this happens at most $O(\mathfrak{a}(\mathcal{D}))$ times.

If none of the aforementioned events happens between the arcs $b_i, b_{i+1}, b_{i+2}, \dots, b_j$, then all these arcs form a bundle (possibly in the reversed order). As there are $O(\mathfrak{a}(\mathcal{D}) + |\mathcal{D}|)$ aforementioned events, the claim is proven.

We are left with the task of partitioning self-loops (in the component multigraph) of an arbitrarily chosen component H into $O(\mathfrak{a}(\mathcal{D}) + |\mathcal{D}|)$ bundles. The arguments are similar to the previous case, but the topology of the situation is a bit different.

We focus on bundle arcs that lie in a single face f of H ; there is one such face in the case of a disc component and two such faces in the case of a ring component. Denote the set of bundle arcs in question by \mathcal{B} . The bundle arcs of \mathcal{B} partition f into a number of areas; let us denote the set of this areas as \mathcal{A} . Let us define a (natural) structure of a tree on the set \mathcal{A} : two areas are adjacent in the tree $T_{\mathcal{A}}$ if they share a bundle arc of \mathcal{B} as a part of their borders. Note that the edges of $T_{\mathcal{A}}$ are in a bijection with the set \mathcal{B} .

First, note that the edge of $T_{\mathcal{A}}$ incident to a leaf in the tree $T_{\mathcal{A}}$ corresponds to a bundle arc $b \in \mathcal{B}$ whose both endpoints are subsequent intersections of \mathcal{B} with a border noose of H . Therefore, the number of leaves of $T_{\mathcal{A}}$ is bounded by $\mathfrak{a}(H) \leq \mathfrak{a}(\mathcal{D})$. Consequently, the same bound holds for the number of vertices of $T_{\mathcal{A}}$ of degree at least three.

Second, as in the case of bundle arcs between two different components, there are:

1. at most $|\mathcal{D}|$ areas $A \in \mathcal{A}$ that contain a component inside;
2. $O(\mathfrak{a}(\mathcal{D}))$ areas $A \in \mathcal{A}$ that are of degree two in $T_{\mathcal{A}}$, but the two bundle arcs on the border of A do not form a bundle, as they are both oriented clockwise or both oriented counter-clockwise around the area A ;

These $O(\mathfrak{a}(\mathcal{D}) + |\mathcal{D}|)$ vertices of $T_{\mathcal{A}}$, together with at most $\mathfrak{a}(\mathcal{D})$ vertices of degree at least 3, partition the set edges of $T_{\mathcal{A}}$ into a set of $O(\mathfrak{a}(\mathcal{D}) + |\mathcal{D}|)$ paths, and the edges on each path form a bundle. This concludes the proof of the lemma. \square

The ability to partition the bundle arcs of a given decomposition into bundles motivates the following definition.

Definition 6.5 (bundle multigraph). For a plane graph G , its decomposition \mathcal{D} and partition \mathbb{B} of bundle arcs into bundles, we define a *bundle multigraph* $G_{\text{bundle}}(G, \mathcal{D}, \mathbb{B})$ as a multigraph obtained from $G_{\text{comp}}(G, \mathcal{D})$ by identifying bundle arcs of one bundle in \mathbb{B} . In other words, $G_{\text{bundle}}(G, \mathcal{D}, \mathbb{B})$ is a multigraph with vertex set \mathcal{D} and edge set \mathbb{B} .

Note that $G_{\text{bundle}}(G, \mathcal{D}, \mathbb{B})$ is planar and the non-standard embedding of $G_{\text{comp}}(G, \mathcal{D})$ naturally imposes a non-standard embedding of $G_{\text{bundle}}(G, \mathcal{D}, \mathbb{B})$ (i.e., where each disc component is represented by a point, and each ring component by a closed curve).

Definition 6.6 (bundled instance). A *bundled instance* is a tuple consisting of a k -DPP input G , its decomposition \mathcal{D} of positive isolation and partition into bundles \mathbb{B} , equipped with an embedding of G into a plane together with corresponding embeddings of $G_{\text{comp}}(G, \mathcal{D})$ and $G_{\text{bundle}}(G, \mathcal{D}, \mathbb{B})$.

We somewhat abuse the notation and denote a bundled instance as $(G, \mathcal{D}, \mathbb{B})$, making the embeddings implicit.

6.2 Bundle words and spirals

The following notion is crucial for the rest of this section.

Definition 6.7 (bundle word). Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance. Then, for a directed walk P in G we define a word $\text{bw}(P)$ over the alphabet \mathbb{B} , called the *bundle word of P* , as follows: we follow P and for each bundle arc e visited by P we append the bundle $B \in \mathbb{B}$ that contains e to $\text{bw}(P)$.

Intuitively, $\text{bw}(P)$ describes precisely how does the path P circulate around the component multigraph $G_{\text{comp}}(G, \mathcal{D})$. Note that it does not distinguish different bundle arcs inside one bundle.

Assume for a moment that \mathcal{D} does not contain any ring components. Consider a solution $(P_i)_{i=1}^k$ to k -DPP in G . The core observation is that if we know $\text{bw}(P_i)$ for each $1 \leq i \leq k$, then the k -DPP problem can be solved using the cohomology algorithm of Schrijver [28]; a generalization of this statement, including some technical difficulties around ring components, is proven in Section 7. The goal of Section 8 is to show that there is only a bounded number of *reasonable* choices for $\text{bw}(P_i)$; in other words (but still ignoring the ring components) we show how to construct an $f(k)$ -sized family of sequences $(p_i)_{i=1}^k$ such that if G is a yes instance to k -DPP, then there exists a solution $(P_i)_{i=1}^k$ and a sequence $(p_i)_{i=1}^k$ such that $\text{bw}(P_i) = p_i$ for each $1 \leq i \leq k$. In this section we prepare some preliminary results for both claims.

We start with an insight into properties of bundle words $(\text{bw}(P_i))_{i=1}^k$ for a solution to k -DPP.

Definition 6.8 (spiral in a bundle word). Let P be a directed path in a bundled instance $(G, \mathcal{D}, \mathbb{B})$. A subword w of $\text{bw}(P)$ is called a *spiral* if it starts and ends with the same bundle B , but all other letters in w are pairwise distinct and different than B .

In other words, a *spiral* corresponds to one ‘turn’ of a path, between two arcs in the same bundle.

Note that a spiral w in $\mathbf{bw}(P)$ corresponds to a closed walk without self-intersections in $G_{\text{bundle}}(G, \mathcal{D}, \mathbb{B})$. Moreover, no arc incident to a terminal may be a part of a spiral, as terminals are of degree one in G , and each terminal is in a single-vertex disc component of \mathcal{D} .

We now make a key observation about the spirals.

Definition 6.9 (spiral cut). Let P be a directed path in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ and let w be a spiral on P . Assume that the first and the last symbol of w is B and let $b^1, b^2 \in B$ be two arcs on P that correspond to these two occurrences of B in w . Then a *spiral cut* associated with the spiral w of the path P is a closed curve γ that consists of a subset of the drawing of P between a midpoint of b^1 and a midpoint of b^2 , and a drawing of a directed path in the dual of G that connects b^1 and b^2 and uses only the arcs of B (an existence of such path is guaranteed by Lemma 6.3).

We would like to note that a spiral cut is degenerate (see Definition 4.1), as it goes along an arc.

Lemma 6.10. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance, let P be a path in G , and let w be a spiral in $\mathbf{bw}(P)$ that starts and ends with a bundle B , and let b^1 and b^2 be two arcs of B that correspond to the two occurrences of B on w . Let γ be the spiral cut associated with the spiral w on P . Then any path P' that does not contain any vertex that lies on γ (i.e., does not contain any inner vertex of the subpath of P that corresponds to the spiral w) intersects γ at most once. Such an intersection, if exists, is contained in an arc of B between b^1 and b^2 and P' traverses γ in the same direction (i.e., either from the inside to the outside or from the outside to the inside) as the path P .*

In particular, the path P does not intersect γ except for the subpath that corresponds to the spiral w .

Proof. The statement follows from the fact that $(\gamma \cap G) \setminus P$ is contained in the arcs of B between b^1 and b^2 ; moreover, all these arcs are directed in the same direction as b^1 and b^2 . \square

Let P be a path in G and let w be a spiral in $\mathbf{bw}(P)$. Assume moreover that B contains only bundle arcs with both endpoints on disc components. Then, knowing only the spiral w (as a bundle word), for each bundle B that is not in w , we know on which side of w it lies, and the notions of bundles *inside* w and *outside* w are well-defined. Note that this statement is not true if w contains a subword B_1B_2 , where B_1 leads to some ring component H^{ring} and B_2 leaves H^{ring} on the same side as B_1 : in this situation from w we cannot deduce on which side of the spiral cut associated with the spiral w on P lies the other side of the ring component H^{ring} .

Note the following.

Corollary 6.11. *Let P be a directed path in a bundled instance $(G, \mathcal{D}, \mathbb{B})$, let w be a spiral in P and let γ be the spiral cut associated with the spiral w on P . Assume $\mathbf{bw}(P) = u_1wu_2$. Then one of the following holds.*

- *The word u_1 contains only bundles from w or outside γ and the word u_2 contains only bundles from w or inside γ . Moreover, for any path P' that is vertex-disjoint with P there exists a partition $\mathbf{bw}(P') = u'_1u'_2$ (where u'_1 or u'_2 may be empty) such that u'_1 contains only bundles from w or outside γ and u'_2 contains only bundles from w or inside γ .*
- *The same situation happens as in the previous point, but with the role of inside and outside exchanged: u_1 and u'_1 are allowed to contain only bundles from w or inside γ , whereas u_2 and u'_2 are allowed to contain only bundles from w and outside γ .*

If w contains only bundles with endpoints in disc components, then the statements ‘inside/outside γ ’ can be replaced with ‘inside/outside w ’.

A similar statement as Corollary 6.11 is also true if we consider paths traversing a single bundle. Consider the following definitions.

Definition 6.12 (bundle profile). Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance, let $B = (b_1, b_2, \dots, b_s)$ be a bundle in G and let $(P_i)_{i=1}^k$ be a solution to k -DPP in G . Then a *bundle profile* of the bundle B , denoted $\mathbf{bp}(B, (P_i)_{i=1}^k)$, is a word over the alphabet $\{1, 2, \dots, k\}$ constructed in the following manner: we inspect arcs b_1, b_2, \dots, b_s of the bundle B in this order and whenever $b_j \in P_i$, we append the symbol i to the word $\mathbf{bp}(B, (P_i)_{i=1}^k)$.

We sometimes shorten $\mathbf{bp}(B)$ if the solution $(P_i)_{i=1}^k$ is clear from the context.

Definition 6.13 (spiral pack). Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance, let $B = (b_1, b_2, \dots, b_s)$ be a bundle in G and let $(P_i)_{i=1}^k$ be a solution to k -DPP in G . A *spiral pack* is a subword w of the word $\mathbf{bp}(B)$ such that w starts and ends with the same symbol, but all other symbols of w are pairwise distinct and different than the starting symbol of w .

Lemma 6.14. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance, let $B = (b_1, b_2, \dots, b_s)$ be a bundle in G and let $(P_i)_{i=1}^k$ be a solution to k -DPP in G . Let $1 \leq j(i, 1), j(i, 2), \dots, j(i, J(i)) \leq s$ be such indices that $b_{j(i, 1)}, b_{j(i, 2)}, \dots, b_{j(i, J(i))}$ are precisely the arcs of B that appear on P_i , in the order of their appearance on P_i . Then:*

1. for any $1 \leq i \leq k$ and $1 \leq \iota < J(i)$, the subword of $\mathbf{bw}(P_i)$ between the occurrences of B that correspond to $j(i, \iota)$ and $j(i, \iota + 1)$ is a spiral;
2. $j(i, 1) < j(i, 2) < j(i, 3) < \dots < j(i, J(i))$ or $j(i, 1) > j(i, 2) > j(i, 3) > \dots > j(i, J(i))$;
3. if for some $i \neq i'$, the order from the previous point is different (i.e., $J(i), J(i') > 1$, but $j(i, 1) < j(i, 2)$ and $j(i', 1) > j(i', 2)$) then either $j(i, \iota) > j(i', \iota')$ for any ι, ι' or $j(i, \iota) < j(i', \iota')$ for any ι, ι' .

Proof. Let $1 \leq i \leq k$ be such an index that $J(i) > 1$ and let w be a subword of $\mathbf{bw}(P_i)$ that starts and ends with B , but contains the symbol B only twice (i.e., w corresponds to a part of the path P_i between two consecutive arcs on B). We first prove that B is a spiral. Assume otherwise: as w does not contain B except for the first and last symbol, w needs to contain a different bundle twice, thus w needs to contain a spiral w' . Moreover, this spiral does not contain B . However, $\mathbf{bw}(P_i)$ contains B both before the spiral w' and after it, a contradiction to Corollary 6.11.

We infer that, by Lemma 6.10, the path P_i does not intersect the spiral cut associated with w except for the spiral w itself, and this spiral cut separates the arcs of B that lie on P_i after w from the arcs of B that lie before w . This proves that $j(i, 1), j(i, 2), \dots, j(i, J(i))$ are ordered monotonically.

Take any other path $P_{i'}$, $i' \neq i$, and assume that $P_{i'}$ contains an arc of B that lies on B between the arcs of B on the spiral w . In this situation, $P_{i'}$ intersects the spiral cut associated with w . By Lemma 6.10, $P_{i'}$ contains the arcs of B in the same order as P_i , and the lemma is proven. \square

Corollary 6.15. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance, let $B = (b_1, b_2, \dots, b_s)$ be a bundle in G and let $(P_i)_{i=1}^k$ be a solution to k -DPP in G . Let w be a spiral pack in $\mathbf{bp}(B)$ and assume $\mathbf{bp}(B) = u_1 w u_2$. Then if a symbol $\iota \in \{1, 2, \dots, k\}$ appears both in u_1 and in u_2 , then ι appears in w .*

Proof. By Lemma 6.14, if i is the first and last symbol of w , then the subpath of P_i that connects the arcs that correspond to the first and the last symbol of w is a spiral. If a symbol $i' \neq i$ does not appear in w , the path $P_{i'}$ does not intersect the spiral cut associated with the aforementioned spiral on P_i and cannot contain arcs of B both inside and outside this spiral. \square

We now investigate a situation where the same spiral appears on $\mathbf{bw}(P)$ several times in a row.

Definition 6.16 (spiraling ring). Let P be a directed path in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ and let u be a word over \mathbb{B} , such that u contains each symbol of \mathbb{B} at most once, u contains only bundles with both endpoints in disc components, and $\mathbf{bw}(P)$ contains $u^r B$ as a subword for some $r \geq 2$, where B is the first symbol of u . Let w_1 and w_2 be the first and the last of the r spirals uB of $\mathbf{bw}(P)$ that appear in the subword $u^r B$ and let γ_1 and γ_2 be the spiral cuts associated with these spirals; in case $r = 2$ we slightly modify γ_1 and γ_2 around the midpoint of the second arc of $P \cap B$ so that they do not intersect. The closed area of the plane contained between γ_1 and γ_2 is called a *spiraling ring*. The curve γ_1 is the *input border* of the spiraling ring, and the curve γ_2 is the *output border*.

Note that, by Lemma 6.10, the input and output borders do not intersect each other, and any spiraling ring is homeomorphic to a closed ring.

Lemma 6.17. *Let P be a directed path in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ and let u be a word over \mathbb{B} , such that u contains each symbol of \mathbb{B} at most once, u contains only bundles with both endpoints in disc components, and $\mathbf{bw}(P)$ contains $u^r B$ as a subword for some $r \geq 2$, where B is the first symbol of u . Let A_R be the spiraling ring associated with $u^r B$ and let γ_1 and γ_2 be its input and output borders. Let P' be a directed path in G that is contained in A_R , except for possibly the first and last arc that may belong to the bundle B intersecting γ_1 and γ_2 , respectively. Then*

1. $\mathbf{bw}(P')$ is a subword of $u^{r'}$ for some integer r' ;
2. if P' starts and ends with an arc of B intersecting γ_1 and γ_2 , respectively, then $\mathbf{bw}(P') = u^{r'} B$ for some integer $r' \geq 1$;
3. if, additionally to the previous point, we assume that P' is vertex-disjoint with P , then $r' = r - 1$, i.e., $\mathbf{bw}(P') = u^{r-1} B$.

Proof. Let $u = B_1 B_2 \dots B_s$, where $B_1 = B$. For $1 \leq j \leq s$ let H_j be the component of \mathcal{D} that contains the ending points of arcs in the bundle B_j . Note that, although the bundles B_j are pairwise distinct, the components H_j may be equal for different indices j . For $1 \leq j \leq s$, denote by G_j the subgraph of H_j contained between the subpath of P between arcs on B_j and B_{j+1} (with $B_{s+1} = B_1 = B$) that is contained in γ_1 , and the subpath of P between arcs on B_j and B_{j+1} that is contained in γ_2 (including paths and vertices on these paths as well). As the bundles B_j are pairwise distinct, the graphs G_j are pairwise disjoint. Moreover, the subgraph of G contained in the spiraling ring A_R consists of the graphs G_j , $1 \leq j \leq s$, and parts of bundles B_j , $1 \leq j \leq s$ that connect G_j with G_{j+1} ; note that here we used the fact that all components H_j are disc components.

Therefore, if some path, that is contained in the subgraph of G contained in A_R , starts in G_j and ends in $G_{j'}$, then its bundleword equals $B_j B_{j+1} \dots B_s u^{r'} B_1 \dots B_{j'-1}$ for some integer r' (or $B_j B_{j+1} \dots B_{j'-1}$ provided that $j \leq j'$). This proves the first two claims of the lemma. For the third claim, we apply the second claim for $r - 1$ subwords $u^2 B$ of the subword $u^r B$ of $\mathbf{bw}(P)$. \square

Corollary 6.18. *Let P be a directed path in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ and let u be a word over \mathbb{B} , such that u contains each symbol of \mathbb{B} at most once, u contains only bundles with both endpoints*

in disc components, and $\mathbf{bw}(P)$ contains u^r as a subword, for some $r \geq 2$. Let B be the first symbol of u and let $w = uB$ be any of the $r - 1$ spirals w contained in the subword u^r of $\mathbf{bw}(P)$. Then for any directed path P' in G that contains arcs from bundles both outside and inside w and is vertex-disjoint with P , the bundle word of P' contains $u^{r-2}B$ as a subword.

Proof. We apply Lemma 6.17 for the spiraling ring that corresponds to the subword $u^{r-1}B$ of $\mathbf{bw}(P)$. \square

6.3 Structure of bundle words and bundle profiles

Corollaries 6.11 and 6.15 imply some structure on bundle words and bundle profiles. To recognize it, we need the following observation.

Definition 6.19 (spiral decomposition of a word). Let u be a word over an alphabet Σ . Let s and $(r_j)_{j=1}^s$ be positive integers and let $(u_j)_{j=1}^s$ be words over Σ . We call $u_1^{r_1}u_2^{r_2}\dots u_s^{r_s}$ a *spiral decomposition* of u if:

1. $u = u_1^{r_1}u_2^{r_2}\dots u_s^{r_s}$;
2. each word u_j contains each symbol of Σ at most once;
3. for each $1 \leq j < s$, there exists a symbol of Σ that appears in exactly one of the two words u_j and u_{j+1} ;
4. for each letter $\sigma \in \Sigma$ that appears in u , there exists $j(\sigma, 1), j(\sigma, 2)$ such that σ appears in u_j if and only if $j(\sigma, 1) \leq j \leq j(\sigma, 2)$.

Lemma 6.20. *Let u be a word over an alphabet Σ and let $u_1^{r_1}u_2^{r_2}\dots u_s^{r_s}$ be a spiral decomposition of u . Then $s \leq 2|\Sigma|$.*

Proof. The claim follows from the fact that, for each $1 \leq i \leq s$, there exists a symbol $\sigma \in \Sigma$, such that u_i is either the first word that contains σ , or the last one. \square

Definition 6.21 (nonnested word). A word u over an alphabet $|\Sigma|$ is called *nonnested* if it satisfies the following property: for any letters $a, b \in \Sigma$, $a \neq b$, and any (possibly empty) words w_1, w_2, w_3 such that $u = w_1aw_2aw_3$, if b appears in w_1 and w_3 , then b appears in w_2 as well.

Lemma 6.22. *Let u be a nonnested word over an alphabet $|\Sigma|$. Then one can in polynomial time compute a spiral decomposition $u_1^{r_1}u_2^{r_2}\dots u_s^{r_s}$ of u .*

Proof. Consider the following greedy algorithm that decomposes u as $v_1v_2\dots v_q$: construct v_j consecutively, taking v_j to be the longest possible word that does not contain the same letter of Σ twice. By construction, the words v_j contain each letter of Σ at most once and v_{j+1} starts with a letter that appears in v_j for each $1 \leq j < q$.

Now assume that $\sigma \in \Sigma$ appears both in v_{i_1} and v_{i_3} , but does not appear in v_{i_2} for some $i_1 < i_2 < i_3$. Let σ' be the first letter of v_{i_2+1} ; by the properties of the words v_i , σ' appears in v_{i_2} and $\sigma' \neq \sigma$. However, this contradicts the assumptions on the word u for $a = \sigma'$ and $b = \sigma$. Therefore, for each letter $\sigma \in \Sigma$ that appears in u , there exist $i(\sigma, 1), i(\sigma, 2)$ such that σ appears in v_i iff $i(\sigma, 1) \leq i \leq i(\sigma, 2)$.

Now assume that, for some $1 \leq i < q$, v_i and v_{i+1} consist of the same letters of Σ , but in different order. Let $v_i = v\sigma_i v'_i$, $v_{i+1} = v\sigma_{i+1} v'_{i+1}$ where $\sigma_i \neq \sigma_{i+1}$, i.e., the first position on which v_i and v_{i+1} differ is $|v| + 1$. However, in this case u contains $v'_i v \sigma_{i+1}$ as a subword that does contain

the letter σ_{i+1} twice and does not contain σ_i , whereas both $v\sigma_i$ and v'_{i+1} contain σ_i , a contradiction with the properties of u .

Therefore, if we define u_1, u_2, \dots, u_s as the sequence of pairwise distinct words of the sequence v_1, v_2, \dots, v_q , and r_j as the number of consecutive occurrences of u_j in the sequence v_1, v_2, \dots, v_q , we obtain a correct spiral decomposition of u . \square

By Corollaries 6.11 and 6.15, both bundle words and bundle profiles are nonnested. We infer the following decomposition statements.

Corollary 6.23 (Bundle word decomposition). *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance and let P be a path in G . Then there exists a spiral decomposition $\mathbf{bw}(P) = u_1^{r_1} u_2^{r_2} \dots u_s^{r_s}$ where $s \leq 2|\mathbb{B}|$. Moreover, if u_j is a cyclic shift of $u_{j'}$ for some $1 \leq j < j' \leq s$ and $r_j, r_{j'} > 1$, then u_j contains a bundle incident to a ring component.*

Proof. The first part of the corollary is straightforward. For the second part, first note that, as u_j and u_{j+1} need to differ on at least one symbol, we have $j' > j + 1$. As u_j and $u_{j'}$ are cyclic shifts, u_{j+1} contains a bundle B that does not appear in u_j nor in $u_{j'}$. However, if $r_j, r_{j'} > 1$ then the terms $u_j^{r_j}$ and $u_{j'}^{r_{j'}}$ contain the same spiral, a contradiction to Lemma 6.10 unless u_j contains a bundle incident to a ring component. \square

Corollary 6.24 (Bundle profile decomposition). *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance, let B be a bundle in G and let $(P_i)_{i=1}^k$ be a solution to k -DPP in G . Then there exists a spiral decomposition $\mathbf{bp}(P) = u_1^{r_1} u_2^{r_2} \dots u_s^{r_s}$ where $s \leq 2|\mathbb{B}|$.*

We also note the following.

Lemma 6.25. *Let w be a word over an alphabet Σ and assume w contains a subword $v^q\sigma$, where v contains each letter of Σ at most once, σ is the first letter of v and $q \geq 3$. Then any spiral decomposition of w needs to contain a term u^r , where u is a cyclic shift of v and $r \geq q - 1$.*

Proof. Consider a spiral decomposition $w = u_1^{r_1} u_2^{r_2} \dots u_s^{r_s}$. Let τ be any letter of the middle $v^{q-2}\sigma$ subword of $v^q\sigma$ in w . Assume that this occurrence of τ appears in u_j in the spiral decomposition. Then, if $|u_j| > |v|$, u_j contains the same symbol twice, and if $|u_j| < |v|$, u_j is a subword of $v^q\sigma$ and the word $v^q\sigma$ contains the same symbol on both sides of the subword u_j , and this symbol does not appear in u_j . In both cases we have reached a contradiction, so $|u_j| = |v|$ and u_j is a cyclic shift of v . As the choice of τ was arbitrary, we infer that $r_j \geq q - 1$. \square

Lemma 6.26. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance. Let P be a path in G with $\mathbf{bw}(P) = u^3B$ where u contains each bundle at most once, u contains only bundles with both endpoints in disc components, and B is the first symbol of u . Moreover, let Q be a path in G with $\mathbf{bw}(Q) = vB'$ where v is a permutation of u that is not a cyclic shift of u and B' is the first symbol of v . Then P and Q share an internal vertex.*

Proof. Assume the contrary: P and Q do not intersect, except possibly for the endpoints.

Let b_1, b_2, b_3, b_4 be the four arcs of $P \cap B$ in the order of their appearance on P and let R be the subpath of P that starts with b_2 and ends with b_3 ; note that $\mathbf{bw}(R) = uB$. Let γ_R and γ_Q be spiral cuts corresponding to R and Q , respectively.

We claim that γ_Q and γ_R do not intersect. Assume the contrary; let $p \in \gamma_Q \cap \gamma_R$. As Q and R are vertex-disjoint except for possibly the endpoints, the point p must belong to $\gamma_Q \setminus Q$ or $\gamma_R \setminus R$ (i.e., to the part of γ_Q inside B' or the part of γ_R inside B). If $p \in \gamma_Q \setminus Q$ then, as γ_R visits each bundle of u exactly once, the midpoints of the two arcs of $Q \cap B'$ lie on different sides of γ_R , and

Q needs to contain an arc of B that lies between b_2 and b_3 . Otherwise, if $p \in (\gamma_R \setminus R) \cap Q$, then, as $\gamma_R \setminus R$ connects b_2 with b_3 inside B , Q again contains an arc of B that lies between b_2 and b_3 .

If Q is contained in the spiraling ring A_P associated with the path P (recall $\mathbf{bw}(P) = u^3B$) then, by the first claim of Lemma 6.17, $\mathbf{bw}(Q)$ is a subword of $u^{r'}$ for some r' , a contradiction to the assumption that v is not a cyclic shift of u . Otherwise, Q contains an arc $b' \in B$, contained either between b_1 and b_2 or between b_3 and b_4 . Consider the first case; the second one is symmetrical. Let P' be the subpath of P between b_1 and b_3 ; note that $\mathbf{bw}(P') = u^2B$. Let A'_R be the spiraling ring associated with P' and let Q' be the subpath of Q contained in A'_R that contains the aforementioned arcs of B between b_1 and b_2 as well as between b_2 and b_3 . By the third claim of Lemma 6.17, applied to the spiraling ring A'_R and the path Q' , $\mathbf{bw}(Q') = uB$, a contradiction. Hence, Q cannot contain an arc $b \in B$ that lies between b_2 and b_3 , and γ_Q and γ_R do not intersect.

As v is a permutation of u , but not a cyclic shift of u , there exist three pairwise distinct bundles B_1, B_R, B_Q and a component D such that γ_R first goes through the part of the plane occupied by B_1 , then through D and then through the part of the plane occupied by B_R , whereas γ_Q goes through the part of the plane occupied by B_1 , then through D and then through B_Q (see Figure 38). Let γ'_Q be the subcurve of γ_Q between the midpoint of the traverse of B_1 and the midpoint of the traverse of B_R that includes the traverse of B_Q . Obtain a curve γ by closing γ'_Q as follows: from the midpoint of the traverse of B_R go along the bundle B_R in the dual of G to the curve γ_R , follow γ_R backwards through G to the midpoint of the traverse of γ_R through B_1 and then go along the bundle B_1 in the dual of G to the start of γ'_Q .

Note that the only intersection of γ with γ_R is the part of γ_R between the midpoint of its traverse of B_1 and the midpoint of its traverse of B_R , containing the traverse of D . As γ_Q lies on the same side of γ_R both in B_1 and in B_R , γ_R leaves γ to the same side in B_1 and in B_R . That is, the start and the end of the part of curve γ_R that traverses B_1 , D and B_R lies on one of the two sides of γ . However, since in B_1 , B_Q and B_R the curve γ_Q needs to lie on the same side of γ_R , the part of the curve γ_R that traverses B_Q lies on the other side of γ , a contradiction. This finishes the proof of the lemma. \square

Corollary 6.27. *Let \mathcal{P} be a family of pairwise vertex-disjoint paths in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ with fixed bundle word decompositions for each $P \in \mathcal{P}$. Assume there exists a path $P_0 \in \mathcal{P}$ and a subword u^rB of $\mathbf{bw}(P_0)$ such that u contains each bundle at most once, u contains only bundles with both endpoints in disc components, B is the first symbol of u and $r \geq 4$. Moreover, assume that any path P starts and ends with an arc that belongs to a bundle that is not contained in u . Then, for any $P \in \mathcal{P}$ either*

1. *the first and the last arc P lie on the same side of the spiral uB and there does not exist any term v^q in the bundle word decomposition of $\mathbf{bw}(P)$ such that $q > 1$ and v is a permutation of u ;*
2. *the first and the last arc of P lie on different sides of the spiral uB and there exists exactly one term v^q in the bundle word decomposition of $\mathbf{bw}(P)$ where $q > 1$ and v is a permutation of u ; moreover, in this case v is a cyclic shift of u and $q \geq r - 2$.*

Proof. First, observe that by Corollary 6.23 there is at most one v^q in the bundle word decomposition of $\mathbf{bw}(P)$, where v is a cyclic shift of u and $q > 1$. Let us prove that for any $P \in \mathcal{P}$ the bundle word $\mathbf{bw}(P)$ does not contain a subword v^2 where v is a permutation of u that is not a cyclic shift of u . This is clearly true for $P \neq P_0$, due to Lemma 6.26 and the assumption of the subword u^rB in $\mathbf{bw}(P_0)$. Moreover, any such subword v^2 of $\mathbf{bw}(P_0)$ would be disjoint with one of the two subwords $u^{r-1}B$ of u^rB , which is in turn a subword of $\mathbf{bw}(P_0)$. As $r \geq 4$, the claim follows from Lemma 6.26,




Figure 38: An illustration of the proof of Lemma 6.26.

as we have two disjoint (except for possibly the endpoints) subpaths of P_0 with bundlewords v^2 and u^3B .

Now let us consider two cases. First assume that $P = P_0$. Then the endpoints of P_0 lie on different sides of the spiral uB , and the bundle word decomposition of $\mathbf{bw}(P_0)$ obviously contains u^r .

Finally, assume that $P \neq P_0$. Note that if the endpoints of $P \in \mathcal{P}$ lie on different sides of the spiral uB then, by the third claim of Lemma 6.17 combined with Lemma 6.25, the bundle word decomposition of P needs to contain the term v^q for v being a cyclic shift of u and $q \geq r - 2 > 1$. In the reverse direction, if $\mathbf{bw}(P)$ contains a subword v^2 where v is a cyclic shift of u , then $\mathbf{bw}(P)$ contains a subword uB and, by Corollary 6.11, P has endpoints on different sides of the spiral uB . \square

6.4 Ring components: basics

The ring components have more complicated topology structure than the disc ones, but, thanks to the isolation property, we are able to prove that, if we are given a YES-instance to k -DPP, there exists a solution with the interaction with ring components that is somehow bounded. In this section we give some basic notions towards proving this statement.

Assume we are given a bundled instance $(G, \mathcal{D}, \mathbb{B})$, such that \mathcal{D} has isolation (Λ, d) , $\Lambda, d > 0$. Consider a ring component H^{ring} with isolation components $H_{IN,\lambda}^{\text{disc}}$ and $H_{OUT,\lambda}^{\text{disc}}$, $1 \leq \lambda \leq \Lambda$. Fix a choice of the d concentric cycles in the subgraph induced by the vertices of $H_{IN,\lambda}^{\text{disc}}$ promised by the definition of an isolating component and let $C_{IN,\lambda}(H^{\text{ring}})$ be the innermost of these cycles. Similarly, let $C_{OUT,\lambda}(H^{\text{ring}})$ be the outermost of a fixed choice of a alternating sequence of cycles in $H_{OUT,\lambda}^{\text{disc}}$.

Definition 6.28 (closure of a ring component). Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance, such that \mathcal{D} has isolation (Λ, d) , $\Lambda, d > 0$. Let H^{ring} be a ring component in \mathcal{D} and let $1 \leq \lambda \leq \Lambda$. Then $\text{cl}^\lambda(H^{\text{ring}})$, the *level- λ closure of H^{ring}* , is the subgraph of G that contains $H^{\text{ring}} \cup \bigcup_{1 \leq \lambda' \leq \lambda} H_{IN, \lambda'}^{\text{disc}} \cup H_{OUT, \lambda'}^{\text{disc}}$ as well as all bundle arcs of G that are contained between $C_{IN, \lambda}(H^{\text{ring}})$ and $C_{OUT, \lambda}(H^{\text{ring}})$ or on one of these cycles. Moreover, we define the closure of H^{ring} , $\text{cl}(H^{\text{ring}})$ as the Λ -level closure of H^{ring} , and the 0-level closure of H^{ring} , $\text{cl}^0(H^{\text{ring}})$, as H^{ring} itself.

Note that, since each disc component may serve as an isolation only to one ring component, closures are pairwise disjoint. By splitting some bundles incident to the isolation components of H^{ring} into smaller bundles, we may assume that either none or all arcs of a single bundle are contained in the closure of a ring component. By Lemma 6.10, a cycle in G may contain at most one arc from a single bundle. Therefore, the aforementioned splitting operation partitions each bundle with arcs with both endpoints in an isolation component into at most three parts, and, consequently, at most triples the number of such bundles.

Definition 6.29 (normal, ring and isolation bundles). Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance, such that \mathcal{D} has isolation (Λ, d) , $\Lambda, d > 0$. A bundle $B \in \mathbb{B}$ is a *ring bundle* if it is contained in a closure of some ring component, and a *normal bundle* otherwise. A *level* of a ring bundle B contained in the closure of a ring component H^{ring} , is the minimum nonnegative integer λ , such that B contains arcs incident to $\text{cl}^\lambda(H^{\text{ring}})$. A ring bundle is an *isolation bundle* if it has both endpoints in the same isolating disc component.

Definition 6.30 (ring part). Let P be a directed path in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ with isolation (Λ, d) , $\Lambda, d > 0$. Then a maximal subword of $\text{bw}(P)$ that consists of only ring bundles and contains at least one level- λ ring bundle for some $\lambda < \Lambda$ is called a *ring part* of $\text{bw}(P)$. The normal bundles preceding and succeeding a ring part of $\text{bw}(P)$ are called the *predecessor* and *successor* of the ring part, respectively.

For a ring part w of $\text{bw}(P)$, a *ring part* of P is maximal possible subpath of P that corresponds to w . The *predecessor* and the *successor* of a ring part of P is an arc that precedes (resp. succeeds) the ring part on the path P .

Note that the predecessor or successor of a ring part may not be defined if P ends or starts in the closure of the ring component; however, they are always well-defined for paths connecting a terminal pair.

Note also that ring parts of a bundle word are pairwise disjoint and there is at least one normal bundle between any ring parts in the bundle word $\text{bw}(P)$. Moreover, as each disc component may serve as an isolation component to at most one ring component and closures of ring components are pairwise disjoint, each ring part of a bundle word contains bundles from a closure of a single ring component and the corresponding ring part of a path is contained in the closure of this ring component.

Definition 6.31 (isolation passage). Let P be a directed path in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ of isolation (Λ, d) , $\Lambda, d > 0$, let H^{ring} be a ring component in \mathcal{D} and let H^{disc} be any isolation component of H^{ring} such that P does not start nor end in H^{disc} . Then an *isolation passage* of P through H^{disc} is any maximal subpath of P that is contained in $\text{cl}(H^{\text{ring}})[V(H^{\text{disc}})]$ (i.e., H^{disc} with arcs of its isolation bundles), but the arc preceding and succeeding the subpath on P lie on different faces of $\text{cl}(H^{\text{ring}})[V(H^{\text{disc}})]$. The level of the isolation passage is the level of H^{disc} .

Note that P starts and ends outside H^{disc} for example, if P connects a terminal pair in G . Moreover, an isolation passage of P through H^{disc} is preceded and succeeded by an arc connecting

H^{disc} with a neighbouring isolation component, a normal bundle (possible if H^{disc} is of level Λ) or H^{ring} (possible if H^{disc} is of level 1).

Definition 6.32 (ring passage, ring visitor). Let P be a directed path in a bundled instance $(G, \mathcal{D}, \mathbb{B})$. Let w be a ring part in $\text{bw}(P)$ in the closure of the ring component H^{ring} , for which the predecessor and the successor are well-defined. Then w is a *ring passage* if the predecessor and the successor of w lie on different faces of the closure of H^{ring} (equivalently, lie on different faces of H^{ring} itself), and *ring visitor* otherwise.

Definition 6.33 (level- λ ring passage). Let P be a directed path in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ with isolation (Λ, d) , $\Lambda, d > 0$. Let $0 \leq \lambda \leq \Lambda$, let w be a ring part in $\text{bw}(P)$ in the closure of the ring component H^{ring} and assume that P starts and ends outside $\text{cl}^\lambda(H^{\text{ring}})$. Then any maximal subpath of a ring part w of P that is contained in $\text{cl}^\lambda(H^{\text{ring}})$ that has a preceding and succeeding arc on P contained in different faces of $\text{cl}^\lambda(H^{\text{ring}})$ is called a *level- λ ring passage*.

Note that a level- Λ ring passage is simply a ring passage.

We also note the following.

Lemma 6.34. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance with positive isolation and let P be an arbitrary path connecting a terminal pair in G . Then any spiral in $\text{bw}(P)$ contains at least one normal bundle or is contained in a ring passage of $\text{bw}(P)$. In particular, any ring visitor of $\text{bw}(P)$ consists of pairwise distinct bundles.*

Proof. Let w be a spiral in $\text{bw}(P)$ that contains only ring bundles and let γ be the corresponding spiral cut. As w contains only ring bundles, the curve γ is completely contained in $\text{cl}(H^{\text{ring}})$ for one ring component H^{ring} . If the curve γ separates the inner face of $\text{cl}(H^{\text{ring}})$ from the outer face, Lemma 6.10 implies that the predecessor and the successor of w lie in different faces of $\text{cl}(H^{\text{ring}})$. In other case, the graph enclosed inside γ is contained in $\text{cl}(H^{\text{ring}})$. This contradicts Lemma 6.10, as P connects a terminal pair in G and there is no terminal contained inside γ . \square

By the definition of ring isolation, all bundle arcs incident to at least one vertex of $\text{cl}^\lambda(H^{\text{ring}})$ (for some $0 \leq \lambda \leq \Lambda$ and some ring component H^{ring}), but not belonging to $\text{cl}^\lambda(H^{\text{ring}})$, lie either in the outer face of $\text{cl}^\lambda(H^{\text{ring}})$, or in one other face of $\text{cl}^\lambda(H^{\text{ring}})$ inside $C_{IN,\lambda}(H^{\text{ring}})$ (or H^{ring} if $\lambda = 0$).

Lemma 6.35. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance of isolation (Λ, d) , $\Lambda, d > 0$. Then in polynomial time we may compute a set of reference curves $\gamma_{\text{ref}}^\lambda(H^{\text{ring}})$ for all $0 \leq \lambda \leq \Lambda$ and H^{ring} being a ring component of \mathcal{D} in such a manner that:*

1. $\gamma_{\text{ref}}^\lambda(H^{\text{ring}})$ connects the outer face of $\text{cl}^\lambda(H^{\text{ring}})$ with the inner face;
2. for each bundle B , either $\gamma_{\text{ref}}^\lambda(H^{\text{ring}})$ does not intersect B , or contains a subcurve that intersects each arc of B and nothing else,
3. $\gamma_{\text{ref}}^\lambda(H^{\text{ring}})$ is a subcurve of $\gamma_{\text{ref}}^{\lambda+1}(H^{\text{ring}})$ for $0 \leq \lambda < \Lambda$ and $\gamma_{\text{ref}}^\lambda(H^{\text{ring}}) \cap \text{cl}^\lambda(H^{\text{ring}}) = \gamma_{\text{ref}}^{\lambda+1}(H^{\text{ring}}) \cap \text{cl}^\lambda(H^{\text{ring}})$;
4. $\gamma_{\text{ref}}^\lambda(H^{\text{ring}}) \cap H^{\text{disc}} = \emptyset$ for any disc component H^{disc} that isolates H^{ring} .

Proof. For each ring component H^{ring} , $\Gamma \in \{IN, OUT\}$ and $1 \leq \lambda \leq \Lambda$, choose an undirected path in the dual of G that connects the outer and inner face of $\text{cl}(H^{\text{ring}})[V(H_{\Gamma,\lambda}^{\text{disc}})]$ and contains bundle arcs only. Connect these paths with parts of (again undirected) cycles in the dual of G that contain

bundle arcs connecting $H_{\Gamma,\lambda}^{\text{disc}}$ with $H_{\Gamma,\lambda-1}^{\text{disc}}$ (or H^{ring} if $\lambda = 1$). The drawings of these undirected paths in the dual of G yield the desired paths $\gamma_{\text{ref}}^\lambda(H^{\text{ring}})$; the paths $\gamma_{\text{ref}}^\lambda(H^{\text{ring}})$ are their appropriate subcurves. \square

In this manner, the subgraph $\text{cl}^\lambda(H^{\text{ring}})$ becomes a rooted ring, and we may use the notion of a winding number with respect to $\gamma_{\text{ref}}^\lambda(H^{\text{ring}})$. In particular, given a path P that connects a terminal pair in G , and a level- λ ring passage w of $\text{bw}(P)$, we define the winding number $\gamma_{\text{ref}}^\lambda(H^{\text{ring}})(w)$ of the level- λ ring passage w as the winding number of the corresponding subpath of the path P .

6.5 Bundle words with holes

In this section we introduce a modification of the definition of bundle words that takes care also of ring components. Informally speaking, if we want to apply Schijver's cohomology algorithm, we need not only to know bundle words of the paths, but also the number of turns a path makes when it makes a ring passage.

Definition 6.36 (bundle word with ring holes). Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance of isolation (Λ, d) for $\Lambda, d > 0$. Let $0 \leq \lambda < \Lambda$. A pair $((p_j)_{j=0}^q, (w_j)_{j=1}^q)$ is called a *bundle word with level- λ ring holes* if:

1. each p_j , $0 \leq j \leq q$ is a bundle word in $(G, \mathcal{D}, \mathbb{B})$ that does not contain two bundles that lie on different sides of $\text{cl}^\lambda(H^{\text{ring}})$ for any ring component H^{ring} ;
2. each w_j , $1 \leq j \leq q$ is an integer;
3. for each $1 \leq j \leq q$, there exists a ring component $H_j^{\text{ring}} \in \mathcal{D}$, such that
 - (a) the last bundle of p_{j-1} is a bundle that contains arcs with ending points in $\text{cl}^\lambda(H_j^{\text{ring}})$ but starting points not in $\text{cl}^\lambda(H_j^{\text{ring}})$;
 - (b) the first bundle of p_j is a bundle that contains arcs with starting points in $\text{cl}^\lambda(H_j^{\text{ring}})$ but ending points not in $\text{cl}^\lambda(H_j^{\text{ring}})$;
 - (c) the two aforementioned bundles lie in different faces of $\text{cl}^\lambda(H_j^{\text{ring}})$ (i.e., one lies in the outer face, and one in the inner face).

We say that a path P is *consistent* with a bundle word with ring holes $((p_j)_{j=0}^q, (w_j)_{j=1}^q)$ if $\text{bw}(P) = p_0 r_1 p_1 r_2 p_2 \dots r_q p_q$, where r_1, r_2, \dots, r_q are exactly the level- λ ring passages of $\text{bw}(P)$, and for each $1 \leq j \leq q$, the subpath of P corresponding to the ring passage r_j has winding number w_j in the level- λ closure of the appropriate ring component.

We first note that, knowing a bundle word with ring holes for some level, we in fact know the bundle words with ring holes for all higher levels.

Definition 6.37 (projection of bundle words with ring holes). Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance of isolation (Λ, d) for $\Lambda, d > 0$ and let $((p_j)_{j=0}^q, (w_j)_{j=1}^q)$ be a bundle word with level- λ ring holes for some $0 \leq \lambda < \Lambda - 1$ such that p_0 starts with a normal bundle and p_q ends with a normal bundle. Let $\lambda < \lambda' < \Lambda$. A *level- λ' projection* of $((p_j)_{j=0}^q, (w_j)_{j=1}^q)$ is a bundle word with level- λ' ring holes $((p'_j)_{j=0}^{q'}, (w'_j)_{j=1}^{q'})$ defined as follows. Let $p = p_0 w_1 p_1 w_2 p_2 \dots w_q p_q$ be a word over alphabet $\mathbb{B} \cup \mathbb{Z}$ and let $p = p'_0 x_1 p'_1 x_2 \dots x_{q'} p'_q$ be such that each x_j , $1 \leq j \leq q'$ is a maximal subword of p that contains at least one integer w_ι , and all the bundles it contains are in $\text{cl}^{\lambda'}(H^{\text{ring}})$ for some ring

component H^{ring} . Let $w'_{j'}$, for $1 \leq j' \leq q'$ be equal to the sum of integers w_j contained in $x_{j'}$ plus $+1$ or -1 for each bundle in $x_{j'}$ that crosses $\gamma_{\text{ref}}^{\lambda'}(H^{\text{ring}})$; the sign depends on the direction of the crossing.

Lemma 6.38. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance of isolation (Λ, d) for $\Lambda \geq 2, d > 0$. Let \mathfrak{p} a bundle word with level- λ ring holes for some $1 \leq \lambda < \Lambda$. and let P be a path consistent with \mathfrak{p} . Assume that \mathfrak{p} starts and ends with a normal bundle. Let $\lambda \leq \lambda' \leq \Lambda$ and let \mathfrak{p}' be the level- λ' projection of \mathfrak{p} . Then P is consistent with \mathfrak{p}' as well.*

Proof. Let $\mathfrak{p} = ((p_j)_{j=0}^q, (w_j)_{j=1}^q)$ and let q', p'_j, x_j and w'_j be as in the definition of the level- λ' projection; in particular $\mathfrak{p}' = ((p'_j)_{j=0}^{q'}, (w'_j)_{j=1}^{q'})$. Let $\mathfrak{bw}(P) = p_0 r_1 p_1 r_2 \dots r_q p_q$ and let r_1, r_2, \dots, r_q be the level- λ ring passages of P . Note that, as p_0 starts with a normal bundle and p_q ends with a normal bundle, by the definition of level- λ' ring passages, $\mathfrak{bw}(P) = p'_0 r'_1 p'_1 \dots r'_q p'_q$ where $r'_1, r'_2, \dots, r'_{q'}$ are exactly the level- λ' ring passages of P . Moreover, as the reference curve $\gamma_{\text{ref}}^{\lambda'}(H^{\text{ring}})$ contains $\gamma_{\text{ref}}^{\lambda}(H^{\text{ring}})$, does not contain any point of $\text{cl}^{\lambda}(H^{\text{ring}})$ outside $\gamma_{\text{ref}}^{\lambda}(H^{\text{ring}})$ and is disjoint with isolating components of H^{ring} , w'_j equals the winding number of r'_j . \square

Finally, we note an easy fact that any path in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ yields a unique bundle word with holes for each fixed level λ .

Lemma 6.39. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance of isolation (Λ, d) for $\Lambda \geq 2, d > 0$. Let P be a path in G , whose both starting and ending points are not contained in any closure of a ring component. Then, for any $0 \leq \lambda < \Lambda$, there exists a unique level- λ bundle word with ring holes that is consistent with P .*

Proof. First, let us construct a level-0 bundle word with ring holes for P . We start with the bundle word $\mathfrak{bw}(P)$. For any level-0 ring passage P^* of P , let b_1 and b_2 be the arcs on P preceding and succeeding P^* and let $B_1, B_2 \in \mathbb{B}$ be such that $b_1 \in B_1, b_2 \in B_2$; take the subword $B_1 \mathfrak{bw}(P^*) B_2$ of $\mathfrak{bw}(P)$ that corresponds to the subpath P^* with arcs b_1 and b_2 and replace $\mathfrak{bw}(P^*)$ with the winding number of P^* in its ring component H^{ring} (note that P^* may contain arcs of bundles with both endpoints in H^{ring}). Performing this operation for each level-0 ring passage of P , we obtain a bundle word with level-0 holes that is consistent with P and, moreover, starts and ends with a normal bundle. By Lemma 6.38, a level- λ projection of this bundle word with level-0 ring holes is a bundle word with level- λ ring holes consistent with P . The uniqueness follows from the fact that in the consistency definition we require that the integers w_j correspond exactly to the level- λ ring passages of P . \square

We also define the following property of bundle words with ring holes that will be useful in the future.

Definition 6.40 (flat bundle word with ring holes). Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance of isolation (Λ, d) for $\Lambda \geq 2, d > 0$. Let \mathfrak{p} a bundle word with level- λ ring holes for some $0 \leq \lambda < \Lambda - 1$ that starts and ends with a normal bundle. We say that \mathfrak{p} is *flat* if the level- $(\lambda + 1)$ projection of \mathfrak{p} does not contain any bundles of level λ or lower.

6.6 Minimal solution

We now define the notion of a *minimal* solution to a bundled instance $(G, \mathcal{D}, \mathbb{B})$. Assume that \mathcal{D} has ring isolation (Λ, d) .

Definition 6.41 (minimal solution). A solution $(P_i)_{i=1}^k$ to k -DPP on a bundled instance $(G, \mathcal{D}, \mathbb{B})$ is called *minimal* if:

1. the words $(\mathbf{bw}(P_i))_{i=1}^k$ have minimal possible total number of bundles that are not isolation bundles;
2. satisfying the above, the words $(\mathbf{bw}(P_i))_{i=1}^k$ have minimal possible total length.

One of the ways we use minimality of a solution is the following.

Lemma 6.42. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance and let $(P_i)_{i=1}^k$ be a minimal solution to k -DPP in G . Assume that for some $1 \leq \iota \leq k$, $\mathbf{bw}(P_\iota)$ contains a subword u^r , where u contains each bundle at most once, u contains only bundles with both endpoints in disc components, and $r \geq 10$. Let A_R be the spiraling ring (as a closed subspace of a plane) associated with the subword u^r in $\mathbf{bw}(P_\iota)$ and borders γ_1 and γ_2 . Let G_R be the subgraph of G consisting of those vertices and arcs of G that lie in $A_R \setminus \gamma_2$. Let f_1 be the face of G_R that contains the part of the plane separated from A_R by γ_1 and let f_2 be defined analogously for γ_2 . Let $I \subseteq \{1, 2, \dots, k\}$ be the set of indices of those paths that intersect G_R . Then, for any set of vertex-disjoint paths $(Q_i)_{i \in I}$ such that for each $i \in I$, Q_i starts in a vertex on f_1 and ends in a vertex on f_2 , the bundle word of each Q_i contains u^{r-10} as a subword.*

Proof. First note that, by Lemma 6.10, the intersection of a path P_i with G_R is either empty or is a single path connecting a vertex on f_1 with a vertex on f_2 .

Let B be the first bundle on u . We treat G_R as a rooted ring with faces f_1 and f_2 , containing the parts of the plane separated by the curves γ_1 and γ_2 from A_R , and with a reference curve being the bundle B in the dual of G .

Assume the contrary: let $(Q_i)_{i \in I}$ be such that $\mathbf{bw}(Q_i)$ does not contain u^{r-10} as a subword. By Lemma 6.17, the winding number of Q_i is at most $r - 10$. By Observation 4.4, the winding number of any Q_i for $i \in I$ is at most $r - 9$.

By Lemmata 6.10 and 6.17, for each $i \in I$, the intersection of P_i with G_R is a path $P_{i,*}$ that connects a vertex on f_1 with a vertex on f_2 such that $B \mathbf{bw}(P_{i,*}) = u^{r-1}$ and, consequently, the winding number of $P_{i,*}$ is $r - 2$. Note that for any $i \notin I$, P_i does not intersect G_R .

By Lemma 4.9, there exists a sequence of vertex-disjoint paths $(P'_{i,*})_{i \in I}$ in G_R such that $P'_{i,*}$ connects the same pair of vertices as $P_{i,*}$, but has winding number at most $r - 3$. Therefore $B \mathbf{bw}(P'_{i,*}) = u_{r(i)}$ for some $r(i) < r - 1$.

For each $1 \leq i \leq k$, construct a path P'_i from P_i by replacing the subpath $P_{i,*}$ with $P'_{i,*}$. As $P_{i,*}$ and $P'_{i,*}$ have the same endpoints, P'_i connects the i -th terminal pair. As $P'_{i,*}$ were pairwise disjoint, and the intersection of $(P_i)_{i=1}^k$ with G_R is exactly $(P_{i,*})_{i \in I}$, the sequence $(P'_i)_{i=1}^k$ is a solution to k -DPP in G . However, $\mathbf{bw}(P_i) = \mathbf{bw}(P'_i)$ for $i \notin I$ and $\mathbf{bw}(P'_i)$ is a proper subsequence of $\mathbf{bw}(P_i)$ for $i \in I$. As $\iota \in I$, this contradicts the minimality of $(P_i)_{i=1}^k$. \square

6.7 Ring components: bound on the interaction

The goal of this section is to show that a minimal solution has a limited interaction with ring components. As a tool, we use the following routing argument.

Lemma 6.43 (one-directional routing). *Let H be a connected graph, embedded in a plane. Let $\ell > 0$ be an integer and let $(p_j)_{j=1}^{2\ell}, (q_j)_{j=1}^{2\ell}$ be pairwise distinct vertices that lie on the outer face of H in the order $p_1, p_2, \dots, p_{2\ell}, q_{2\ell}, q_{2\ell-1}, \dots, q_1$. Moreover, let $(a_j)_{j=1}^\ell, (b_j)_{j=1}^\ell$ be pairwise distinct vertices that lie on the outer face of H in the order $b_1, b_2, \dots, b_\ell, a_\ell, a_{\ell-1}, \dots, a_1$ and such that*

a_1, a_2, \dots, a_ℓ lie on the outer face of H between $p_{2\ell}$ and $q_{2\ell}$ (possibly $p_{2\ell} = a_1$ or $q_{2\ell} = a_\ell$) and b_1, b_2, \dots, b_ℓ lie on the outer face of H between q_1 and p_1 (possibly $p_1 = b_1$ or $q_1 = b_\ell$). Assume that there exists a set of 2ℓ vertex-disjoint paths $(P_j)_{j=1}^{2\ell}$ in H such that for $1 \leq j \leq 2\ell$, P_j starts in p_j and ends in q_j , and a set of ℓ vertex-disjoint paths $(C_j)_{j=1}^\ell$ in H such that for $1 \leq j \leq \ell$, C_j starts in a_j and ends in b_j .

Then there exist a sequence $(R_j)_{j=1}^\ell$ of vertex-disjoint paths in H such that for $1 \leq j \leq \ell$, the path R_j starts in p_j and ends in $q_{j+\ell}$.

Proof. We first note that, by planarity, any set of ℓ vertex-disjoint paths that connects $\{p_1, p_2, \dots, p_\ell\}$ with $q_{\ell+1}, q_{\ell+2}, \dots, q_{2\ell}$ needs to connect p_ι with $q_{\ell+\iota}$ for $1 \leq \iota \leq \ell$. Therefore, by Menger's theorem if the desired paths $(R_j)_{j=1}^\ell$ does not exist, there exists a set $X \subseteq V(H)$, $|X| < \ell$, such that in $H \setminus X$ no vertex $q_{\ell+\iota'}$ is reachable from p_ι for any $1 \leq \iota, \iota' \leq \ell$. However, as $|X| < \ell$, there exists $1 \leq \iota, \iota' \leq \ell$ such that X does not contain any vertex of P_ι , $P_{\ell+\iota}$ and $C_{\iota'}$. Due to the order of the vertices p_j, q_j, a_j and b_j on the outer face of H , the union of these three paths contain a path from p_ι to $q_{\ell+\iota}$, a contradiction. \square

We now use the isolation components $H_{OUT,\lambda}^{\text{disc}}$ and $H_{IN,\lambda}^{\text{disc}}$ to show that there is a bounded number of ring parts in a minimal solution.

Theorem 6.44 (bound on the number of isolation passages). *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance, such that \mathcal{D} has isolation (Λ, d) where $\Lambda > 0$ and $d \geq 2k$. Assume that G is a yes instance to k -DPP and let $(P_i)_{i=1}^k$ be a minimal solution. Then the paths P_i , in total, contain at most $4|\mathbb{B}|^2 k^2$ isolation passages.*

Proof. Consider one ring component $H^{\text{ring}} \in \mathcal{D}$ and one its isolation component H^{disc} . Fix one bundle B^1 that contains arcs with ending points contained in H^{disc} fix one bundle B^2 that contains arcs with starting points in H^{disc} . Consider isolation passages of paths $(P_i)_{i=1}^k$ such that that B^1 contains the arc preceding the isolation part and B^2 contains the arc succeeding it. Note that between B^1 and B^2 the path P_i may contain arcs only in bundles that are isolation bundles of component H^{disc} .

To prove the lemma, we need to show that in a minimal solution there are at most $4k^2$ isolation passages. for a fixed choice of H^{ring} , H^{disc} , B^1 and B^2 . Note that the choice of B^1 determines H^{ring} and H^{disc} , thus there are less than $|\mathbb{B}|^2$ choices.

Consider a choice of H^{ring} , H^{disc} , B^1 and B^2 where there are more than $4k^2$ such isolation passages. Denote all such passages, in the order of their appearance on B^1 , as Q_1, Q_2, \dots, Q_r . Assume Q_j belongs to a path P_{i_j} . Let $\text{prec}(Q_j) \in B^1$ be an arc preceding Q_j on P_{i_j} and $\text{succ}(Q_j) \in B^2$ be an arc succeeding Q_j on P_{i_j} .

Let $H = \text{cl}(H^{\text{ring}})[H^{\text{disc}}]$. By the definition of isolation, all arcs of G that are not contained in H , but are incident to at least one vertex of H , are contained in one of the two faces of H : the one containing H^{ring} and B^2 , and the one containing B^1 . Thus, by planarity, the paths Q_1, Q_2, \dots, Q_r arrive at bundle B^2 in order $Q_{q+1}, Q_{q+2}, \dots, Q_r, Q_1, Q_2, \dots, Q_q$ for some $1 \leq q \leq r$. We assume $q \geq r/2$; the second case is symmetrical. Note that $q \geq r/2$ implies $q > 2k^2$.

Consider an area A_0 of the plane enclosed by:

- the path Q_1 , together with parts of arcs of $\text{prec}(Q_1)$ and $\text{succ}(Q_1)$;
- a drawing of a path in the dual of G , connecting the arc $\text{succ}(Q_1)$ and the arc $\text{succ}(Q_q)$, that intersects only the arcs of B^2 between $\text{succ}(Q_1)$ and $\text{succ}(Q_q)$;
- the path Q_q , together with parts of arcs of $\text{prec}(Q_q)$ and $\text{succ}(Q_q)$;

- a drawing of a path in the dual of G , connecting the arc $\text{prec}(Q_1)$ and the arc $\text{prec}(Q_q)$, that intersects only the arcs of B^1 between $\text{prec}(Q_1)$ and $\text{prec}(Q_q)$.

Note that any path may enter A_0 only via an arc of B^1 between $\text{prec}(Q_1)$ and $\text{prec}(Q_q)$ and leave A_0 only via an arc of B^2 between $\text{succ}(Q_1)$ and $\text{succ}(Q_q)$. Moreover, A_0 does not contain any terminal, as it contains only elements of H . Therefore the paths Q_1, Q_2, \dots, Q_q (with preceding and succeeding arcs) are the only intersections of the solution $(P_i)_{i=1}^k$ with the closure of the area A_0 .

By Corollary 6.24, the word $i_1 i_2 \dots i_q$ over alphabet $\{1, 2, \dots, k\}$ can be decomposed into powers of at most $2k$ words, such that each of these words contains each symbol at most once. As $q > 2k^2$, there exist a subword w or $i_1 i_2 \dots i_q$ of the form u^2 , where u contains each symbol at most once. Assume $|u| = \ell$ and $w = i_j i_{j+1} \dots i_{j+2\ell-1}$, $i_\iota = i_{\iota+\ell}$ for $j \leq \iota < j + \ell$.

Consider now a subarea A of A_0 , defined similarly as A_0 , but with paths Q_j and $Q_{j+2\ell-1}$ as borders. The intersection of the solution $(P_i)_{i=1}^k$ with the closure of A is exactly the set of paths $Q_j, Q_{j+1}, \dots, Q_{j+2\ell-1}$, together with parts of their preceding and succeeding arcs of B^N and B^R . For $j \leq \iota < j + 2\ell$, let x_ι and y_ι be the first and the last point of Q_ι , respectively. For $j \leq \iota < j + \ell$, the paths Q_ι and $Q_{\iota+\ell}$ belong to P_{i_ι} and i_ι are pairwise distinct for $j \leq \iota < j + \ell$.

Let H_A be a subgraph of H consisting of vertices and edges that are contained entirely in the closure of A . Assume that Q_j appears earlier on the path P_{i_j} than $Q_{j+\ell}$; in the opposite case, we may consider a mirror image of H_A . By Lemma 6.14, for any $j \leq \iota < j + \ell$, the path Q_ι appears on P_{i_ι} earlier than $Q_{\iota+\ell}$.

We now note that we may apply Lemma 6.43 to the graph H_A with paths Q_ι connecting $(x_\iota)_{\iota=j}^{j+2\ell-1}$ with $(y_\iota)_{\iota=j}^{j+2\ell-1}$. Indeed, recall that $d \geq 2k$, thus H contains $2k$ alternating cycles and each of these cycles intersects all paths $(Q_\iota)_{\iota=j}^{j+2\ell-1}$. Therefore k of these cycles (in one of the directions) yield the promised paths $(C_j)_{j=1}^\ell$. We infer that in H_A there exist a sequence of vertex-disjoint paths $(R_\iota)_{\iota=j}^{j+\ell-1}$ such that R_ι starts in x_ι and ends in $y_{\iota+\ell}$.

Consider the following set of paths $(P'_i)_{i=1}^k$. For each $j \leq \iota < j + \ell$, we remove from P_{i_ι} a subpath starting from x_ι and ending at $y_{\iota+\ell}$ and replace it with R_ι . Since the intersection of H_A with the solution $(P_i)_{i=1}^k$ consists of the paths $\{Q_\iota : j \leq \iota < j + 2\ell\}$ only, $(P'_i)_{i=1}^k$ is a solution to k -DPP on G . Moreover, each path P'_i contains strictly less arcs of non-isolating bundles than P_{i_ι} , as we removed from P_{i_ι} an arc $\text{prec}(Q_{\iota+\ell})$, and the paths R_ι contain only isolation arcs of H^{disc} . This contradicts the minimality of $(P_i)_{i=1}^k$ and concludes the proof of theorem. \square

We now show that a minimal solution cannot oscillate between isolation layers.

Theorem 6.45 (no oscillators in the ring). *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance, such that \mathcal{D} has isolation (Λ, d) where $\Lambda \geq 1$ and $d \geq f(k, k) + 4$, where $f(k, t) = 2^{O(kt)}$ is the bound on the type- t bend promised by Lemma 3.7. Let H^{disc} be an isolation component of a ring component H^{ring} , and let f_1 and f_2 be the two faces of $\text{cl}(H^{\text{ring}})[V(H^{\text{disc}})]$ that contain vertices and edges of $G \setminus H^{\text{disc}}$. Let $(P_i)_{i=1}^k$ be a minimal solution to k -DPP on G and assume that, for some $1 \leq \iota \leq k$, the path P_ι contains a subpath P that starts with an arc e_1 with ending point in H^{disc} , ends with an arc e_2 with starting point in H^{disc} , is contained in $\text{cl}(H^{\text{ring}})$ and both e_1 and e_2 lie in f_1 . Then P does not contain any arc that lies in f_2 .*

Proof. Assume the contrary, and the path P contains some arc in f_2 ; in particular, P contains a bundle arc e of bundle B leading from H^{disc} to another disc or ring component of $\text{cl}(H^{\text{ring}})$ that lies in f_2 . Note that B is not an isolating component. Our goal is to apply the bound from Lemma 3.7 to reroute $(P_i)_{i=1}^k$ so that it does not use the arc e . In this way we would strictly decrease

the number of non-isolation bundles of the solution, a contradiction to its minimality. However, to reuse Lemma 3.7, we need to make arguments similar to the proof of Theorem 3.2 using Lemma 3.4.

Let $C_0^*, C_1^*, \dots, C_{d-1}^*$ be a fixed choice of alternating cycles in $\text{cl}(H^{\text{ring}})[V(H^{\text{disc}})]$, separating f_1 from f_2 . Without loss of generality assume that f_1 is the infinite face of $\text{cl}(H^{\text{ring}})[V(H^{\text{disc}})]$, and C_i^* encloses C_j^* whenever $i < j$ (i.e., the cycle C_0^* is close to f_1 and C_{d-1}^* is close to f_2). Let H^* be a subgraph of G consisting of the solution $(P_i)_{i=1}^k$ and the cycles $(C_j^*)_{j=0}^{d-1}$. Let H be a graph obtained from H^* by contracting any arc that belongs both to some path P_i and to some cycle C_j^* . Each path P_i projects to a path Q_i in H , P projects to Q , a subpath of Q_i , and each cycle C_j^* projects to a cycle C_j ; note that C_j may consist of a single arc (self-loop at some vertex), but does not disappear completely. Note that $(Q_i)_{i=1}^k$ is a solution to k -DPP on H and the cycles $(C_j)_{j=0}^{d-1}$ are free.

Let x_0 and y_0 be the two points of $Q \cap C_0$ closest (on Q) to the chosen edge e . As P is contained in $\text{cl}(H^{\text{ring}})$, for one of the two subpaths of C_0 between x_0 and y_0 , denote it C'_0 , we have that the cycle $C'_0 \cup Q[x_0, y_0]$ does not enclose any point of f_2 and, therefore, does not enclose any terminal. By Lemma 3.9, there exists subpaths C'_i of C_i with endpoints x_i and y_i for $1 \leq i < d-1$, such that $(Q[x_0, y_0], C'_0, C'_1, \dots, C'_{d-2})$ is a $(d-2)$ -bend.

Let H^1 be a subgraph of H that consists of the paths $(Q_i)_{i=1}^k$ and the chords $(C'_j)_{j=0}^{d-2}$. Clearly, $(Q_i)_{i=1}^k$ is a solution to k -DPP on H^1 as well, and $(R, C'_0, C'_1, \dots, C'_{d-2})$ is a $(d-2)$ -bend in H^1 . Let $(Q_i^\circ)_{i=1}^k$ be a solution to k -DPP on H^1 that uses minimum possible number of edges that do not lie on the chords $(C'_j)_{j=0}^{d-2}$. We claim the following.

Claim 6.46. *The paths $(Q_i^\circ)_{i=1}^k$ do not use the arc e .*

Proof. Construct a graph H^2 from H^1 as follows: first, for any $0 \leq j \leq d-2$, connect x_j with y_j outside the bend $(R, C'_0, C'_1, \dots, C'_{d-2})$ with an arc a_j in a direction such that $C_j^2 := C'_j \cup \{a_j\}$ is a directed cycle. Note that this can be done such that H^2 is planar and without changing the embedding of H^1 , as the vertices x_j and y_j lie in a good order along the path $Q[x_0, y_0]$: simply draw the arc a_j parallelly to the path $Q[x_0, y_0]$. In this manner we obtain that $(C_j^2)_{j=0}^{d-2}$ is a set of alternating concentric cycles not enclosing any terminal.

Now construct a graph H^3 from H^2 by first removing any arc that does not belong to any cycle C_j^2 nor any path Q_i° and then by contracting any arc that belongs both to some path Q_i° and some cycle C_j^2 . Note that a_j does not become contracted in this manner, and the cycle C_j^2 projects to a cycle C_j^3 in H^3 . Moreover, the solution $(Q_i^\circ)_{i=1}^k$ projects to a solution $(Q_i^3)_{i=1}^k$ to k -DPP in H^3 and the cycles C_j^3 are free with respect to this solution.

Assume that $e \in Q_\eta^\circ$ for some $1 \leq \eta \leq k$; as e does not lie on any cycle C'_j , $e \in Q_\eta^3$ as well. Note that C'_0 separates e from the terminals in H^1 , and thus C_0^3 separates e from the terminals in H^3 . Thus, there exist vertices $x_0^3, y_0^3 \in Q_\eta^3 \cap C_0^3$ that are closest to e on Q_η^3 ; let C_0^{3+} be the subpath of C_0^3 between x_0^3 and y_0^3 that does not contain a_0 . Note that the cycle $C_0^{3+} \cup Q_\eta^3[x_0^3, y_0^3]$ does not enclose any terminal nor any arc a_j , contains a vertex of C_{d-2}^3 . By Lemma 3.9, there is a $(d-4)$ -bend on Q_η^3 that does not enclose any arc a_j , $0 \leq j \leq d-2$.

As $d \geq f(k, k) + 4$, by Lemma 3.7, $(Q_i^3)_{i=1}^k$ is not a minimal solution to k -DPP in $H^3 \setminus \{a_j : 0 \leq j \leq d-2\}$. Let $(R_i^3)_{i=1}^k$ be a different solution. Let $E_R = \bigcup_{i=1}^k E(R_i^3)$, $E_Q = \bigcup_{i=1}^k E(Q_i^3)$ and $E_C = \bigcup_{j=0}^{d-2} E(C_j^3)$. Note that $E(H^3) = E_Q \cup E_C$, and E_Q is a set of disjoint paths. Hence $E_R \setminus E_C$ is a proper subset of E_Q .

Similarly as in the proof of Theorem 3.2, the solution $(R_i^3)_{i=1}^k$ yields a solution $(R_i^\circ)_{i=1}^k$ in H^1 . Indeed, H^3 is created from H^1 by removing some arcs, adding arcs a_j (not used by any path R_i^3)

and contracting some arcs: however, as H^2 consisted of cycles $(C_j^2)_{j=0}^{d-2}$ and paths $(Q_i^\circ)_{i=1}^k$, if we uncontract an arc xy the arc xy is the only arc leaving x and entering y , and any path going through the image of xy in H^3 can be redirected via the arc xy in H^2 . However, the solution $(R_i^\circ)_{i=1}^k$ uses strictly less arcs outside the chords $(C_j^2)_{j=0}^{d-2}$ than $(Q_i^\circ)_{i=1}^k$, a contradiction to the choice of $(Q_i^\circ)_{i=1}^k$. This finishes the proof of the claim. \square

Clearly, $(Q_i^\circ)_{i=1}^k$ is a solution to k -DPP on both H and H^1 that does not use the arc e . Again, by the construction of the graph H from H^* , the solution $(Q_i^\circ)_{i=1}^k$ in H yields a solution $(P_i^\circ)_{i=1}^k$ in H^* that does not use the arc e . Hence, by the construction of H^* , the bundle arcs of non-isolation bundles of the solution (P_i°) is a proper subset of the bundle arcs of non-isolation bundles of the solution $(P_i)_{i=1}^k$, a contradiction to the minimality of $(P_i)_{i=1}^k$. \square

Corollary 6.47 (no ring visitors). *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance, such that \mathcal{D} has isolation (Λ, d) where $\Lambda \geq 1$ and $d \geq f(k, k) + 4$, where $f(k, t) = 2^{O(kt)}$ is the bound on the type- t bend promised by Lemma 3.7. Assume that G is a YES-instance to k -DPP and let $(P_i)_{i=1}^k$ be a minimal solution. Then there are no ring visitors on any path P_i .*

Proof. We apply Theorem 6.45 for $H^{\text{disc}} = H_{\Gamma, \Lambda}^{\text{disc}}$ for $\Gamma \in \{IN, OUT\}$ and $H_{\Gamma, \Lambda}^{\text{disc}}$ be an isolating component of a ring component H^{ring} . \square

Corollary 6.48 (bundle words with ring holes for paths in a minimal solution). *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance, such that \mathcal{D} has isolation (Λ, d) where $\Lambda \geq 2$ and $d \geq f(k, k) + 4$, where $f(k, t) = 2^{O(kt)}$ is the bound on the type- t bend promised by Lemma 3.7. Assume that G is a YES-instance to k -DPP and let $(P_i)_{i=1}^k$ be a minimal solution. Then, for any $0 \leq \lambda < \Lambda - 1$ and $1 \leq i \leq k$, the unique bundle word with level- λ ring holes that is consistent with P_i is flat, that is, the its level- $(\lambda + 1)$ projection does not contain any arcs of bundles of level λ or lower.*

Proof. As a bundle word with level- $(\lambda + 1)$ ring holes does not contain any symbol corresponding to an arc that belongs to a level- $(\lambda + 1)$ ring passage, any level- λ ring bundle in such a bundle word with ring holes needs to correspond to an arc on a structure forbidden by Theorem 6.45. \square

7 From bundle words to disjoint paths

In this section we prove the following theorem that will be the main tool both for measuring the good guesses for winding numbers, as well as for seeking the ultimate solution.

Theorem 7.1. *For any bundled instance $(G, \mathcal{D}, \mathbb{B})$ of ring isolation (Λ, d) , $\Lambda, d > 0$, any $1 \leq \lambda < \Lambda$ and any sequence $(\mathbf{p}_i)_{i=1}^k$ of bundle words with level- λ ring holes that do not contain any level-0 bundles, one may in polynomial time either:*

1. *correctly conclude that there is no solution $(P_i)_{i=1}^k$ to k -DPP on G such that P_i is consistent with \mathbf{p}_i for each $1 \leq i \leq k$; or*
2. *compute a solution $(P_i)_{i=1}^k$ to k -DPP on G with the following property: for any $1 \leq i \leq k$ and any bundle B that is a normal bundle or has at least one endpoint in an isolation bundle of level higher than λ , the number of appearances of B in $\text{bw}(P_i)$ is not greater than the number of appearances of B in \mathbf{p}_i .*

The rest of this Section is devoted to the proof of Theorem 7.1. Intuitively, the argument is based on designing an instance of cohomology problem in a natural manner, and then running the algorithm of Theorem 4.16. However, formal construction of the instance and proof that the output

of the algorithm of Theorem 4.16 is meaningful requires technical effort. In particular, we need to check that bundle words with level- λ ring holes $(\mathbf{p}_i)_{i=1}^k$ contain already enough information to deduce the topology of path network in the graph.

Apart from the decomposition \mathcal{D} , we will be working as well on a modified decomposition \mathcal{D}' where for every ring component H^{ring} , all the isolation components of H^{ring} up to level λ are included into H^{ring} . In other words, for every ring component H^{ring} , we identify component H^{ring} together with first λ layers of IN- and OUT- isolation into one ring component $\text{cl}^\lambda(H^{\text{ring}})$. We also define \mathbb{B}' to be the set of these bundles $B \in \mathbb{B}$ that are normal bundles or have at least one endpoint in a isolation component of level higher than λ .

In decomposition \mathcal{D}' , for every component H choose an arbitrary vertex $v(H) \in V(H)$, and for every bundle $B \in \mathbb{B}'$ choose an arbitrary arc $a(B) \in B$. Moreover, if B connects components H_1 and H_2 , let us choose arbitrary undirected paths $P(B, H_1)$ and $P(B, H_2)$ in H_1, H_2 that connect $v(H_1), v(H_2)$ with $a(B)$, respectively. Note that this is possible as every component is weakly connected. Moreover, we can choose the paths so that they do not cross each other in the planar embedding (though may use the same arcs). In addition, for every ring component $H^{\text{ring}} \in \mathcal{D}'$ choose an arbitrary undirected cycle $C(H^{\text{ring}})$ of homotopy 1 that passes through $v(H^{\text{ring}})$, again in such a manner that this cycle does not cross any path $P(B, H^{\text{ring}})$ for any bundle B adjacent to H^{ring} .

Now, given a sequence $(\mathbf{p}_i)_{i=1}^k$ of bundle words with level- λ ring holes in decomposition \mathcal{D} , we construct undirected paths $(Q_i)_{i=1}^k$ connecting respective terminals. Intuitively, paths Q_i are going to model seeked paths P_i in the solution. Consider one bundle word with level- λ holes $\mathbf{p}_i = ((p_{i,j})_{j=0}^{q(i)}, (w_{i,j})_{j=1}^{q(i)})$. We say that a bundle is *deep* if both its endpoints are either in ring components or in isolation components of level lower or equal to λ , i.e., it does not appear in \mathbb{B}' ; otherwise the bundle is *shallow*. Note that by our assumption on $(\mathbf{p}_i)_{i=1}^k$, every deep bundle of each \mathbf{p}_i has both endpoints in an isolation component of level lower or equal to λ . In each bundle word $p_{i,j}$, distinguish maximal subwords consisting of deep bundles only. Then (note that in the following we refer to components in decomposition \mathcal{D}):

- For every two consecutive shallow bundles B_1, B_2 on $p_{i,j}$, such that B_1 is directed toward a disc component H^{disc} while B_2 is directed from H^{disc} , take concatenation of paths $P(B_1, H^{\text{disc}})$ and $P(B_2, H^{\text{disc}})$ (note here that $H^{\text{disc}} \in \mathcal{D}'$).
- Consider now any two shallow bundles B_1, B_2 on $p_{i,j}$ such that the subword v between B_1, B_2 consists of deep bundles only. Then bundles B_1, B_2 are on the same side of H_*^{ring} , and the subword v between them can visit only isolation components of H^{ring} of levels smaller or equal than λ on the same side as B_1, B_2 . Examine word v and compute the signed number w of how many times bundles from v cross the reference curve of $\text{cl}^\lambda(H^{\text{ring}})$. Observe that, by the construction of the reference curve of $\text{cl}^\lambda(H^{\text{ring}})$, any path consistent with bundle word $B_1 v B_2$ will have winding number w in $\text{cl}^\lambda(H^{\text{ring}})$. Note that any such path leads from one face of $\text{cl}^\lambda(H^{\text{ring}})$ to the same face of $\text{cl}^\lambda(H^{\text{ring}})$, so its winding number with respect to the reference curve of $\text{cl}^\lambda(H^{\text{ring}})$ must belong to the set $\{-1, 0, 1\}$. Hence, if $w \notin \{-1, 0, 1\}$, then we can conclude that there cannot be any solution consistent with bundle words with level- λ ring holes $(\mathbf{p}_i)_{i=1}^k$, and we terminate the whole algorithm returning the negative answer. Otherwise, we take the concatenation of paths $P(B_1, \text{cl}^\lambda(H^{\text{ring}}))$, $C(\text{cl}^\lambda(H^{\text{ring}}))^{w'}$ and $P(B_2, \text{cl}^\lambda(H^{\text{ring}}))$, where w' is chosen such that the resulting path has winding number w in $\text{cl}^\lambda(H^{\text{ring}})$. Note here that of course B_1 and B_2 are adjacent to $\text{cl}^\lambda(H^{\text{ring}})$ in \mathcal{D}' .
- For every two bundles B_1, B_2 such that B_1 is the last bundle of $p_{i,j}$ and B_2 is the first bundle of $p_{i,j+1}$, perform the following construction. Let H^{ring} be the ring component such that

B_1, B_2 are on different sides of $\text{cl}^\lambda(H^{\text{ring}})$ and let $w := w_{i,j+1}$ be the winding number between them. Take concatenation of $P(B_1, \text{cl}^\lambda(H^{\text{ring}}))$, $C(\text{cl}^\lambda(H^{\text{ring}}))^{w'}$ and $P(B_2, \text{cl}^\lambda(H^{\text{ring}}))$, where w' is chosen such that the resulting path has winding number w .

Finally, obtain Q_i by concatenating constructed undirected paths in a natural order imposed by $(\mathfrak{p}_i)_{i=1}^k$, and gluing them in between by arcs $a(B)$ for $B \in \mathfrak{p}_i$. If some arc $a(B)$ is traversed in the constructed path in a wrong direction, we may immediately terminate the computation providing a negative answer: the corresponding path of the solution would also need to traverse the same bundle in the wrong direction, which is impossible.

We have already constructed models Q_i for paths P_i , but to finish the construction of the instance of the cohomology problem, we need to choose consistently the order of paths on shared arcs.

Consider two paths Q_i, Q_j (possibly $i = j$) and let M_i, M_j be some (possibly consisting of a single vertex) maximal fragments of Q_i, Q_j that are common, i.e., M_i and M_j are the same path, but the arcs p_i, p_j preceding M_i and M_j on Q_i, Q_j , respectively, are different, as well as arcs s_i, s_j succeeding M_i and M_j on Q_i, Q_j are different. Contract fragments M_i, M_j to one vertex v and consider the positioning of arcs p_i, p_j, s_i, s_j around v . Observe that if there is a solution $(P_i)_{i=1}^k$ consistent with $(\mathfrak{p}_i)_{i=1}^k$, then p_i and s_i cannot separate p_j from s_j : otherwise, the parts of paths P_i, P_j between the last bundles preceding M_i, M_j and the first bundles succeeding P_i, P_j would need to cross. Therefore, if we think of paths Q_i, Q_j as curves on the plane, it is possible to spread them slightly along the common path M_i, M_j so that the curves do not intersect along this path. The same reasoning can be performed when paths Q_i, Q_j traverse the common path in different directions, i.e., M_i, M_j are two maximal subpaths such that M_i is M_j reversed, the arc preceding M_i on Q_i is different than the arc succeeding M_j on Q_j , and the arc preceding M_j on Q_j is different than the arc succeeding M_i on Q_i .

We perform this reasoning for every pair of indices i, j , every two maximal common fragments, and both directions of traversal; if for any pair of fragments we obtain inconsistency, we terminate the algorithm providing an answer that no solution can be found. As a result, for every arc that is traversed by some path at least two times, we obtain a constraint between every two traversals, which traversal should be spread towards the face on one side of the arc, and which should be spread to the other side. It can be easily seen that this constraints are transitive, i.e., if traversal a must be left to traversal b and b must be to the left to traversal c , then traversal a must be to the left to traversal c ; if this is not the case, we can again terminate the algorithm providing a negative answer. Hence, we may order the traversals of every arc in a linear order, so that slightly spreading all the paths $(Q_i)_{i=1}^k$ according to these orders on arcs yields a family of pairwise non-intersecting curves, denoted $(\tilde{Q})_{i=1}^k$.

We now consider a free group Λ on k generators g_1, g_2, \dots, g_k , corresponding to paths P_1, P_2, \dots, P_k . Define a function $\phi : E(G) \rightarrow \Lambda$ by putting for every arc a the product of generators or their inverses corresponding to traversals of this arc: if P_i traverses a respecting direction of a then we take g_i , if disrespecting then we put g_i^{-1} , and we multiply these group elements according to the linear order of traversals on this arc, found as in the previous paragraph. As the family $(\tilde{Q})_{i=1}^k$ is pairwise non-intersecting, it follows that for every non-terminal vertex v , if we compute the product of group elements on the arcs incident to v in the order imposed by the planar embedding, where the value of every arc is taken as its inverse if the arc is incoming to v and as the value itself otherwise, then we obtain 1_Λ .

Let G^* be the dual of G and let G^+ be the extended dual of G . Clearly, as $E(G) = E(G^*)$, we may consider ϕ also as a function from arcs of G^* to Λ . Moreover, we may naturally extend ϕ from $E(G^*)$ to ϕ^+ defined on $E(G^+)$, as in the proof of Theorem 4.19.

We are now going to define an instance of cohomology feasibility problem similarly as in the proof of Theorem 4.19. For every arc $a \in E(G^+)$, let $H(a) = \{1, g_1, g_2, \dots, g_k\}$ if $a \in E(G)$, and let $H(a) = \{1, g_1, g_2, \dots, g_k, g_1^{-1}, g_2^{-1}, \dots, g_k^{-1}\}$ if $a \in E(G^+) \setminus E(G)$, i.e., a was added in the extension. Moreover, let S be the set of end-faces in decomposition \mathcal{D}' ; we would like to stress that we consider decomposition \mathcal{D}' , hence we do not put any restrictions on faces that are contained in level- λ closures of ring components. Thus we have defined an instance of cohomology feasibility problem $(G^+, \Lambda, \phi, H, S)$, to which we can apply the algorithm of Theorem 4.16. The rest of the proof of Theorem 7.1 follows from the two lemmata that show equivalence of the cohomology instance and the statement of Theorem 7.1.

Lemma 7.2. *If there is a sequence of disjoint paths $(P_i)_{i=1}^k$ consistent with $(\mathfrak{p}_i)_{i=1}^k$, then instance $(G^+, \Lambda, \phi, H, S)$ has some solution ψ .*

Proof. Assume that we have a sequence of disjoint paths $(P_i)_{i=1}^k$ consistent with $(\mathfrak{p}_i)_{i=1}^k$. Similarly to paths $(Q_i)_{i=1}^k$ and function ϕ , sequence $(P_i)_{i=1}^k$ also naturally defines a function $\psi : E(G^+) \rightarrow \Lambda$, as in the proof of Theorem 4.19. Moreover, ψ satisfies all the constraints imposed by H . It remains to prove that ψ is cohomologous to ϕ via a function F that assigns 1_Λ to all the end-faces in decomposition \mathcal{D}' .

For every path P_i , decompose it into sequence $(b_{i,0}, P_{i,1}, b_{i,1}, P_{i,2}, \dots, b_{i,t}, P_{i,t}, b_{i,t+1})$, where $P_{i,j}$ are maximal subpaths contained in decomposition \mathcal{D}' and $b_{i,j}$ are bundle arcs from bundles of decomposition \mathcal{D}' that appear on \mathfrak{p}_i (note that not every bundle appearing in \mathfrak{p}_i is represented by some arc $b_{i,j}$, as we omit the bundles that are not in \mathbb{B}'). As both P_i and Q_i are consistent with \mathfrak{p}_i and \mathfrak{p}_i is a level- λ bundle word with holes, path Q_i can be similarly decomposed into sequence $(b'_{i,0}, Q_{i,1}, b'_{i,1}, Q_{i,2}, \dots, b'_{i,t}, Q_{i,t}, b'_{i,t+1})$, where arcs $b_{i,j}$ and $b'_{i,j}$ belong to the same bundle for every j , while $P_{i,j}$ and $Q_{i,j}$ traverse the same component from the same bundle to the same bundle. Moreover, if $P_{i,j}$ and $Q_{i,j}$ are connecting sides of the same ring component H^{ring} (regardless whether the same or different), then $P_{i,j}$ and $Q_{i,j}$ have the same winding number in H^{ring} .

It follows that for every component H , all the traversals of paths $(Q_i)_{i=1}^k$ via H can be simultaneously shifted to $(P_i)_{i=1}^k$ using a by a continuous shift of the interior of H (formally, by a continuous function $f : U(H) \times [0, 1] \rightarrow U(H)$ such that $f(\cdot, 0)$ is identity and in $f(\cdot, 1)$ paths Q_i are mapped to respective paths P_i). Note that we may choose this shift so that end-faces of bundles adjacent to H stay invariant, as ends of paths in a particular bundle must be mapped to ends of paths in the same bundle. Now observe that this shift can be modelled by a function F defined at vertices of G^+ , which measures which paths are shifted over a given face and in which order. If we define this function F for each component separately, then F will be defined consistently in the same manner on the mortar faces of \mathcal{D}' . Moreover, invariance of the end-faces in the homeomorphisms means that end-faces will be assigned 1_Λ . Hence ϕ and ψ are cohomologous via a cohomology that fixes end-faces. \square

Lemma 7.3. *If the instance $(G^+, \Lambda, \phi, H, S)$ has some solution ψ , then there is a solution $(P_i)_{i=1}^k$ to k -DPP on G with property (2) defined in the statement of Theorem 7.1.*

Proof. Assume that $\psi : E(G^+) \rightarrow \Lambda$ is a function that is consistent with H and is cohomologous to ψ via a function F that assigns 1_Λ to all the end-faces of \mathcal{D}' . For $i = 1, 2, \dots, k$, let G_i be the subgraph of G consisting of all the arcs a of G such that $\psi(a^*) = g_i$. Note that conditions imposed in the cohomology instance imply that:

- subgraphs G_i are vertex-disjoint;

- for every $i = 1, 2, \dots, k$, all the vertices of G_i have the same in- and outdegrees, apart from s_i and t_i that have exactly one outgoing arc in G_i and exactly one incoming arc in G_i , respectively.

By the standard degree counting argument, we infer that for each $i = 1, 2, \dots, k$ vertices s_i and t_i must be in the same weakly connected component of G_i . Moreover, this weakly connected component must contain an eulerian tour from s_i to t_i . By shortcutting this eulerian tour appropriately, we obtain a simple path P_i in G_i leading from s_i to t_i . As subgraphs G_i are vertex-disjoint, so do paths P_i . It remains to prove that paths P_i satisfy the property that for each shallow bundle B , the number of occurrences of B in $\mathbf{bw}(P_i)$ is not larger than the number of occurrences of B on \mathfrak{p}_i . To this end, we will prove that the subgraph G_i in total contains exactly the same number of arcs of B as the number of occurrences of B on \mathfrak{p}_i ; as P_i is a subgraph of G_i , the lemma statement follows.

Consider any component H of \mathcal{D}' and any bundle $B \in \mathbb{B}'$ adjacent to H . Let P_B^* be the directed path or directed cycle in G^+ traversing the bundle B , whose existence is guaranteed by Lemma 6.3. Let f_0, f_1, \dots, f_r are the consecutive faces of G visited by P_B^* and a_1, a_2, \dots, a_r are consecutive arcs. Suppose for a moment that f_0 and f_r are end-faces. We claim that then the products of elements assigned by ψ and by ϕ to the arcs $a_1^*, a_2^*, \dots, a_r^*$ are equal. Indeed, we have that

$$\prod_{i=1}^r \psi(a_i^*) = \prod_{i=1}^r F^{-1}(f_{i-1})\phi(a_i^*)F(f_i) = F^{-1}(f_0) \left(\prod_{i=1}^r \phi(a_i^*) \right) F(f_r) = \prod_{i=1}^r \phi(a_i^*),$$

where the last equality follows from the fact that F assigns 1_Λ to all the end-faces. If now f_0 and f_r are not end-faces, then $f_0 = f_r$, B is the only bundle adjacent to H , and $\prod_{i=1}^r \psi(a_i^*)$ and $\prod_{i=1}^r \phi(a_i^*)$ are conjugate using $F(f_0) = F(f_r)$.

For $\iota = 1, 2, \dots, k$, consider a homomorphism $h_\iota : \Lambda \rightarrow \mathbb{Z}$ defined on generators by $h_\iota(g_\iota) = 1$ and $h_\iota(g_{\iota'}) = 0$ for $\iota' \neq \iota$; this homomorphism just counts the number of g_ι -s. It follows that in both cases $h_\iota(\prod_{i=1}^r \psi(a_i^*)) = h_\iota(\prod_{i=1}^r \phi(a_i^*))$: either simply $\prod_{i=1}^r \psi(a_i^*) = \prod_{i=1}^r \phi(a_i^*)$, or

$$\begin{aligned} h_\iota \left(\prod_{i=1}^r \psi(a_i^*) \right) &= h_\iota \left(F^{-1}(f_0) \left(\prod_{i=1}^r \phi(a_i^*) \right) F(f_r) \right) \\ &= h_\iota \left(\prod_{i=1}^r \phi(a_i^*) \right) + h_\iota(F(f_r)) - h_\iota(F(f_0)) = h_\iota \left(\prod_{i=1}^r \phi(a_i^*) \right). \end{aligned}$$

Note that all the elements $\psi(a_i^*)$ are simply generators, as ψ respects constraints imposed by H , and $\phi(a_i)$ are only multiplications of positive powers of generators, as we excluded the cases where paths Q_i traverse bundles in the wrong direction. Hence, $h_\iota(\prod_{i=1}^r \psi(a_i^*))$ is the number of arcs of G_ι contained in the bundle B , and $h_\iota(\prod_{i=1}^r \phi(a_i^*))$ is the number of passages of path Q_ι via arc $a(B)$, which, by the construction of Q_ι , is equal to the number of occurrences of B on \mathfrak{p}_ι . \square

8 Guessing bundle words and their winding numbers

The aim of this section is to prove that, for a given bundled instance $(G, \mathcal{D}, \mathbb{B})$ with sufficient isolation, the number of reasonable bundle words with holes for a minimal solution on G is bounded by a function of $|\mathbb{B}|$, $|\mathcal{D}|$ and k . Note that, as each terminal lies in its own component of \mathcal{D} and G is weakly connected, we may assume $2k \leq |\mathcal{D}| \leq |\mathbb{B}| + 1$.

Formally, we prove the following theorem.

Theorem 8.1. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance of isolation (Λ, d) , where $\Lambda \geq 3$, $d \geq \max(2k, f(k, k) + 4)$, and $f(k, t) = 2^{O(kt)}$ is the bound on the type- t bend promised by Lemma 3.7. Then in $2^{O(k^2|\mathbb{B}|^2 \log |\mathbb{B}|)} |G|^{O(1)}$ time one can compute a family of at most $2^{O(k^2|\mathbb{B}|^2 \log |\mathbb{B}|)}$ sequences $(\mathbf{p}_i)_{i=1}^k$ of bundle words with level-2 ring holes, not containing any level-0 bundle, such that if $(G, \mathcal{D}, \mathbb{B})$ is a YES-instance to k -DPP, then there exists a solution $(P_i)_{i=1}^k$ to k -DPP on $(G, \mathcal{D}, \mathbb{B})$ and a sequence $(\mathbf{p}_i)_{i=1}^k$ in the generated set such that P_i is consistent with \mathbf{p}_i for each $1 \leq i \leq k$.*

The proof of Theorem 8.1 consists of two steps. In Section 8.2 we show that there is only a bounded number of reasonable word parts of the promised bundle words with ring holes. Then, in Section 8.3, we show that we may pick winding numbers from a bounded set of candidates.

8.1 Zooms

In this section we introduce a toolbox for measuring spiraling properties of parts of the graph G . Informally speaking, Lemma 6.42 implies that in a minimal solution to k -DPP, the length of any spiral is (close to) minimum possible in G — otherwise, we should be able to reroute the solution to a shorter spiral, contradicting the minimality of the solution. We now introduce some gadgets and auxiliary graphs that allow us to measure this minimum possible size of spirals in G .

Definition 8.2 (zoom). Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance with isolation (Λ, d) . A *zoom* is a pair $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ such that $\hat{\mathbb{B}}$ is a subset of \mathbb{B} that does not contain any bundle with an arc incident to a terminal, $\hat{\mathcal{D}} \subseteq \mathcal{D}$ and $H \in \hat{\mathcal{D}}$ if and only if there exists $B \in \hat{\mathbb{B}}$ whose arcs start or end in H .

We say that a zoom $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ is *level- λ safe* for some $0 \leq \lambda \leq \Lambda$, if for any ring component $H^{\text{ring}} \in \mathcal{D}$, either the set $\hat{\mathbb{B}}$ contains all bundles with arcs with both endpoints in $\text{cl}^\lambda(H^{\text{ring}})$ or $\hat{\mathcal{D}}$ is disjoint with the set of components of $\text{cl}^\lambda(H^{\text{ring}})$.

Definition 8.3 (zoom pass, pack of zoom passes). Let $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ be a zoom in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ with isolation (Λ, d) that is λ -safe for some $0 \leq \lambda < \Lambda$. A *level- λ zoom pass* in $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ is a bundle word with level- λ ring holes $\mathbf{p} = ((p_j)_{j=0}^q, (w_j)_{j=1}^q)$ in $(G, \mathcal{D}, \mathbb{B})$ such that:

1. $q \geq 1$ or $|p_0| \geq 2$, i.e., the bundle words of \mathbf{p} contain at least two symbols in total;
2. the first bundle of p_0 (called *the first bundle of \mathbf{p}*) does not belong to $\hat{\mathbb{B}}$, but its arcs end in a component belonging to $\hat{\mathcal{D}}$;
3. the last bundle of p_q (called *the last bundle of \mathbf{p}*) does not belong to $\hat{\mathbb{B}}$, but its arcs start in a component belonging to $\hat{\mathcal{D}}$;
4. all other bundles of p_j , $0 \leq j \leq q$, belong to $\hat{\mathbb{B}}$.

. A *pack of level- λ zoom passes* in $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ is a pair $((\mathbf{p}_\tau)_{\tau \in I}, (\psi_{B,\alpha})_{B \in \mathbb{B}, 1 \leq \alpha \leq 2})$ where

1. $(\mathbf{p}_\tau)_{\tau \in I}$ is a sequence of level- λ zoom passes for some index set I ;
2. $\psi_{B,1}$ is a permutation of those indices $\tau \in I$ for which \mathbf{p}_τ starts with B , and $\psi_{B,2}$ is a permutation of those indices $\tau \in I$ for which \mathbf{p}_τ ends with B .

Informally speaking, a zoom is a part of a graph G which we investigate, and zoom passes are bundle words with ring holes of chosen parts of a solution that pass through a zoom. The permutations $\psi_{B,1}$ and $\psi_{B,2}$ are intended to correspond to the order of the chosen parts on bundle B when entering and leaving the zoom, respectively. Note that we do not require $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ to be an induced subgraph of $(\mathcal{D}, \mathbb{B})$ and, consequently, we allow a zoom pass p to start or end with a bundle with arcs with both endpoints in components of $\hat{\mathcal{D}}$ (but this bundle cannot belong to $\hat{\mathbb{B}}$).

We now note that the number of choices for $\psi_{B,\alpha}$ is bounded.




Figure 39: An illustration of the zoom pass starting gadget, attached to bundle B with endpoints in component D . The gray squares represent the new terminals, and the dashed lines represent the arcs of the bundle B .

Lemma 8.4. *Let $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ be a zoom in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ with isolation (Λ, d) and let $(\mathbf{p}_\tau)_{\tau \in I}$ be a sequence of level- λ zoom passes in $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ for some $0 \leq \lambda < \Lambda$. Then there exist at most $(|I|!)^2$ sequences $(\psi_{B,\alpha})_{B \in \mathbb{B}, 1 \leq \alpha \leq 2}$ such that $((\mathbf{p}_\tau)_{\tau \in I}, (\psi_{B,\alpha})_{B \in \mathbb{B}, 1 \leq \alpha \leq 2})$ is a pack of zoom passes in $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$.*

Proof. For $B \in \mathbb{B}$, let $I_{B,1}$ be a set of indices $\tau \in I$ such that \mathbf{p}_τ starts with B , and $I_{B,2}$ be a set of indices $\tau \in I$ such that \mathbf{p}_τ ends with B . By definition, $\psi_{B,\alpha}$ is a permutation of $I_{B,\alpha}$ for any $B \in \mathbb{B}$, $1 \leq \alpha \leq 2$. The lemma follows from the convexity of the factorial function and the fact that

$$\sum_{B \in \mathbb{B}} |I_{B,1}| = \sum_{B \in \mathbb{B}} |I_{B,2}| = |I|.$$

□

We now define how a zoom with a pack of zoom passes defines a sub-instance of the original k -DPP bundled instance $(G, \mathcal{D}, \mathbb{B})$.

Definition 8.5 (Zoom pass starting and ending gadget). Consider a bundled instance $(G, \mathcal{D}, \mathbb{B})$ and a bundle $B \in \mathbb{B}$. Let I be an index set and let ψ be a permutation of I .

The *zoom pass starting gadget* is constructed as follows. First, take a bi-directional grid of size $|I| \times |B|$, that is, take a set of vertices $\{v_{\tau,b} : \tau \in I, b \in B\}$ and connect $v_{\tau,b}$ and $v_{\tau,b'}$ with arcs in both directions for $\tau \in I$ and b, b' being two consecutive arcs of B , as well as $v_{\tau,b}$ and $v_{\tau',b}$ for $b \in B$ and τ, τ' being two consecutive indices in the permutation ψ of I . Let $\tau_1 \in I$ be the first element of I in the permutation ψ and b_1 be the first arc of B . For each $b \in B$, make an arc from $v_{\tau_1,b}$ to the endpoint of b in G . For each $\tau \in I$, create a new source terminal s_τ and connect it with v_{τ,b_1} . See Figure 39 for an illustration.

The *zoom pass ending gadget* is constructed similarly as the zoom pass starting gadget, except for three differences: the arcs are reversed, the constructed terminals are sink terminals, and the gadget is attached to the starting points of the arcs of B .

Definition 8.6 (zoom auxiliary graph). Let $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ be a zoom in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ and let $((\mathbf{p}_\tau)_{\tau \in I}, (\psi_{B,\alpha})_{B \in \mathbb{B}, 1 \leq \alpha \leq 2})$ be a pack of zoom passes in $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$.

The *zoom auxiliary graph* is constructed as follows. First, we take a subgraph of G consisting of all components of $\hat{\mathcal{D}}$ and all arcs of bundles of $\hat{\mathcal{D}}$. Then, for each bundle $B \in \mathbb{B}$ that is the first symbol of some bundle word with ring holes \mathbf{p}_τ , we define $I_{B,1}$ to be the set of those indices $\tau \in I$ for which \mathbf{p}_τ starts with B , and we construct a zoom pass starting gadget for bundle B , index set $I_{B,1}$ and permutation $\psi_{B,1}$. Similarly, for each bundle $B \in \mathbb{B}$ that is the last symbol of some \mathbf{p}_τ , we define $I_{B,2}$ to be the set of those indices $\tau \in I$ for which \mathbf{p}_τ ends with B , and we construct a zoom pass ending gadget for bundle B , index set $I_{B,2}$ and permutation $\psi_{B,2}$.

We note that a zoom auxiliary graph is always planar: any bundle B for which we construct a starting or ending gadget does not belong to $\hat{\mathbb{B}}$, and we can embed these gadgets in the plane in the space occupied by B in the embedding of the graph G . Note that this claim remains true also if a bundle B appears both as the first bundle of some zoom pass and the last bundle of some (possibly other) zoom pass.

Moreover, note that the aforementioned embedding imposes a natural decomposition of the zoom auxiliary graph. Formally, as a set of components we take $\hat{\mathcal{D}}$, the set of bundles $\hat{\mathbb{B}}$ and we add the following components and bundles:

1. for each zoom pass starting gadget for a bundle B , we take a disc component $H_{B,1}^{\text{disc}}$ that contains the bidirectional grid, and a bundle $B_{B,1}$ that contains arcs connecting $H_{B,1}^{\text{disc}}$ with the endpoints of the bundle B ;
2. similarly, for each zoom pass ending gadget we define disc component $H_{B,2}^{\text{disc}}$ and a bundle $B_{B,2}$;
3. each introduced terminal is embedded in its own disc component, and its incident arc is embedded in its own bundle.

In this manner we define at most $4|I|$ new components and bundles, and we create a bundled instance with $|I|$ terminal pairs indexed by I . Let us call this instance a *zoom auxiliary instance*.

Let $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ be a level- λ safe zoom in a bundled instance $(G, \mathcal{D}, \mathbb{B})$ with isolation (Λ, d) for some $0 \leq \lambda < \Lambda$. Let $(P_i)_{i=1}^k$ be a solution to k -DPP on G and assume $\hat{\mathcal{D}}$ does not contain any component with a terminal. For one path P_i , we say that its subpath P_τ is a *zoom incident* (with respect to the zoom $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$) if P_τ :

1. starts with an arc $b_{\tau,1}$ that belongs to a bundle $B_{\tau,1} \notin \hat{\mathbb{B}}$, but the ending point of $b_{\tau,1}$ belongs to a component of $\hat{\mathcal{D}}$;
2. ends with an arc $b_{\tau,2}$ that belongs to a bundle $B_{\tau,2} \notin \hat{\mathbb{B}}$, but the starting point of $b_{\tau,2}$ belongs to a component of $\hat{\mathcal{D}}$;
3. all other arcs of P_τ either belong to a bundle of $\hat{\mathbb{B}}$ or a component of $\hat{\mathcal{D}}$.

Note that all zoom incidents of a path P_i are pairwise disjoint, except for possibly overlapping the first and last arcs; let $P_{i,1}, P_{i,2}, \dots, P_{i,q(i)}$ be the sequence of all zoom incidents on P_i (with respect to $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$) in the order of they appearance on P_i . Let $I = \{(i, j) : 1 \leq i \leq k, 1 \leq j \leq q(i)\}$.

For $\tau = (i, j) \in I$, let \mathbf{p}_τ be a bundle word with level- λ ring holes constructed as follows: we take $\text{bw}(P_\tau)$ and for each level- λ ring passage of P_τ through $\text{cl}^\lambda(H^{\text{ring}})$ for some $H^{\text{ring}} \in \hat{\mathcal{D}}$, we replace the subword corresponding to this ring passage with its winding number. The fact that \mathbf{p}_τ

satisfies all properties of a bundle word with level- λ ring holes follows from the assumption that the zoom $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ is level- λ safe.

For any $B \in \mathbb{B}$ and $1 \leq \alpha \leq 2$, let $I_{B,\alpha} = \{\tau \in I : b_{\tau,\alpha} \in B\}$ and $\psi_{B,\alpha}$ be the order of the indices $\tau \in I_{B,\alpha}$ in which arcs $b_{\tau,\alpha}$ appear on B . Observe the following.

Observation 8.7. *The pair $((\mathfrak{p}_\tau)_{\tau \in I}, (\psi_{B,\alpha})_{B \in \mathbb{B}, 1 \leq \alpha \leq 2})$ defined above is a pack of zoom passes in $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$. Moreover, the zoom auxiliary instance constructed for this pack of zoom passes has a solution $(P_\tau^*)_{\tau \in I}$, where P_τ^* is constructed from P_τ by modifying its first and last arcs: we replace $b_{\tau,1}$ with a path from s_τ through vertices $v_{\tau,b}$ up to index $b_{\tau,1}$ and then through vertices $v_{\tau,b_{\tau,1}}$ up to $v_{\tau_1,b_{\tau,1}}$ and to the ending point of $b_{\tau,1}$; a similar replacement is made for $b_{\tau,2}$. Moreover, the bundle word with level- λ ring holes for P_τ^* in the created instance is equal to \mathfrak{p}_τ^* , except for the prefix and suffix that correspond to the passage in the zoom pass starting and ending gadgets.*

8.2 Deducing bundle words

In this section we present the first half of the algorithm of Theorem 8.1: we aim to guess the bundle word part of the required bundle words with holes; the winding number part is guessed in the next subsection.

Consider a minimal solution $(P_i)_{i=1}^k$ to the bundled instance $(G, \mathcal{D}, \mathbb{B})$. By Corollary 6.23, the bundle word $\mathbf{bw}(P_i)$ can be decomposed into $u_{i,1}^{r_{i,1}} u_{i,2}^{r_{i,2}} \dots u_{i,s(i)}^{r_{i,s(i)}}$ for some $s(i) \leq 2|\mathbb{B}|$. Note that there is only a bounded number of choices for the words $u_{i,j}$; the difficult part is to guess the exponents $r_{i,j}$ for those j where $r_{i,j}$ is large. In this section we guess only exponents $r_{i,j}$ for which $u_{i,j}$ contains only bundles with both endpoints in disc components (i.e., does not contain any level-0 ring bundles); in the next section we ‘hide’ the unknown exponents by going to the level-1 projection of the bundle words and using results of Section 6.7.

However, if some path P_i has subpath with bundle word u^r for some large integer r and u not containing any level-0 ring bundles, we have a spiraling ring A and we may use Lemma 6.42: we are not able to route the same paths through the spiraling ring using less than $r - 10$ turns. Lemma 6.10 gives us an answer how many times the paths $(P_i)_{i=1}^k$ cross the spiraling ring A and in which direction. Our approach is to create a zoom containing the spiral u , attach to it all paths that want to cross the spiraling ring A and measure the minimal number of turns they need to make along the spiral u .

However, there is a significant difficulty with this approach — it is unclear where to attach the paths crossing A to the zoomed spiral u . To create a zoom instance, we need that each zoom pass starts and ends in a bundle that does not appear in the zoom (in our case, does not appear on the spiral u). However, any path P_i , before spiraling on the spiral u , may go along a bundle word v^q , where v contains a proper subset of the symbols of u . Thus, we need to include such spirals as well in our zoom instance. Luckily, we can assume that the exponent q is already known to us, if we assume that we guess exponents $r_{i,j}$ in the order of increasing lengths $|u_{i,j}|$.

The approach from the previous paragraph, although resolves the issue of where to attach the zoom passes in the created zoom instance, creates a new problem. If we apply Theorem 7.1 to the created zoom instance, the returned paths follow our guessed bundle words in a very relaxed manner. As we have included in our zoom instance all parts of the solution that cross the spiraling ring A , by Lemma 6.42 in the returned solution the path P_i needs to spiral at least $r - 10$ times in this ring. However, it is no longer true for other terms v^q where v contains a proper subset of the symbols of u : the solution returned by Theorem 7.1 may ‘borrow’ some bundles from such terms, in order to spiral along u significantly *more* times than r .

To circumvent this problem, we add even more passes to our zoom instance: for each path P_i we add a zoom pass that corresponds to a maximal subpath of P_i that does not contain arcs from

bundles that do not appear in u , and contains a subword v_0^2 for some v_0 . Corollary 6.27 ensures that the only unknown exponent in our zoom instance is the number of turns along the spiral u , while the presence of all these zoom passes allow us to use Lemma 6.42 for any term v^q in $\text{bw}(P_t)$, not only to u^r . Hence, in the solution returned by Theorem 7.1 the simulated part of the path P_t needs to behave very similarly as in the original solution $(P_i)_{i=1}^k$, and we are able to guess the exponent r .

Let us now proceed to the formal argumentation. We perform multiple branching steps, one for each unknown exponent $r_{i,j}$. We start with the following set of definitions that formalizes the state of this branching procedure.

Definition 8.8 (potential spiral). Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance. Let u be a nonempty word over alphabet \mathbb{B} that contains each letter of \mathbb{B} at most once. Then u is a *potential spiral* if, for each two consecutive bundles $B_1 B_2$ on u , the components where the arcs of B_1 end and the arcs of B_2 start are equal. Moreover, u is a *potential closed spiral* if the components where the arcs of the last bundle of u end and the arcs of the first bundle of u start are equal.

Definition 8.9 (potential long spiral). Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance of isolation (Λ, d) where $\Lambda \geq 2$, $d \geq \max(2k, f(k, k) + 4)$, and $f(k, t) = 2^{O(kt)}$ is the bound on the type- t bend promised by Lemma 3.7. A potential closed spiral u is a *potential long spiral* if either:

1. u contains at least one normal bundle and does not contain any ring bundle of level $\Lambda - 1$ or lower; or
2. u contains only ring bundles of $\text{cl}(H^{\text{ring}})$ for some ring component H^{ring} , and does not contain non-isolation bundles of different levels.

Definition 8.10 (partial bundle word). Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance of isolation (Λ, d) where $\Lambda \geq 2$, $d \geq \max(2k, f(k, k) + 4)$, and $f(k, t) = 2^{O(kt)}$ is the bound on the type- t bend promised by Lemma 3.7. A *partial bundle word* is a (formal) string $\pi = u_1^{\rho_1} u_2^{\rho_2} \dots u_s^{\rho_s}$, where

- $s \leq 2|\mathbb{B}|$;
- each u_j is a potential spiral;
- for each $1 \leq j < s$, there exists a symbol of \mathbb{B} that appears in exactly one of the two words u_j and u_{j+1} ;
- if a bundle $B \in \mathbb{B}$, if B appears in u_{j_1} and in u_{j_2} , then B appears in all words u_j for $j_1 \leq j \leq j_2$;
- each ρ_j is either a positive integer or a symbol $?$; moreover, if $\rho_j = ?$ then u_j is a potential long spiral and if $\rho_j \neq 1$ then u_j is a potential closed spiral.

Definition 8.11 (consistent with partial bundle word). A bundle word u is *consistent* with a partial bundle word $u_1^{\rho_1} u_2^{\rho_2} \dots u_s^{\rho_s}$ if there exists positive integers r_1, r_2, \dots, r_s such that $u = u_1^{r_1} u_2^{r_2} \dots u_s^{r_s}$ and, for each $1 \leq j \leq s$, either $r_j = \rho_j$, or $\rho_j = ?$; in the second case we require that $r_j > 4$.

Definition 8.12 (semi-complete partial bundle word). A partial bundle word $u_1^{\rho_1} u_2^{\rho_2} \dots u_s^{\rho_s}$ is *semi-complete* if $\rho_j = ?$ implies that u_j contains at least one level-0 ring bundle.

Note that, if $\rho_j = ?$ in some semi-complete partial bundle word, then (as $(G, \mathcal{D}, \mathbb{B})$ has isolation (Λ, d) with $\Lambda \geq 2$) the word u_j contains only level-0, level-1 and level-2 ring bundles and does not contain any normal bundle.

We now note the following about a minimal solution to k -DPP.

Lemma 8.13. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance of isolation (Λ, d) where $\Lambda \geq 2$, $d \geq \max(2k, f(k, k) + 4)$, and $f(k, t) = 2^{O(kt)}$ is the bound on the type- t bend promised by Lemma 3.7. Assume $(G, \mathcal{D}, \mathbb{B})$ is a YES-instance to k -DPP and let $(P_i)_{i=1}^k$ be a minimal solution. For each $1 \leq i \leq k$, let $u_{i,1}^{r_{i,1}} u_{i,2}^{r_{i,2}} \dots u_{i,s(i)}^{r_{i,s(i)}}$ be a decomposition of $\mathbf{bw}(P_i)$ given by Corollary 6.23. Let $\rho_{i,j} = r_{i,j}$ if $r_{i,j} \leq 4$ or $u_{i,j}$ is not a potential long spiral, and $\rho_{i,j} = ?$ otherwise. Moreover, let $\rho'_{i,j} = ?$ if $\rho_{i,j} = ?$ and $u_{i,j}$ contains at least one level-0 ring bundle, and $\rho'_{i,j} = r_{i,j}$ otherwise. Then, for each $1 \leq i \leq k$:*

1. $u_{i,1}^{\rho_{i,1}} u_{i,2}^{\rho_{i,2}} \dots u_{i,s(i)}^{\rho_{i,s(i)}}$ is a partial bundle word consistent with $\mathbf{bw}(P_i)$;
2. $u_{i,1}^{\rho'_{i,1}} u_{i,2}^{\rho'_{i,2}} \dots u_{i,s(i)}^{\rho'_{i,s(i)}}$ is a semi-complete partial bundle word consistent with $\mathbf{bw}(P_i)$;
3. for each $1 \leq j \leq s(i)$, if $\rho_{i,j} \neq ?$ then $\rho'_{i,j} = \rho_{i,j} = r_{i,j}$;
4. for each $1 \leq j \leq s(i)$, if $\rho_{i,j} = ?$ then $\rho_{i,j} \leq 4|\mathbb{B}|^2 k^2 + 1$.

Proof. Each word $u_{i,j}$ is a potential spiral due to the properties promised by Corollary 6.23 and the fact that they are subwords of bundle words of paths in G . Moreover, Corollary 6.23 as well as the definitions of $\rho_{i,j}$ and $\rho'_{i,j}$ directly imply Claims 1, 2. Claim 3 follows directly from the definitions. What remains is to show Claim 4.

To this end, fix a choice of $1 \leq i \leq k$ and $1 \leq j \leq s(i)$ and assume that $\rho_{i,j} = r_{i,j}$. If $r_{i,j} \leq 4$, the claim is obvious, so let $r_{i,j} > 4$ which, by the definition of $\rho_{i,j}$, implies that $u_{i,j}$ is not a potential long spiral. By the definition of a potential long spiral, $u_{i,j}$ contains at least one arc of bundle of level $\Lambda - 1$ or lower. If $u_{i,j}$ contains at least one normal bundle, then the word $u_{i,j}^{r_{i,j}}$ (which is a subword of $\mathbf{bw}(P_i)$) contains at least $r_{i,j} - 1$ pairwise equal disjoint subwords, each corresponding to an isolation passage or a ring visitor. By Theorem 6.44 and Corollary 6.47, $r_{i,j} \leq 4|\mathbb{B}|^2 k^2 + 1$, as desired.

In the other case, assume that all bundles of $u_{i,j}$ are ring bundles of some ring component H^{ring} , and the subpath of P_i that corresponds to the bundle word $u_{i,j}^{r_{i,j}}$ is completely contained in $\text{cl}(H^{\text{ring}})$. If $u_{i,j}$ contains bundle arcs from non-isolation bundles of different levels, then a situation forbidden by Theorem 6.45 appears on the path P_i . This completes the proof of Claim 4. \square

Lemma 8.13 allows us to perform the following branching step.

Lemma 8.14. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance of isolation (Λ, d) where $\Lambda \geq 2$ and $d \geq \max(2k, f(k, k) + 4)$, where $f(k, t) = 2^{O(kt)}$ is the bound on the type- t bend promised by Lemma 3.7. Then in $O(2^{O(k|\mathbb{B}|^2 \log |\mathbb{B}|)} |G|^{O(1)})$ time one can enumerate a set of at most $2^{O(k|\mathbb{B}|^2 \log |\mathbb{B}|)}$ sequences $(\pi_i)_{i=1}^k$ of partial bundle words, such that for any minimal solution $(P_i)_{i=1}^k$ to k -DPP there exists a generated sequence $(\pi_i)_{i=1}^k$ such that $\mathbf{bw}(P_i)$ is consistent with π_i for all $1 \leq i \leq k$.*

Proof. By Lemma 8.13, it is sufficient to prove that for any $\beta > 1$ there are at most $2^{O(|\mathbb{B}|^2 \log |\mathbb{B}| + |\mathbb{B}| \log \beta)}$ partial bundle words $u_1^{\rho_1} u_2^{\rho_2} \dots u_s^{\rho_s}$ with $s \leq 2|\mathbb{B}|$ and $\rho_j \leq \beta$ whenever $\rho_j \neq ?$ (recall that $2k \leq |\mathbb{B}| - 1$) and they can be generated in $O(2^{O(|\mathbb{B}|^2 \log |\mathbb{B}| + |\mathbb{B}| \log \beta)} |G|^{O(1)})$ time.

To this end, note that there are less than $|\mathbb{B}| \cdot |\mathbb{B}|!$ words u over alphabet \mathbb{B} that do not contain a symbol twice, and can be enumerated with polynomial delay. As we require $s \leq 2|\mathbb{B}|$, there are only $2^{O(|\mathbb{B}|^2 \log |\mathbb{B}|)}$ choices for the values of s and the strings u_j , $1 \leq j \leq s$. Moreover, there are at most $\beta + 2$ choices of each value ρ_j . Note that one can verify in polynomial time if a given choice of s , u_j s and ρ_j s yields indeed a partial bundle word. \square

Lemma 8.14 enables us to guess partial bundle words that are consistent with bundle words of a minimal solution to the given bundled instance $(G, \mathcal{D}, \mathbb{B})$. Thus, henceforth we assume that,

apart from a bundled instance $(G, \mathcal{D}, \mathbb{B})$, we are given a set $(\pi_i)_{i=1}^k$ of partial bundle words and we look for a minimal solution $(P_i)_{i=1}^k$ such that $\mathbf{bw}(P_i)$ is consistent with π_i for each $1 \leq i \leq k$. By Lemma 8.14, there are at most $2^{O(k|\mathbb{B}|^2 \log |\mathbb{B}|)}$ subcases to consider.

Our goal now is to show a branching procedure that yields a bounded in $|\mathcal{D}|$, k and $|\mathbb{B}|$ number of subcases of evaluating all values $\rho_{i,j}$ for which $\rho_{i,j} = ?$ but $u_{i,j}$ does not contain any level-0 ring bundle. In other words, we aim to produce a bounded number of semi-complete partial bundle words and seek for minimal solutions consistent with one of them.

To achieve this goal, we analyze exponents $\rho_{i,j} = ?$ one by one, and prove that there is only a bounded number of choices for $r_{i,j}$, given the guesses made so far.

Formally, let $\pi_i = u_{i,1}^{\rho_{i,1}} u_{i,2}^{\rho_{i,2}} \dots u_{i,s(i)}^{\rho_{i,s(i)}}$. Assume that $1 \leq \iota \leq k$ and $1 \leq \eta \leq s(\iota)$ are chosen that $\rho_{\iota,\eta} = ?$, $u_{\iota,\eta}$ does not contain any level-0 ring bundle, but for any i and j such that $|u_{i,j}| < |u_{\iota,\eta}|$, if the exponent $\rho_{i,j} = ?$ then $u_{i,j}$ contains least one level-0 ring bundle. In other words, we guess the exponents $\rho_{\iota,\eta}$, starting from the shorter strings $u_{\iota,\eta}$.

In one step, we identify a set of possible values of for the exponent $r_{\iota,\eta}$, whose size is bounded as a function of $|\mathcal{D}|$, $|\mathbb{B}|$ and k , and branch into a number of subcases, replacing the value of $\rho_{\iota,\eta}$ with one of the elements of the identified set.

Note that a bundle word of a path P_i needs to start with the bundle that consists of the arc incident to the first terminal of the i -th terminal pair, and ends with a bundle that consists of the arc incident to the second terminal of the i -th terminal pair. Moreover, these arcs appear only once in the bundle word of P_i . Therefore, we may assume that for each i , $\rho_{i,1} = \rho_{i,s(i)} = 1$ and the corresponding words $u_{i,1}$ and $u_{i,s(i)}$ start and end respectively with the appropriate bundle; if that is not the case, we may terminate the current branch.

Let $(P_i)_{i=1}^k$ be a (hypothetical) solution to k -DPP on G , such that $\mathbf{bw}(P_i)$ is consistent with π_i for all $1 \leq i \leq k$, i.e., let $\mathbf{bw}(P_i) = u_{i,1}^{r_{i,1}} u_{i,2}^{r_{i,2}} \dots u_{i,s(i)}^{r_{i,s(i)}}$. As $\rho_{\iota,\eta} = ?$, $r_{\iota,\eta} > 4$ and $u_{\iota,\eta}^{r_{\iota,\eta}}$ contains a spiral $u_{\iota,\eta} B$, where B is the first symbol of $u_{\iota,\eta}$. Note that $u_{\iota,\eta}$ needs to be a potential long spiral in the bundle graph (as $\rho_{\iota,\eta} = ?$), and it does not contain any level-0 ring bundles. In particular, we may use Corollary 6.11 and speak of bundles inside and outside $u_{\iota,\eta}$.

Directly from Corollary 6.11 we have the following.

Lemma 8.15. *Consider a term $u_{\iota,\zeta}^{r_{\iota,\zeta}}$ in $\mathbf{bw}(P_\iota)$ for some $1 \leq \zeta \leq s(\iota)$ where $u_{\iota,\zeta}$ does not contain any level-0 ring bundles and $r_{\iota,\zeta} > 4$. Then $u_{\iota,\zeta}$ is a potential long spiral and the terminals of the ι -th pair lie on different sides of the closed walk $u_{\iota,\zeta}$ in the bundle graph $(\mathcal{D}, \mathbb{B})$. Moreover, for any $1 \leq i \leq k$ either:*

1. *both terminals of the i -th pair and all bundles of $\mathbf{bw}(P_i)$ lie inside the closed walk $u_{\iota,\zeta}$;*
2. *both terminals of the i -th pair and all bundles of $\mathbf{bw}(P_i)$ lie outside the closed walk $u_{\iota,\zeta}$;*
3. *the terminals of the i -th pair lie on different sides of the closed walk $u_{\iota,\zeta}$ and the starting terminal of the i -th pair lies inside $u_{\iota,\zeta}$ if and only if the starting terminal of the ι -th pair does.*

Let I_ζ^{\rightarrow} be the set of those indices i for which P_i satisfies the last option in Lemma 8.15 for the term $u_{\iota,\zeta}^{r_{\iota,\zeta}}$.

Let $u_{i,j}^{r_{i,j}}$ be a term of the bundle word decomposition of P_i where $r_{i,j} \geq 2$ and $u_{i,j}$ does not contain any level-0 ring bundles. Let $B_{i,j,1}$ be the last bundle on $\mathbf{bw}(P_i)$ that lies on the same side of $u_{i,j}$ as the starting terminal of the i -th pair, and $B_{i,j,2}$ be the first bundle of $\mathbf{bw}(P_i)$ that lies on the same side of $u_{i,j}$ as the ending terminal of the i -th pair. Let $u_{i,\eta(i,j,1)}$ be the word that contains the last occurrence of $B_{i,j,1}$ on $\mathbf{bw}(P_i)$ and $u_{i,j,\eta(i,2)}$ be the word that contains the first occurrence of

$B_{i,j,2}$ on $\mathbf{bw}(P_i)$. Let $P_{i,j}$ be the subpath of P_i between the ending point of the arc corresponding to the last occurrence of $B_{i,j,1}$ and the starting point of the arc corresponding to the first occurrence of $B_{i,j,2}$ in $\mathbf{bw}(P_i)$.




Figure 40: A nontrivial situation around spiral $u_{i,j}^{r_{i,j}}$ on P_i . Path P_i first spirals two times using four bundles (red spiral), then makes the $u_{i,j}^{r_{i,j}}$ spiral by spiraling two times using six bundles (pink spiral), and finally makes one turn of a spiral using three bundles (blue spiral). Note that the set of bundles used in the red spiral and in the blue spiral is a subset of the set of bundles used in the pink spiral, which we aim to measure. The crucial observation of this section is that these sets of bundles must be in fact **proper** subsets of the set used by the pink spiral.

Lemma 8.16. *Let $u_{i,j}^{r_{i,j}}$ be a term of the bundle word decomposition of P_i where $r_{i,j} \geq 2$ and $u_{i,j}$ does not contain any level-0 ring bundles. Then $B_{i,j,1}$ occurs on $\mathbf{bw}(P_i)$ before $B_{i,j,2}$ and $\eta(i, j, 1) < j < \eta(i, j, 2)$. Moreover, for any $j' \neq j$, $\eta(i, j, 1) < j' < \eta(i, j, 2)$ we have $|u_{i,j'}| < |u_{i,j}|$.*

Proof. The first claim follows directly from Corollary 6.11. As for the second claim, by the properties of the bundle word decomposition, if $\eta(i, j, 1) + 1 < j$, there exists a symbol $\bar{B}_{i,j,1}$ that appears in $u_{i,j}$ but not in $u_{i,j'}$ for any $\eta(i, j, 1) < j' < j$. Thus, as for each such j' the set of symbols of $u_{i,j'}$ is not only a subset of the set of symbols of $u_{i,j}$, but also a proper subset, and consequently $|u_{i,j'}| < |u_{i,j}|$. A symmetrical argument holds for the case $j' > j$ and a symbol $\bar{B}_{i,j,2}$ is missing from all words $u_{i,j'}$ for $j < j' < \eta(i, j, 2)$. \square

Lemma 8.17. Consider a term $u_{\iota, \zeta}^{r_{\iota, \zeta}}$ in $\mathbf{bw}(P_\iota)$ for some $1 \leq \zeta \leq s(\iota)$ where $u_{\iota, \zeta}$ does not contain any level-0 ring bundles and $r_{\iota, \zeta} > 4$. Let $B_{\iota, \zeta}$ be the first symbol of $u_{\iota, \zeta}$. Then, for any $i \in I_\zeta^\rightarrow$

1. there exists a unique index $1 \leq j(i, \zeta) \leq s(i)$ such that $r_{i, j(i, \zeta)} \geq 2$ and $u_{i, j(i, \zeta)}$ is a cyclic shift of $u_{\iota, \zeta}$;
2. $|r_{i, j(i, \zeta)} - r_{\iota, \zeta}| \leq 3$.

Proof. The first claim, as well as inequality $r_{i, j(i, \zeta)} \geq r_{\iota, \zeta} - 3$ follows from Corollary 6.27 applied for $P_0 = P_\iota$ and $\mathcal{P} = \{P_0, P_i\}$. On the other hand, assume for the sake of contradiction that $r_{\iota, \zeta} < r_{i, j(i, \zeta)} - 3$. If we apply Corollary 6.27 to $P_0 = P_i$ and $\mathcal{P} = \{P_i, P_\iota\}$ (note that we can do it as $r_{i, j(i, \zeta)} \geq 5$), we infer that $\mathbf{bw}(P_\iota)$ must contain a term v^q where v is a cyclic shift of $u_{i, j(i, \zeta)}$ and $q \geq r_{i, j(i, \zeta)} - 3$. As $\mathbf{bw}(P_\iota)$ already contains $u_{\iota, \zeta}^{r_{\iota, \zeta}}$, this implies that $v = u_{\iota, \zeta}$, $q = r_{\iota, \zeta}$, and, consequently $r_{\iota, \zeta} \geq r_{i, j(i, \zeta)} - 3$. This is a contradiction. \square

Note that the values of I_ζ^\rightarrow , $j(i, \zeta)$ as well as $B_{i, j, 1}$, $B_{i, j, 2}$, $\eta(i, j, 1)$ and $\eta(i, j, 2)$ for valid values of i , j and ζ , are known, given the partial bundle words $(\pi_i)_{i=1}^k$: the values of $r_{i, j}$ for $\rho_{i, j} = ?$ are not necessary to compute these values. If there is some inconsistency in the current branch (say, there is no unique candidate for $j(i, \zeta)$), we terminate the current branch.

Moreover, $\mathbf{bw}(P_{i, j})$ is consistent with the partial bundle word

$$\hat{u}_{i, \eta(i, j, 1)} u_{i, \eta(i, j, 1)+1}^{\rho_{i, \eta(i, j, 1)+1}} \cdots u_{i, \eta(i, j, 2)-1}^{\rho_{i, \eta(i, j, 2)-1}} \hat{u}_{i, \eta(i, j, 2)},$$

where $\hat{u}_{i, \eta(i, j, 1)}$ is the maximal suffix of $u_{i, \eta(i, j, 1)}$ that does not contain $B_{i, j, 1}$ and $\hat{u}_{i, \eta(i, j, 2)}$ is the maximal prefix of $u_{i, \eta(i, j, 2)}$ that does not contain $B_{i, j, 2}$ (note that any of these two words may be empty).

Recall that we aim to guess $r_{\iota, \eta}$. We claim that, if there exists a minimal solution $(P_i)_{i=1}^k$ such that P_i is consistent with π_i for each $1 \leq i \leq k$, then for any $i \in I_\eta^\rightarrow$ and any $\eta(i, j(i, \eta), 1) < j < \eta(i, j(i, \eta), 2)$ the exponent $\rho_{i, j}$ may be equal to ? only for $j = j(i, \eta)$. Indeed, take any such j . As $u_{\iota, \eta}$ does not contain any level-0 ring bundles, $u_{i, j}$ does not contain as well. By Lemma 8.16, $|u_{i, j}| < |u_{\iota, \eta}|$ unless $j = j(i, \eta)$. The claim follows from our chosen order of guessing of the exponents $\rho_{i, j}$. Therefore we may safely terminate branches where the claim is not satisfied.

Moreover, by Lemma 8.17, there exist integers $(\alpha_i)_{i \in I_\eta^\rightarrow}$, $-3 \leq \alpha_i \leq 3$, $\alpha_\iota = 0$, and a single integer $4 < \aleph \leq n$ such that $\rho_{i, j(i, \eta)} = \aleph + \alpha_i$ for any $i \in I_\eta^\rightarrow$. We branch into at most 7^{k-1} options, guessing the values of α_i for $i \in I_\eta^\rightarrow$. If for any $i \in I_\eta^\rightarrow$, the value $\rho_{i, j(i, \eta)}$ does not equal ?, the value of $\rho_{\iota, \eta}$ is determined. Thus, henceforth we assume that this is not the case.

To choose a good value for $\rho_{\iota, \eta}$, we construct a zoom and a pack of zoom passes. Let $\hat{\mathbb{B}}$ be the set of bundles that appear in $u_{\iota, \eta}$, and let $\hat{\mathcal{D}}$ be the set of components that contain endpoints of arcs of bundles of $\hat{\mathbb{B}}$. Clearly, $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ is a zoom. As $u_{\iota, \eta}$ does not contain any level-0 ring bundles, $\hat{\mathcal{D}}$ does not contain any ring components, and hence is level-0 safe.

Fix $4 < \aleph \leq n$. We are to construct a pack of zoom passes, parameterized by \aleph . Intuitively, we want to reproduce what happens in all the spirals $u_{i, \zeta}^{r_{i, \zeta}}$ for $\eta(i, \eta, 1) < \zeta < \eta(i, \eta, 2)$, so that the equivalent of the path $P_{\iota, \eta}$ will behave in our zoom instance in a very similar way to the original (unknown to us) path $P_{\iota, \eta}$.

To this end, for each partial bundle word π_i we say that a pair (a, b) , $1 \leq a \leq b \leq s(i)$ is *relevant* if

1. each $u_{i, j}$, $a \leq j \leq b$ contains only symbols that appear in $u_{\iota, \eta}$;
2. at least one exponent $\rho_{i, j}$, $a \leq j \leq b$, does not equal 1.

A pair (a, b) is a *maximal relevant* pair if neither $(a - 1, b)$ nor $(a, b + 1)$ is a relevant pair. Recall that $u_{i,1}$ and $u_{i,s(i)}$ contains bundles incident to terminals, and thus $a > 1$ and $b < s(i)$ for each relevant pair (a, b) and, consequently, for any relevant pair (a, b) there exists a unique maximal relevant pair (a', b') with $a' \leq a \leq b \leq b'$.

By definition, for any $i \in I_\eta^\rightarrow$, the pair $(\eta(i, j(i, \eta), 1) + 1, \eta(i, j(i, \eta), 2) - 1)$ is a maximal relevant pair in π_i . We now note that, in the parts of the partial bundle words π_i that correspond to relevant pairs, almost every exponent $\rho_{i,j}$ is known.




Figure 41: An example showing that there may be relevant pairs that satisfy the second option in Lemma 8.18. The path P_1 spirals multiple times using five bundles (the red spiral), while path P_2 spirals using a subset of these bundles (blue spiral) creating a relevant pair. Note that this happens even though both of the terminals of P_2 are enclosed by the red spiral.

Lemma 8.18. *Let (a, b) be a maximal relevant pair in π_i . Then exactly one of the following holds:*

1. $i \in I_\eta^\rightarrow$, $a = \eta(i, j(i, \eta), 1) + 1$, $b = \eta(i, j(i, \eta), 2) - 1$ (in particular, $\rho_{i,j} = ?$ for $a \leq j \leq b$ if and only if $j = j(i, \eta)$); or
2. for any $a \leq j \leq b$, $|u_{i,j}| < |u_{i,\eta}|$ or $\rho_{i,j} = 1$ (in particular, $\rho_{i,j} \neq ?$ for $a \leq j \leq b$).

Proof. Note that if the first option is satisfied, then there exists $j = j(i, \eta)$, for which $\rho_{i,j} = ?$ and $|u_{i,j}| = |u_{i,\eta}|$, so the second option is not satisfied. We are left with proving that if the second option is not satisfied, then the first is.

Let (a, b) be a maximal relevant pair in π_i that does not satisfy the second option from the statement of the lemma. That is, there exists j , $a \leq j \leq b$, such that $\rho_{i,j} \neq 1$ (and, consequently, $r_{i,j} > 1$) but $|u_{i,j}| = |u_{i,\eta}|$. Thus, $u_{i,j}$ is a permutation of $u_{i,\eta}$. By Corollary 6.27, $u_{i,j}$ is a cyclic shift of $u_{i,\eta}$, $i \in I_\eta^\rightarrow$ and, by the uniqueness of $j(i, \eta)$, $j = j(i, \eta)$. \square

For a fixed value of $4 < \aleph \leq n$, for any π_i and any maximal relevant pair (a, b) in π_i , we define the following bundle word:

$$q_{i,(a,b)}(\aleph) = B'_{i,(a,b),1} \hat{u}_{i,(a,b),1} u_{i,a}^{\rho'_{i,a}} u_{i,a+1}^{\rho'_{i,a+1}} \cdots u_{i,b}^{\rho'_{i,b}} \hat{u}_{i,(a,b),2} B'_{i,(a,b),2},$$

where:

1. $\rho'_{i,j(i,\eta)} = \aleph + \alpha_i$ and $\rho'_{i,j} = \rho_{i,j}$ for $j \neq j(i, \eta)$;
2. $\hat{u}_{i,(a,b),1}$ is the maximal suffix of $u_{i,a-1}$ that contains only symbols that appear in $u_{\iota,\eta}$ and $B'_{i,(a,b),1}$ is the symbol of $u_{i,a-1}$ immediately preceding $\hat{u}_{i,(a,b),1}$;
3. symmetrically, $\hat{u}_{i,(a,b),2}$ is the maximal prefix of $u_{i,b+1}$ that contains only symbols that appear in $u_{\iota,\eta}$ and $B'_{i,(a,b),2}$ is the symbol of $u_{i,b+1}$ immediately succeeding $\hat{u}_{i,(a,b),2}$.

Note that both $\hat{u}_{i,(a,b),1}$ and $\hat{u}_{i,(a,b),2}$ may be empty.

We observe that, by Lemma 8.18, $q_{i,(a,b)}(\aleph)$ is a bundle word: all exponents are integers. As $\hat{\mathcal{D}}$ does not contain any ring components, $q_{i,(a,b)}(\aleph)$ is also a level-0 zoom pass in $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$. Let I be the set of pairs $(i, (a, b))$ where $1 \leq i \leq k$ and (a, b) is a maximal relevant pair in π_i . As $s(i) \leq 2|\mathbb{B}|$ for each i , we infer that $|I| \leq 2|\mathbb{B}|k$.

Using Lemma 8.4, branch into at most $(|I|!)^2 \leq ((2|\mathbb{B}|k)!)^2$ subcases, guessing the permutations $(\psi_{B,\alpha})_{B \in \mathbb{B}, 1 \leq \alpha \leq 2}$ for which $((q_\tau(\aleph))_{\tau \in I}, (\psi_{B,\alpha})_{B \in \mathbb{B}, 1 \leq \alpha \leq 2})$ is a pack of level-0 zoom passes in $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ for any $4 < \aleph \leq n$ (note that the set of possible options for permutations $\psi_{B,\alpha}$ does not depend on \aleph). For $B \in \mathbb{B}$ and $1 \leq \alpha \leq 2$, define $I_{B,\alpha} = \{\tau \in I : B_{\tau,\alpha} = B\}$; the permutation $\psi_{B,\alpha}$ permutes $I_{B,\alpha}$.

Construct the zoom auxiliary graph and instance for this pack of zoom passes in $(\hat{\mathcal{D}}, \hat{\mathbb{B}})$ and denote it $(H, \hat{\mathcal{D}}_H, \hat{\mathbb{B}}_H)$.

The discussion in Section 8.1 concluded with Observation 8.7 immediately yields the following.

Lemma 8.19. *If there exists a solution $(P_i)_{i=1}^k$ to the bundled instance $(G, \mathcal{D}, \mathbb{B})$ such that*

1. *for each $1 \leq i \leq k$, $\mathbf{bw}(P_i)$ is consistent with π_i , $\mathbf{bw}(P_i) = u_{i,1}^{r_{i,1}} u_{i,2}^{r_{i,2}} \dots u_{i,s(i)}^{r_{i,s(i)}}$;*
2. *for each $i \in I$, $r_{i,j(\eta)} = \alpha_i + r_{\iota,\eta}$;*
3. *if we denote for $\tau = (i, (a, b)) \in I$ by Q_τ the subpath of P_i that corresponds to the subword $q_\tau(r_{\iota,\eta})$ of $\mathbf{bw}(P_i)$, then for each $B \in \mathbb{B}$ the permutation $\psi_{B,1}$ is equal to the order of first arcs of paths Q_τ on B for $\tau \in I_{B,1}$ and the permutation $\psi_{B,2}$ is equal to the order of the last arcs of paths Q_τ on B for $\tau \in I_{B,2}$;*

then there exists a solution $(Q'_\tau)_{\tau \in I}$ to the constructed zoom auxiliary instance such that $\mathbf{bw}(Q'_\tau)$ equals $q_\tau(r_{\iota,\eta})$ up to the prefix and suffix that corresponds to the part of the path contained in the zoom starting and ending gadgets.

Let us solve the constructed zoom auxiliary instance $(H, \hat{\mathcal{D}}_H, \hat{\mathbb{B}}_H)$ using Theorem 7.1, (note that $\hat{\mathcal{D}}$ and $\hat{\mathcal{D}}_H$ do not contain any ring components). For a fixed choice of the permutations $\psi_{B,\alpha}$, $B \in \mathbb{B}$, $1 \leq \alpha \leq 2$, we find a minimum \aleph_0 such that Theorem 7.1 returns a solution for bundled instance $(H, \mathcal{D}_H, \mathbb{B}_H)$ and bundle words $(q_\tau(\aleph_0))_{\tau \in I}$. If there is no such \aleph_0 , by Lemma 8.19, we may terminate the current branch. Otherwise we note the following.

Lemma 8.20. *Let $(P_i)_{i=1}^k$ be a solution to the bundled instance $(G, \mathcal{D}, \mathbb{B})$ as in Lemma 8.19, and suppose that $(P_i)_{i=1}^k$ is minimal. Then $\aleph_0 \leq r_{\iota,\eta} \leq \aleph_0 + 32|\mathbb{B}|$.*

Proof. The inequality $\aleph_0 \leq r_{\iota,\eta}$ is straightforward by the choice of \aleph_0 and Lemma 8.19.

Let $(Q'_\tau)_{\tau \in I}$ be the family of paths returned by Theorem 7.1 for bundle words $(q_\tau(\aleph_0))_{\tau \in I}$. Let $t = (\iota, (\eta(\iota, \eta, 1) + 1, \eta(\iota, \eta, 2) - 1)) \in I$. We claim that for each $\eta(\iota, \eta, 1) < \zeta < \eta(\iota, \eta, 2)$ such that $r_{\iota,\zeta} > 10$, the bundle word $\mathbf{bw}(Q'_t)$ contains $u_{\iota,\zeta}^{r_{\iota,\zeta}-10}$ as a subword and those subwords are pairwise disjoint for different choices of ζ .

Indeed, consider the subword $u_{\iota,\zeta}^{r_{\iota,\zeta}}$ of $\mathbf{bw}(P_\iota)$ and the corresponding spiraling ring A_ζ associated with the subpath of P_ι corresponding to $u_{\iota,\zeta}^{r_{\iota,\zeta}}$ with borders $\gamma_{\zeta,1}$ and $\gamma_{\zeta,2}$ and faces $f_{\zeta,1}$ and $f_{\zeta,2}$. As $r_{\iota,\zeta} > 10$, for any $i \in I_\zeta^-$ the index $j(i, \zeta)$ is defined, $u_{i,j(i,\zeta)}$ is a cyclic shift of $u_{\iota,\zeta}$ and $r_{i,j(i,\zeta)} \geq r_{\iota,\zeta} - 3 > 7$. Therefore there exists an element $\tau_i = (i, (a_i, b_i)) \in I$ such that $a_i \leq j(i, \zeta) \leq b_i$, the terminals of the pair τ_i in the bundled instance $(H, \hat{\mathcal{D}}_H, \hat{\mathbb{B}}_H)$ lie on different sides of the spiral $u_{\iota,\zeta}$ and the path Q'_{τ_i} contains a subpath R'_{τ_i} that starts in a vertex on $\gamma_{\zeta,1}$ and ends in a vertex on $\gamma_{\zeta,2}$. As $(P_i)_{i=1}^k$ is a minimal solution, by Lemma 6.42 the bundle word of each path R'_{τ_i} contains $u_{\iota,\zeta}^{r_{\iota,\zeta}-10}$ as a subword. Since the spiraling rings A_ζ are disjoint for different choices of ζ , the paths R'_{τ_i} are edge-disjoint for different choices of ζ . As $t \in I_\zeta^-$ for any choice of ζ , the claim is proven.

We infer that

$$|\mathbf{bw}(Q'_t)| \geq \sum_{\zeta=\eta(\iota,\eta,1)+1}^{\eta(\iota,\eta,2)-1} (r_{\iota,\zeta} - 10)|u_{\iota,\zeta}|.$$

On the other hand, as $\mathbf{bw}(Q'_t)$ contains a subset (as a multiset) of the symbols of $q_t(\aleph_0)$, we have that

$$|\mathbf{bw}(Q'_t)| \leq |q_t(\aleph_0)| = 2 + |\hat{u}_{t,1}| + |\hat{u}_{t,2}| + \sum_{\zeta=\eta(\iota,\eta,1)+1}^{\eta(\iota,\eta,2)-1} \rho'_{\iota,\zeta} |u_{\iota,\zeta}|.$$

Recall $\rho'_{\iota,\zeta} = r_{\iota,\zeta}$ for $\zeta \neq \eta$ and $\rho'_{\iota,\eta} = \aleph_0$. Moreover, $|u_{\iota,\zeta}| \leq |u_{\iota,\eta}|$ and $|\hat{u}_{t,\alpha}| \leq |u_{\iota,\eta}|$ for $\alpha = 1, 2$. We infer that

$$(r_{\iota,\eta} - 10 - \aleph_0)|u_{\iota,\eta}| \leq 2 + 10 \cdot (2 + \eta(\iota, \eta, 2) - \eta(\iota, \eta, 1) - 2)|u_{\iota,\eta}|.$$

As $\eta(\iota, \eta, 2) - \eta(\iota, \eta, 1) < s(\iota) \leq 2|\mathbb{B}|$, we have $r_{\iota,\eta} \leq \aleph_0 + 32|\mathbb{B}|$, as desired. \square

Lemma 8.20 allows us to conclude with the following lemma that summarizes the branching steps made in this section.

Lemma 8.21. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance of isolation (Λ, d) where $\Lambda \geq 2$ and $d \geq \max(2k, f(k, k) + 4)$, where $f(k, t) = 2^{O(kt)}$ is the bound on the type- t bend promised by Lemma 3.7. Then in $2^{O(k^2|\mathbb{B}|^2 \log |\mathbb{B}|)} |G|^{O(1)}$ time one can compute a family of at most $2^{O(k^2|\mathbb{B}|^2 \log |\mathbb{B}|)}$ semi-complete sequences of partial bundle words $(\pi_i)_{i=1}^k$ such that for any minimal solution $(P_i)_{i=1}^k$ to k -DPP on $(G, \mathcal{D}, \mathbb{B})$, there exists a generated sequence $(\pi_i)_{i=1}^k$ in the set such that P_i is consistent with π_i for each $1 \leq i \leq k$.*

Proof. We first branch into $2^{O(k|\mathbb{B}|^2 \log |\mathbb{B}|)}$ subcases, guessing the initial partial bundle word π_i for each $1 \leq i \leq k$, using Lemma 8.14. Then, for each unknown exponent $\rho_{\iota,\eta}$, in the order of increasing lengths of $|u_{\iota,\eta}|$, we guess the value of $\rho_{\iota,\eta}$. Recall that this includes guessing the values α_i (at most 7^{k-1} options) permutations $(\psi_{B,\alpha})_{B \in \mathbb{B}, 1 \leq \alpha \leq 2}$ (at most $(|I|!)^2 \leq ((2|\mathbb{B}|k)!)^2$ options) and a value $r_{\iota,\eta}$ between \aleph_0 and $\aleph_0 + 32|\mathbb{B}|$. Therefore we have at most $2^{O(k|\mathbb{B}| \log |\mathbb{B}|)}$ subcases for each exponent $\rho_{\iota,\eta}$ to guess.

Recall that in each π_i we have $s(i) \leq 2|\mathbb{B}|$. Therefore, we perform the aforementioned guessing step at most $2|\mathbb{B}|k$ times. The promised bound follows.

Finally, note that if $\rho_{i,j} = ?$ implies that $u_{i,j}$ contains a level-0 ring bundle for any $1 \leq i \leq k$, $1 \leq j \leq s(i)$, then $(\pi_i)_{i=1}^k$ are semi-complete by the definition. \square

8.3 Ring components: deducing winding numbers

In the previous section we have shown that there is a bounded number of semi-complete partial bundle words to consider. Here our goal is to change this semi-complete partial bundle words into bundle words with ring holes. The main difficulty is to find a set of good candidates for paths' winding numbers in the ring components. To cope with this, we use Lemma 4.8: if we know which parts of paths traverse a ring component, and we find one way to route them through a ring component, there exists a solution that winds in the ring component similarly as the way we have found.

However, there are two main technical problems with this approach. First, the paths may visit an isolation of a ring component, but do not traverse the ring component itself (i.e., there are ring visitors). These visitors block space for rerouting: we cannot use Lemma 4.8 directly to a ring component or some fixed closure of it. Here the rescue comes from results developed in Section 6.7 that help us control the behaviour of a minimal solution in the closure of a ring component.

A second problem is that, if we ask Theorem 7.1 to provide us with some canonical way to route ring passages through (a closure of) a ring component, the returned solution follows our guidelines (i.e., bundle words with ring holes) in a quite relaxed way. To cope with that, we employ a similar line of reasoning as in the previous subsection: if in a minimal solution a ring passage spirals along a bundle word u^r , for some large r , then Lemma 6.42 forces any canonical way found by Theorem 7.1 to spiral at least $r - 10$ times (i.e., to contain u^{r-10} in its bundle word). Hence, the solution returned by Theorem 7.1 can differ from the minimal solution only by a limited number bundles, which implies that their winding numbers also do not differ much.

In this section we assume that the isolation of our decomposition is (Λ, d) for $\Lambda \geq 3$ and $d \geq \max(2k, f(k, k) + 4)$. The assumption $d \geq \max(2k, f(k, k) + 4)$ allows us to use the results of Section 6.7. The assumed 3 layers of isolation gives us space to carefully extract the ring on which Lemma 4.8 is applied. It is worth noticing that all essential argumentation happens in layers 1 and 2; the last layer are added only for the sake of clarity of the presentation (for example, we do not need to care about normal bundles with both endpoints in the same level- Λ isolation component etc.).

Let us now proceed with a formal argumentation. We first note that a semi-complete partial bundle word contains more information than the bundle word part of a bundle word with level-1 ring holes.

Lemma 8.22. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance with isolation (Λ, d) where $\Lambda \geq 2$ and $d \geq \max(2k, f(k, k) + 4)$ and $f(k, t) = 2^{O(kt)}$ is the bound on the type- t bend promised by Lemma 3.7. Let π be a semi-complete partial bundle word in $(G, \mathcal{D}, \mathbb{B})$. Then there exists a unique sequence of bundle words $(p_j)_{j=0}^h$ such that the following holds: $(p_j)_{j=0}^h$ does not contain any level-0 bundles and for any path P that connects a terminal pair, is consistent with π and its unique bundle word with level-1 ring holes does not contain any level-0 ring bundle, there exists a choice of integers $(w_j)_{j=1}^h$ such that $((p_j)_{j=0}^h, (w_j)_{j=1}^h)$ is a bundle word with level-1 ring holes consistent with P . Moreover, the sequence $(p_j)_{j=0}^h$ can be computed in polynomial time, given π .*

Proof. Let $\pi = u_1^{\rho_1} u_2^{\rho_2} \dots u_s^{\rho_s}$. Assume P is consistent with π and let $\mathbf{bw}(P) = u_1^{r_1} u_2^{r_2} \dots u_s^{r_s}$. Note that, as P connects a terminal pair, $r_1 = r_s = 1$, u_1 starts with a bundle that contains the arc incident to the starting terminal of P and u_s ends with a bundle that contains the arc incident to the ending terminal of P .

Consider now an index j for which $\rho_j = ?$. As π is semi-complete, u_j contains at least one level-0 ring bundle, is a potential long spiral and, consequently, does not contain non-isolation bundles of level different than 0. From the assumption that the bundle word with level-1 ring holes of P does

not contain any level-0 ring bundle, we infer that the subpath of P that corresponds to the subword $u_j^{r_j}$ is a part of a level-1 ring passage of P . As the choice of j is arbitrary, we infer that the bundle word part of the bundle word with level-1 ring holes of P does not depend on the choice of P , but only on π .

Moreover, the aforementioned argument yields a polynomial-time algorithm to compute the bundle words $(p_j)_{j=0}^h$ from π . We compute p defined as a bundle word created from π by evaluating each $\rho_j = ?$ to a fixed positive integer. Then we compute the decomposition $p = p_0 r_1 p_1 r_2 p_2 \dots r_h p_h$, where $(r_j)_{j=1}^h$ are all level-1 ring passages in p , and output the sequence $(p_j)_{j=0}^h$. \square

Lemma 8.23. *Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance with isolation (Λ, d) where $\Lambda \geq 2$ and $d \geq \max(2k, f(k, k) + 4)$, $f(k, t) = 2^{O(kt)}$ is the bound on the type- t bend promised by Lemma 3.7. Let $(P_i)_{i=1}^k$ be a minimal solution to k -DPP on G such that P_i is consistent with π_i for each $1 \leq i \leq k$. Let $(p_{i,j})_{j=0}^{h(i)}$ be a sequence computed by Lemma 8.22 for π_i . Then $\sum_{i=1}^k h(i) \leq 4|\mathbb{B}|^2 k^2$ and there exist integers $(w_{i,j})_{1 \leq i \leq k, 1 \leq j \leq h(i)}$ such that for each $1 \leq i \leq k$ the pair $((p_{i,j})_{j=0}^{h(i)}, (w_{i,j})_{j=1}^{h(i)})$ is a bundle word with level-1 ring holes consistent with P_i .*

Proof. The bound on $\sum_{i=1}^k h(i)$ follows from Theorem 6.44. As for the second claim, note that by Corollary 6.48, for any $1 \leq i \leq k$, the bundle word with level-1 ring holes consistent with P_i does not contain any bundle of level 0. \square

From this point, we assume that the isolation of the bundled instance $(G, \mathcal{D}, \mathbb{B})$ satisfies $\Lambda \geq 3$ and $d \geq (2k, f(k, k) + 4)$, as in the assumptions of Theorem 8.1.

Fix $1 \leq i \leq k$ and $1 \leq j \leq h(i)$. Let $B_{i,j,1}$ be the last bundle of $p_{i,j-1}$ and $B_{i,j,2}$ be the first bundle of $p_{i,j}$. Directly from Lemma 8.23 we obtain the following observation.

Observation 8.24. *If there exists a minimal solution $(P_i)_{i=1}^k$ such that P_i is consistent with π_i for each $1 \leq i \leq k$, then for each $1 \leq i \leq k$ and $1 \leq j \leq h(i)$ there exists a ring component $H_{i,j}^{ring}$ such that*

1. $B_{i,j,1}$ contains arcs leading from a level-2 isolation component of $H_{i,j}^{ring}$ to a level-1 isolation component;
2. $B_{i,j,2}$ contains arcs leading from a level-1 isolation component of $H_{i,j}^{ring}$ to a level-2 isolation component;
3. $B_{i,j,1}$ and $B_{i,j,2}$ lie on different sides of $H_{i,j}^{ring}$.

Thus, if this is not the case, we may terminate the current branch.

For each ring component $H^{ring} \in \mathcal{D}$, we define $I(H^{ring}) = \{(i, j) : H^{ring} = H_{i,j}^{ring}\}$. Note the following, due to Lemma 8.23.

Observation 8.25. *If there exists a minimal solution $(P_i)_{i=1}^k$ to k -DPP on G such that P_i is consistent with π_i for each $1 \leq i \leq k$, we have*

$$\sum_{H^{ring} \in \mathcal{D}} |I(H^{ring})| = \sum_{i=1}^k h(i) \leq 4|\mathbb{B}|^2 k^2.$$

Again, if this is not the case, we terminate the current branch.

Observe that, by Theorem 6.45, we obtain the following.

Observation 8.26. *If there exists a minimal solution $(P_i)_{i=1}^k$ to k -DPP on G such that P_i is consistent with π_i , then for each $1 \leq i \leq k$ and $0 \leq j \leq h(i)$, the bundle word $p_{i,j}$ contains at least one normal bundle.*

Proof. The claim is obvious for $j = 0$ or $j = h(i)$, as then $p_{i,j}$ contains a bundle with an arc incident to a terminal. Assume that the claim is not true for some $1 \leq i \leq k$ and $0 < j < h(i)$. Then the subpath of P_i between the arcs corresponding to the last symbol of $p_{i,j-1}$ and the first symbol $p_{i,j+1}$ contains the structure forbidden by Theorem 6.45. \square

Again, if this is not the case, we terminate the current branch.

Recall that $\Lambda \geq 3$. For each $1 \leq i \leq k$ and $1 \leq j \leq h(i)$, we define $B_{i,j,1}^\circ$ to be the last bundle of $p_{i,j-1}$ that contains arcs leading from the level-3 isolation component of $H_{i,j}^{\text{ring}}$ to level-2 one, and $B_{i,j,2}^\circ$ to be the first bundle of $p_{i,j}$ that contains arcs leading from the level-2 isolation component of $H_{i,j}^{\text{ring}}$ to level-3 one. Let $p_{i,j,1}^\circ$ be the suffix of $p_{i,j-1}$ starting with $B_{i,j,1}^\circ$ and $p_{i,j,2}^\circ$ be the prefix of $p_{i,j}$ ending with $B_{i,j,2}^\circ$. Moreover, for $1 \leq i \leq k$, $0 \leq j \leq h(i)$ let $q_{i,j}$ be the subword of $p_{i,j}$ between $B_{i,j,2}^\circ$ and $B_{i,j+1,1}^\circ$, where $B_{i,0,2}^\circ$ is the first symbol of $p_{i,0}$ and $B_{i,h(i)+1,1}^\circ$ is the last symbol of $p_{i,h(i)}$.




Figure 42: An exemplary beginning of the j -th ring passage on path P_i . The blue circle depicts bundle $B_{i,j,1}^\circ$, where $q_{i,j-1}$ ends and $p_{i,j,1}^\circ$ starts. The orange circle depicts the last bundle of $p_{i,j-1}$ (thus also of $p_{i,j,1}^\circ$), where the level-1 hole starts.

By Observation 8.26 and again Theorem 6.45 we have the following.

Observation 8.27. *If there exists a minimal solution $(P_i)_{i=1}^k$ to k -DPP on G such that P_i is consistent with π_i , then the bundles $B_{i,j,1}^\circ$ and $B_{i,j,2}^\circ$ are well defined and lie on the opposite sides of $H_{i,j}^{\text{ring}}$. Moreover, for $1 \leq \alpha \leq 2$, the bundle word $p_{i,j,\alpha}^\circ$, except for the symbol $B_{i,j,\alpha}^\circ$, contains only bundles with arcs with both endpoints in level- λ , $1 \leq \lambda \leq 2$ isolation components of $H_{i,j}^{\text{ring}}$ that lie on the same side of $H_{i,j}^{\text{ring}}$ as $B_{i,j,\alpha}^\circ$.*

Again, if this is not the case, we terminate the current branch.

Now note that for ring passages in the same ring component H^{ring} , by Observation 4.4, the winding numbers cannot differ too much.

Observation 8.28. *Let $(P_i)_{i=1}^k$ be a minimal solution to k -DPP on G such that P_i is consistent with π_i for each $1 \leq i \leq k$. For $1 \leq i \leq k$, let $((p_{i,j})_{j=0}^{h(i)}, (w_{i,j})_{j=1}^{h(i)})$ be the bundle word with level-1 ring holes of P_i . Then, if for some (i, j) and (i', j') we have $H_{i,j}^{\text{ring}} = H_{i',j'}^{\text{ring}}$, then $|w_{i,j} - w_{i',j'}| \leq 1$.*

Observation 8.28 motivates us to the following branch. For each $1 \leq i \leq k$ and for each $1 \leq j \leq h(i)$ we branch into three subcases, picking an integer $-1 \leq \alpha_{i,j} \leq 1$. By Observation 8.25, this step leads to at most $3^{4|\mathbb{B}|^2 k^2}$ subcases. We say that a solution $(P_i)_{i=1}^k$ is *consistent with the current branch* for each ring component H^{ring} there exists an integer $w(H^{\text{ring}})$ such that for each $1 \leq i \leq k$, the bundle word with level-1 ring holes of P_i equals

$$((p_{i,j})_{j=0}^{h(i)}, (w(H_{i,j}^{\text{ring}}) + \alpha_{i,j})_{j=1}^{h(i)}).$$

By Observation 8.28 we obtain the following.

Observation 8.29. *If there exists a minimal solution $(P_i)_{i=1}^k$ to k -DPP on G such that P_i is consistent with π_i for each $1 \leq i \leq k$, then there exists a subcase with a choice of integers $\alpha_{i,j}$ such that $(P_i)_{i=1}^k$ is consistent with this branch.*

Our goal now is, for a fixed branch with integers $\alpha_{i,j}$ and for a fixed ring component $H^{\text{ring}} \in \mathcal{D}$ such that $I(H^{\text{ring}}) \neq \emptyset$, to compute a set of bounded size of possible candidates for $w(H^{\text{ring}})$.

For each ring component $H^{\text{ring}} \in \mathcal{D}$ such that $I(H^{\text{ring}}) \neq \emptyset$ and for each integer $-n \leq w \leq n$ we construct an zoom auxiliary instance $(H(H^{\text{ring}}, w), \hat{\mathcal{D}}_H(H^{\text{ring}}, w), \mathbb{B}_H(H^{\text{ring}}, w))$ as follows. First we take a zoom $(\hat{\mathcal{D}}(H^{\text{ring}}), \hat{\mathbb{B}}(H^{\text{ring}}))$ that includes all components in the level-2 closure of H^{ring} and bundles with arcs with both endpoints in these components. Then, for each $(i, j) \in I(H^{\text{ring}})$ we construct bundle word with level-2 ring holes $\mathfrak{p}_{i,j}^\circ = ((p_{i,j,1}^\circ, p_{i,j,2}^\circ), w + \alpha_{i,j})$; note that, by the definition of $p_{i,j,1}^\circ$ and $p_{i,j,2}^\circ$, this pair is indeed a bundle word with level-2 ring holes and a level-2 zoom pass in $(\hat{\mathcal{D}}(H^{\text{ring}}), \hat{\mathbb{B}}(H^{\text{ring}}))$ as well. Finally, using Lemma 8.4, we branch into at most $(|I(H^{\text{ring}})|)^2$ subcases, guessing, for each $B \in \mathbb{B}$ and $1 \leq \alpha \leq 2$, a permutation $\psi_{B,\alpha}(H^{\text{ring}})$ of those indices $(i, j) \in I(H^{\text{ring}})$ for which $B_{i,j,\alpha} = B$. Thus, $((\mathfrak{p}_{i,j}^\circ)_{(i,j) \in I(H^{\text{ring}})}, (\psi_{B,\alpha}(H^{\text{ring}}))_{B \in \mathbb{B}, 1 \leq \alpha \leq 2})$ is a pack of zoom passes in $(\hat{\mathcal{D}}(H^{\text{ring}}), \hat{\mathbb{B}}(H^{\text{ring}}))$. The zoom auxiliary instance $(H(H^{\text{ring}}, w), \hat{\mathcal{D}}_H(H^{\text{ring}}, w), \mathbb{B}_H(H^{\text{ring}}, w))$ is defined as the zoom auxiliary instance for this pack of zoom passes.

The discussion in Section 8.1 concluded with Observation 8.7 immediately yields the following.

Lemma 8.30. *If there exists a minimal solution $(P_i)_{i=1}^k$ to k -DPP on G such that*

1. *for each $1 \leq i \leq k$, P_i is consistent with π_i ;*
2. *for each ring component H^{ring} where $I(H^{\text{ring}}) \neq \emptyset$ there exists an integer $w(H^{\text{ring}})$ such that for each $1 \leq i \leq k$ and $1 \leq j \leq h(i)$ the level-1 ring passage that corresponds to the part of the bundle word of P_i between $p_{i,j-1}$ and $p_{i,j}$ has winding number $w(H^{\text{ring}}) + \alpha_{i,j}$;*
3. *for each ring component H^{ring} , $B \in \mathbb{B}$ and $1 \leq \alpha \leq 2$, the order of the arcs $b_{i,j,\alpha}$ for $(i, j) \in I(H^{\text{ring}})$ and $b_{i,j,\alpha} \in B$ is equal to $\psi_{B,\alpha}(H^{\text{ring}})$,*

then, for each H^{ring} where $I(H^{\text{ring}}) \neq \emptyset$ the zoom auxiliary instance constructed for the pair $(H^{\text{ring}}, w(H^{\text{ring}}))$ has a solution $(P_\tau^)_{\tau \in I(H^{\text{ring}})}$, where the bundle word with level-1 ring holes of P_τ^* equals \mathfrak{p}_τ° up to a prefix and suffix that corresponds to the subpath in the zoom starting and ending gadgets.*

For each ring component H^{ring} where $I(H^{\text{ring}}) \neq \emptyset$ and for each $-n \leq w \leq n$ we apply Theorem 7.1 to the zoom auxiliary instance $(H(H^{\text{ring}}, w), \hat{\mathcal{D}}_H(H^{\text{ring}}, w), \mathbb{B}_H(H^{\text{ring}}, w))$. For each H^{ring} , let $w(H^{\text{ring}})$ be an integer for which Theorem 7.1 returned a solution; if such an integer does not exist, by Lemma 8.30 we may safely terminate the current branch. Let $(P_\tau^*)_{\tau \in I(H^{\text{ring}})}$ be the solution returned by Theorem 7.1 and let w_τ^* be the winding number of P_τ^* in the level-2 closure of H^{ring} (which is a subgraph of both G and $H(H^{\text{ring}}, w(H^{\text{ring}}))$).

We now prove the following crucial claim.

Lemma 8.31. *If there exists a minimal solution $(P_i)_{i=1}^k$ to k -DPP on G such that*

1. *for each $1 \leq i \leq k$, P_i is consistent with π_i , and*
2. *for each ring component H^{ring} , $B \in \mathbb{B}$ and $1 \leq \alpha \leq 2$, the order of the arcs $b_{i,j,\alpha}$ for $(i, j) \in I(H^{\text{ring}})$ and $b_{i,j,\alpha} \in B$ is equal to $\psi_{B,\alpha}(H^{\text{ring}})$,*

then, there exists integers $(x_{i,j})_{1 \leq i \leq k, 1 \leq j \leq h(i)}$ such that $|x_{i,j} - w_{i,j}^| \leq 40|\mathbb{B}|^2 + 2|\mathbb{B}| + 8$ for each $1 \leq i \leq k$, $1 \leq j \leq h(i)$ and a solution $(P'_i)_{i=1}^k$ to k -DPP on G (not necessarily minimal) such that P'_i is consistent with bundle word with level-2 ring holes $((q_{i,j})_{j=0}^{h(i)}, (x_{i,j})_{j=1}^{h(i)})$ for each $1 \leq i \leq k$.*

Proof. Fix $H^{\text{ring}} \in \mathcal{D}$ for which $I(H^{\text{ring}}) \neq \emptyset$. Our goal is to modify $(P_i)_{i=1}^k$ in the level-2 closure of H^{ring} so that the winding numbers of passages of H^{ring} , indexed with $\tau \in I(H^{\text{ring}})$, do not differ from w_τ^* much.

First, we slightly modify the graph $\text{cl}^2(H^{\text{ring}})$, so that further topological arguments become cleaner. For each bundle $B = (b_1, b_2, \dots, b_s)$ contained in $\text{cl}^2(H^{\text{ring}})$ we first subdivide each arc b_j twice, introducing vertices $v_{j,1}$ and $v_{j,2}$, and then, for each $1 \leq j < s$ and $1 \leq \alpha \leq 2$, we add a vertex $z_{j,\alpha}$ and arcs $(v_{j,\alpha}, z_{j,\alpha})$ and $(v_{j+1,\alpha}, z_{j,\alpha})$ inside the face of G between b_j and b_{j+1} . We do it in such a manner that, if the reference curve $\gamma_{\text{ref}}^2(H^{\text{ring}})$ crosses B , it crosses arcs $(v_{j,1}, v_{j,2})$ for each $1 \leq j \leq s$, i.e., is contained between the undirected paths $v_{1,\alpha}z_{1,\alpha}v_{2,\alpha}z_{2,\alpha} \dots v_{s,\alpha}$, $1 \leq \alpha \leq 2$. Note that this operation does not change the answer to k -DPP on any supergraph of $\text{cl}^2(H^{\text{ring}})$, as the added arcs are useless from the point of view of the directed paths (vertices $z_{j,\alpha}$ are sinks). In the new graph the bundle B consists of arcs $(v_{j,1}, v_{j,2})$ for $1 \leq j \leq s$, the vertices $v_{j,1}, z_{j,1}$ belong to the component where the arcs of B originally start, and the vertices $v_{j,2}, z_{j,2}$ belong to the component where the arcs of B originally end.

Now we define the graph G^\sharp . We start with the subgraph of G induced by the vertices of the level-2 closure of H^{ring} . Then we repeatedly take maximal subpaths of paths $(P_i)_{i=1}^k$ that go through vertices of $\text{cl}^2(H^{\text{ring}})$ and, if such a path starts and ends on the same side of $\text{cl}^2(H^{\text{ring}})$, we remove from G^\sharp all arcs and vertices that lie on the subpath or on the different side of the chosen subpath than H^{ring} . As $(P_i)_{i=1}^k$ is a minimal solution, by Theorem 6.45, any such path contains only vertices of level 2, and G^\sharp contains the subgraph of G induced by the vertices of $\text{cl}^1(H^{\text{ring}})$.

We identify two faces f_1 and f_2 of G^\sharp that contain the outer and inner face of $\text{cl}(H^{\text{ring}})$, respectively. We want to choose a subcurve of $\gamma_{\text{ref}}^2(H^{\text{ring}})$ to be a reference curve in G^\sharp . Let f'_1 be the last face crossed by $\gamma_{\text{ref}}^2(H^{\text{ring}})$, contained in f_1 , and let f'_2 be the first face crossed by $\gamma_{\text{ref}}^2(H^{\text{ring}})$. By the construction of $\gamma_{\text{ref}}^2(H^{\text{ring}})$ (Lemma 6.35), f'_1 appears on $\gamma_{\text{ref}}^2(H^{\text{ring}})$ earlier than f'_2 . We choose γ_{ref} to be subcurve of $\gamma_{\text{ref}}^2(H^{\text{ring}})$ between leaving f'_1 and entering f'_2 . Note that γ_{ref} is a reference curve in G^\sharp , as it travels from the boundary of f_1 to the boundary of f_2 .

In the rest of the proof, we often measure winding numbers of different paths in G^\sharp with respect to either γ_{ref} or $\gamma_{\text{ref}}^2(H^{\text{ring}})$; note that, although the latter may not necessarily be a proper reference curve in G^\sharp (as it may visit f_1 and f_2 several times), the notion of winding number is properly defined. However, Lemma 4.8 requires us to use a proper reference curve γ_{ref} , for this reason we need to translate winding numbers between these two curves.

For each $\tau \in I(H^{\text{ring}})$, the path P_τ^* contains a subpath in G^\sharp connecting a vertex on f_1 with a vertex on f_2 (in one of the directions). Denote the first such path as Q_τ^* . Let $P_\tau^{*,1}$ be the subpath of P_τ^* from the start of P_τ^* up to the beginning of Q_τ^* and $P_\tau^{*,2}$ be the subpath of P_τ^* from the end of Q_τ^* to the end of P_τ^* .

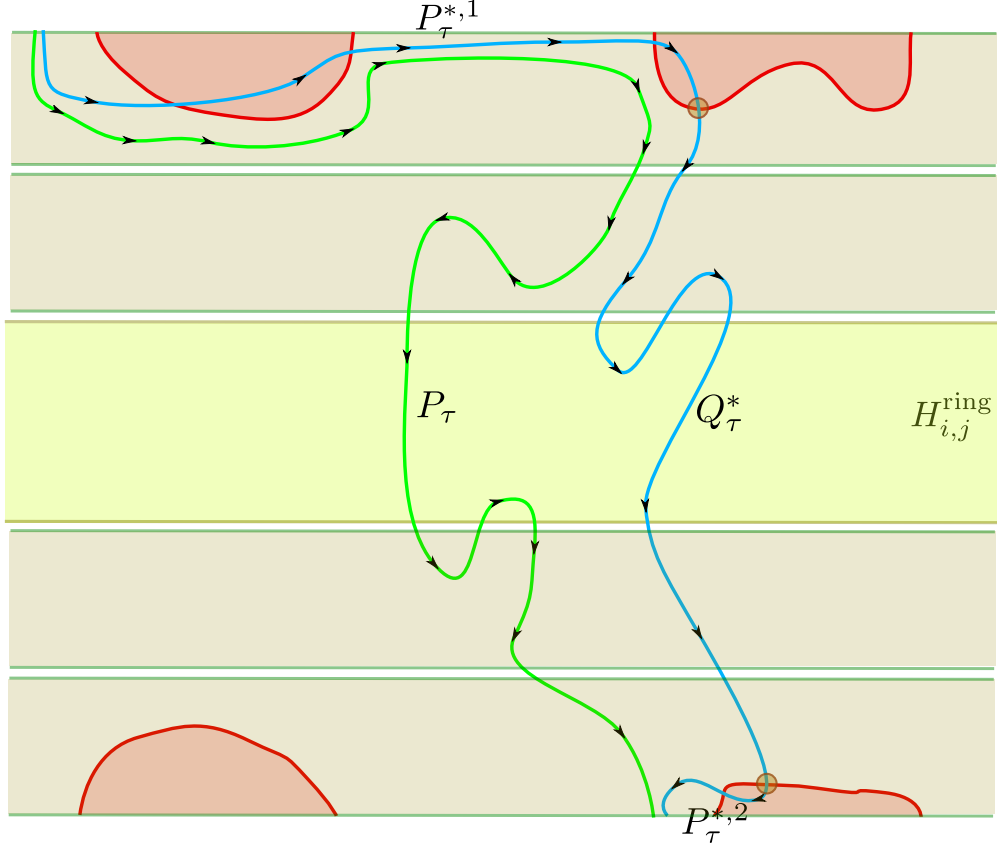


Figure 43: Situation in the proof of Lemma 8.31 for one passage index τ . The red parts depict parts of the graph removed while constructing G^\sharp , that is, parts separated from the ring component by paths from the minimal solution. The blue path is the path P_τ^* found by Theorem 7.1, while the green path is the passage induced by the minimal solution. The orange circles depict points where P_τ^* is split into its three parts: $P_\tau^{*,1}$, Q_τ^* , and $P_\tau^{*,2}$. The reader may verify that this picture can be realized so that in level 2, the blue path uses a subset of bundles used by the green path.

For $\alpha = 1, 2$, let $\mathbb{B}_\alpha^\circ \subseteq \hat{\mathbb{B}}(H^{\text{ring}})$ be the set of bundles consisting of all bundles with at least one endpoint in level-2 isolation component of H^{ring} that lie on the same side of H^{ring} as f_α . Moreover, let $\mathbb{B}^\circ = \mathbb{B}_1^\circ \cup \mathbb{B}_2^\circ$.

For a path R and a set $S \subseteq \mathbb{B}$, by $|\text{bw}(R) \cap S|$ we denote the number of appearances of a bundle from S in $\text{bw}(R)$. We claim the following.

Claim 8.32.

$$|\text{bw}(P_\tau^{*,1}) \cap \mathbb{B}^\circ| + |\text{bw}(P_\tau^{*,2}) \cap \mathbb{B}^\circ| \leq 20|\mathbb{B}|^2.$$

Proof. Consider two cases. First, assume that there are two paths $P_{\tau_1}^*$ and $P_{\tau_2}^*$, $\tau_1, \tau_2 \in I$, that go in different direction, that is, $P_{\tau_1}^*$ starts on f_1 and ends on f_2 and $P_{\tau_2}^*$ starts on f_2 and ends on f_1 . Then any path P_τ^* for $\tau \in I(H^{\text{ring}})$ cannot contain two arcs of the same bundle, as otherwise its

bundle word would contain a spiral and the existence of both paths $P_{\tau_1}^*$ and $P_{\tau_2}^*$ would contradict Lemma 6.10. Therefore, P_τ^* contains at most $|\mathbb{B}|$ arcs that are bundle arcs and the claim is proven.

In the other case, all paths P_τ^* , $\tau \in H^{\text{ring}}$ go in the same direction; without loss of generality, let us assume that they start in a vertex on f_1 and end in a vertex on f_2 . Note that, by the choice of Q_τ^* , any path Q_τ^* also starts in a vertex on f_1 and ends in a vertex on f_2 . Fix $\tau = (i, j) \in I(H^{\text{ring}})$ and let P_τ be a level-2 ring passage of H^{ring} of the path P_i (from the solution $(P_i)_{i=1}^k$) that corresponds to the bundle word with level-1 ring holes \mathfrak{p}_τ° . Let $B \in \mathbb{B}^\circ$ and assume B appears in $\text{bw}(P_\tau)$. Then, as \mathfrak{p}_τ° is a bundle word with level-1 ring holes, B appears in some term u^r of a bundle word decomposition of \mathfrak{p}_τ° . As \mathfrak{p}_τ° does not contain any level-0 bundle, neither does u .

Assume for a moment that $r > 10$. Note that, by the definition of the ring G^\sharp , the spiraling ring A associated with the term u^r in the path P_i is contained in G^\sharp . By Lemma 6.10, each path $Q_{\tau'}^*$ for $\tau' \in I(H^{\text{ring}})$ traverses A . By Lemma 6.42, $\text{bw}(Q_{\tau'}^*)$ needs to contain u^{r-10} as a subword. Moreover, as the spiraling rings A are disjoint for different terms u^r , the subwords u^{r-10} of $\text{bw}(Q_{\tau'}^*)$ are pairwise disjoint for different terms u^r .

By the properties of the solution $(P_\tau^*)_{\tau \in I(H^{\text{ring}})}$ returned by Theorem 7.1, the number of appearances of B in the bundle word of the path P_τ^* is not greater than the number of appearances of B in \mathfrak{p}_τ° . There are at most $2|\mathbb{B}|$ terms in a bundle word decomposition of \mathfrak{p}_τ° , and for each term u^r , a subword $u^{\min(0, r-10)}$ appears in $\text{bw}(Q_\tau^*)$, and these subwords are pairwise disjoint for different terms u^r . We infer that B appears at most $2|\mathbb{B}| \cdot 10$ times in $\text{bw}(P_\tau^{*,1})$ and $\text{bw}(P_\tau^{*,2})$ in total. This finishes the proof of the claim. \square

We also need the following observations.

Claim 8.33. *For any $\alpha = 1, 2$ and for any maximal subpath R of $P_\tau^{*,\alpha}$ that does not contain any bundles of \mathbb{B}_α° the winding number of R with respect to $\gamma_{\text{ref}}^2(H^{\text{ring}})$ equals $-1, 0$ or $+1$.*

Proof. Recall that the reference curve $\gamma_{\text{ref}}^2(H^{\text{ring}})$ has properties promised by Lemma 6.35. If R is contained in the level-2 isolation component of H^{ring} that lies on the same side of H^{ring} as $B_{\tau,\alpha}$, then R does not cross the reference curve and its winding number is 0.

Otherwise, recall that both endpoints of Q_τ^* , and thus $P_\tau^{*,\alpha}$ as well, lie outside $\text{cl}^1(H^{\text{ring}})$, since $\text{cl}^1(H^{\text{ring}})$ is contained in G^\sharp . Therefore the endpoints of R lie on the inside or outside face of $\text{cl}^1(H^{\text{ring}})$, and, as we exclude only bundles from \mathbb{B}_α° for either $\alpha = 1$ or $\alpha = 2$, they lie on the same side of R . Obtain a closed curve γ from R by connecting the endpoints of R using parts of arcs that precede and succeed R on $P_\tau^{*,\alpha}$ and the arcs connecting the level-1 isolation and level-2 isolation component of H^{ring} on the same side as $B_{\tau,\alpha}$ in the dual of G , in such a manner that γ does not separate the sides of $\text{cl}(H^{\text{ring}})$. Note that this implies that the winding number of γ is 0, whereas, by the properties of $\gamma_{\text{ref}}(H^{\text{ring}})$, $\gamma \setminus R$ has winding number $+1, 0$ or -1 . The claim follows. \square

Claim 8.34. *Let P be a path in G^\sharp connecting f_1 with f_2 . Then the winding number of P with respect to curve γ_{ref} differs by at most $|\mathbb{B}|$ from the winding number of P with respect to $\gamma_{\text{ref}}^2(H^{\text{ring}})$.*

Proof. Let $f^1, f^2, \dots, f^s = f_1'$ be faces of $\text{cl}^2(H^{\text{ring}})$ crossed by $\gamma_{\text{ref}}^2(H^{\text{ring}})$, contained in f_1 , in the order of their appearance on $\gamma_{\text{ref}}^2(H^{\text{ring}})$. For $1 \leq j < s$, let $\gamma_{\text{ref}}^{R,j}$ be the subcurve of $\gamma_{\text{ref}}^2(H^{\text{ring}})$ between f^j and f^{j+1} .

Let B be a bundle crossed by $\gamma_{\text{ref}}^2(H^{\text{ring}})$. We claim that at most one curve $\gamma_{\text{ref}}^{R,j}$ may intersect arcs of B . Recall the construction of G^\sharp ; let P be one of the maximal subpaths of a solution $(P_i)_{i=1}^k$ that goes through the vertices of $\text{cl}^2(H^{\text{ring}})$ and has both endpoints on the same side of H^{ring} as the face f_1 . Let $B = (b_1, b_2, \dots, b_s)$. Due to the subdivision of bundles we performed at the beginning of the proof, any maximal subpath of P that consists of vertices and edges incident to the union

of faces between arcs b_j, b_{j+1} , $1 \leq j < s$, in fact consists of a single arc b_η for some $1 \leq \eta \leq s$. However, due to Lemma 6.10, P does not traverse B twice, and arcs b_j for $j \geq \eta$ or $j \leq \eta$ are removed from G^\sharp . Hence, the number of curves $\gamma_{\text{ref}}^{R,j}$, $s-1$, is not larger than the number of bundles that lie on the same side of H^{ring} as f_1 .

For each $1 \leq j < s$, close the curve $\gamma_{\text{ref}}^{R,j}$ inside f_1 to obtain a closed curve. The winding number of P with regards to the closed curve $\gamma_{\text{ref}}^{R,j}$ is 0, +1 or -1. By performing the same analysis on the side of H^{ring} that contains f_2 , the claim follows. \square

We may now conclude with the following statement.

Claim 8.35. *The winding number of Q_τ^* with respect to γ_{ref} differs from w_τ^* by at most $40|\mathbb{B}|^2 + |\mathbb{B}| + 2$.*

Proof. By Claim 8.33, we infer that the winding number of $P_\tau^{*,1}$, with respect to $\gamma_{\text{ref}}^2(H^{\text{ring}})$, is at most $2|\text{bw}(P_\tau^{*,1}) \cap \mathbb{B}_1^\circ| + 1$, and a similar claim holds for $P_\tau^{*,2}$. Hence, by Claim 8.32, the winding number of Q_τ^* with respect to $\gamma_{\text{ref}}^2(H^{\text{ring}})$ does not differ from w_τ^* by more than $40|\mathbb{B}|^2 + 2$. By Claim 8.34, the winding numbers of Q_τ^* with respect to $\gamma_{\text{ref}}^2(H^{\text{ring}})$ and γ_{ref} differ by at most $|\mathbb{B}|$. The claim follows by pipelining the above three bounds. \square

Recall that, by the definition of the graph G^\sharp , the intersection of the solution $(P_i)_{i=1}^k$ with G^\sharp is a set of paths $(P_\tau)_{\tau \in I(H^{\text{ring}})}$; each such path is a level-2 ring passage consistent with \mathfrak{p}_τ° , with the first and last bundle removed. Moreover, note the following.

Claim 8.36. *The (circular) orders of the starting and ending vertices of $(P_\tau)_{\tau \in H^{\text{ring}}}$ on the faces f_1 and f_2 of G_R is exactly the same as the order of the starting and ending vertices of the paths $(Q_\tau^*)_{\tau \in H^{\text{ring}}}$.*

Proof. Recall that for each ring component H^{ring} , $B \in \mathbb{B}$ and $1 \leq \alpha \leq 2$, the order of the arcs $b_{i,j,\alpha}$ for $(i,j) \in I(H^{\text{ring}})$ and $b_{i,j,\alpha} \in B$ is equal to $\psi_{B,\alpha}(H^{\text{ring}})$. By the construction of the zoom auxiliary graph $H(H^{\text{ring}}, w(H^{\text{ring}}))$, in $\text{cl}^2(H^{\text{ring}})$, the orders of the starting and ending vertices of $(P_\tau)_{\tau \in H^{\text{ring}}}$ and the paths $(P_\tau^*)_{\tau \in H^{\text{ring}}}$ are equal.

Consider now the following graph $G^{\sharp/2}$, constructed similarly as G^\sharp , but the removing procedure is performed only on the side of f_2 . We start with the subgraph of G induced by the vertices of the level-2 closure of H^{ring} . Then we repeatedly take maximal subpaths of paths $(P_i)_{i=1}^k$ that go through vertices of $\text{cl}^2(H^{\text{ring}})$ and, if such a path starts and ends on the same side of $\text{cl}^2(H^{\text{ring}})$ as the face f_2 , we remove from $G^{\sharp/2}$ all arcs and vertices that lie on the subpath or on the different side of the chosen subpath than H^{ring} . Note that $G^{\sharp/2}$ is a supergraph of G^\sharp and a subgraph of $\text{cl}^2(H^{\text{ring}})$. One of its faces is f_2 , and the other face is one of the faces of $\text{cl}^2(H^{\text{ring}})$ that is contained in f_1 ; let us denote it f'_1 .

Note that for any $\tau \in I(H^{\text{ring}})$, since Q_τ^* is the *first* subpath of P_τ^* that connects f_1 with f_2 , $P_\tau^{*,1} \cup Q_\tau^*$ connects f'_1 with f_2 inside $G^{\sharp/2}$. Hence, the order of the starting vertices of the paths P_τ^* on f'_1 is equal to the order of the ending vertices of the paths Q_τ^* on f_2 . On the other hand, if we look at G^\sharp , the order of the starting vertices of the paths Q_τ^* on f_1 is equal to the order of the ending vertices of Q_τ^* on f_2 . This concludes the proof of the claim. \square

Claim 8.36 allows us to apply Lemma 4.8 for paths $(P_\tau)_{\tau \in I(H^{\text{ring}})}$ and $(Q_\tau^*)_{\tau \in I(H^{\text{ring}})}$ in the rooted ring G_R with reference curve γ_{ref} , obtaining a sequence of vertex-disjoint paths $(P'_\tau)_{\tau \in I(H^{\text{ring}})}$, such that for each $\tau \in I(H^{\text{ring}})$ the path P'_τ has the same starting and ending vertex as P_τ , but the winding numbers of P'_τ and Q_τ^* with respect to γ_{ref} differ by at most 6. Let x_τ be the winding number of P'_τ with respect to $\gamma_{\text{ref}}^2(H^{\text{ring}})$.

Claim 8.37. For each $\tau \in I(H^{\text{ring}})$ we have $|x_\tau - w_\tau^*| \leq 40|\mathbb{B}|^2 + 2|\mathbb{B}| + 8$.

Proof. By Claim 8.35, the winding number of Q_τ^* with respect to γ_{ref} and w_τ^* , differ by at most $40|\mathbb{B}|^2 + |\mathbb{B}| + 2$. By Lemma 4.8, the winding numbers of Q_τ^* and P'_τ with respect to γ_{ref} differ by at most 6. By Claim 8.34, the winding number of P'_τ with respect to γ_{ref} and x_τ differ by at most $|\mathbb{B}|$. The claim follows by pipelining the above three bounds. \square

Recall that the paths $(P'_\tau)_{\tau \in I(H^{\text{ring}})}$ are vertex-disjoint, are contained in G_R , and $(P_\tau)_{\tau \in I(H^{\text{ring}})}$ are the only parts of $(P_i)_{i=1}^k$ in G_R . Thus, if we conduct the same argument for each ring component H^{ring} with $I(H^{\text{ring}}) \neq \emptyset$ and replace in the solution $(P_i)_{i=1}^k$ each subpath $P_{i,j}$ with $P'_{i,j}$ for $1 \leq i \leq k$, $1 \leq j \leq h(i)$, we obtain another solution $(P'_i)_{i=1}^k$ to k -DPP on G . Moreover, as we modified only subpaths $P_{i,j}$ for $1 \leq i \leq k$, $1 \leq j \leq h(i)$, each path P'_i for $1 \leq i \leq k$ is consistent with the bundle word with level-2 ring holes $((q_{i,j})_{j=0}^{h(i)}, (x_{i,j})_{j=1}^{h(i)})$, as desired. This completes the proof of the lemma. \square

We can now summarize with the following lemma.

Lemma 8.38. Let $(G, \mathcal{D}, \mathbb{B})$ be a bundled instance of isolation (Λ, d) where $\Lambda \geq 3$, $d \geq \max(2k, f(k, k) + 4)$, and $f(k, t) = 2^{O(kt)}$ is the bound on the type- t bend promised by Lemma 3.7. Assume we are given a sequence $(\pi_i)_{i=1}^k$ of semi-complete partial bundle words. Then in $O(2^{O(k^2|\mathbb{B}|^2 \log |\mathbb{B}|)} |G|^{O(1)})$ time one can compute a family of at most $2^{O(k^2|\mathbb{B}|^2 \log |\mathbb{B}|)}$ sequences $(\mathbf{p}_i)_{i=1}^k$ of bundle words with level-2 ring holes such that if there exists a minimal solution $(P_i)_{i=1}^k$ to k -DPP on G such that P_i is consistent with π_i for each $1 \leq i \leq k$, then there exists a solution $(P'_i)_{i=1}^k$ to k -DPP on G and generated sequence $(\mathbf{p}_i)_{i=1}^k$ such that P_i is consistent with \mathbf{p}_i for each $1 \leq i \leq k$.

Proof. First recall that we branch into at most $3^{4|\mathbb{B}|^2 k^2}$ subcases guessing the values $\alpha_{i,j}$ and further at most

$$\left(\left(\sum_{i=1}^k h(i) \right)! \right)^2 \leq ((4|\mathbb{B}|^2 k^2)!)^2 = 2^{O(|\mathbb{B}|^2 k^2 \log |\mathbb{B}|)}$$

subcases when we choose permutations $(\psi_{B,\alpha}(H^{\text{ring}}))_{B \in \mathbb{B}, 1 \leq \alpha \leq 2}$ for each $H^{\text{ring}} \in \mathcal{D}$ with $I(H^{\text{ring}}) \neq \emptyset$. Finally, once we compute integers $(w_{i,j})_{1 \leq i \leq k, 1 \leq j \leq h(i)}$, we may guess the values $(x_{i,j})_{1 \leq i \leq k, 1 \leq j \leq h(i)}$; by Lemma 8.31 for each $1 \leq i \leq k$, $1 \leq j \leq h(i)$ there are $80|\mathbb{B}|^2 + 4|\mathbb{B}| + 17$ possible values for $x_{i,j}$. As $\sum_{i=1}^k h(i) \leq 4|\mathbb{B}|^2 k^2$, we end up with the promised running time and number of subcases. Correctness follows from Lemmata 8.30 and 8.31. \square

By pipelining Lemma 8.21 with Lemma 8.38 we finish the proof of Theorem 8.1. Note that the output bundle words with level-2 ring holes does not contain any bundle of level 0 due to the application of Lemma 8.22.

9 Summary: proof of the main result

We are now ready to summarize the results of the previous sections by formally proving Theorem 1.2.

Proof of Theorem 1.2. Given an instance G of k -DPP with k terminal pairs, we first apply the Decomposition Theorem (Theorem 5.8) on G with constants $\Lambda = 3$, $d = \max(2k, f(k, k) + 4)$ and $r = d(k) + 1$, where $f(k, t) = 2^{O(kt)}$ is the bound on the type- t bend of Lemma 3.7 and $d(k) = 2^{O(k^2)}$ is the bound on the number of concentric cycles of Theorem 3.2. If Theorem 5.8 returns a set of

r concentric cycles, we delete any vertex of the innermost cycle and restart the algorithm. The correctness follows from Theorem 3.2, and the algorithm is restarted at most $|V(G)|$ times.

Otherwise, the algorithm of Theorem 5.8, in time $O(2^{O(\Lambda(d+r)k^2)}|G|^{O(1)}) = O(2^{2^{O(k^2)}}|G|^{O(1)})$ returns a set of $2^{O(\Lambda(d+r)k^2)} = 2^{2^{O(k^2)}}$ pairs (G_i, \mathcal{D}_i) ; by Theorem 5.8 it suffices to check if any graph G_i is a YES-instance to k -DPP. Thus, from this point we investigate one graph G_i with decomposition \mathcal{D}_i . Note that $|\mathcal{D}_i| = O(k^2)$ and the alternation of \mathcal{D}_i is $O(\Lambda(d+r)k^2) = 2^{O(k^2)}$.

We first apply the bundle recognition algorithm of Lemma 6.4 to obtain a bundled instance $(G_i, \mathcal{D}_i, \mathbb{B}_i)$ with $|\mathbb{B}_i| = 2^{O(k^2)}$. Then we apply the algorithm of Theorem 8.1; note that the values of d and Λ are large enough to allow this step. We obtain a family \mathcal{F}_i of sequences $(\mathbf{p}_i)_{i=1}^k$ of bundle words with level-2 ring holes with a promise that, if G_i is a YES-instance to k -DPP, then there exists a solution consistent with one of the sequences. As $|\mathbb{B}_i| = 2^{O(k^2)}$, the size of the family \mathcal{F}_i is bounded by $2^{2^{O(k^2)}}$ and the running time of the algorithm of Theorem 8.1 is bounded by $O(2^{2^{O(k^2)}}|G|^{O(1)})$.

Moreover, Theorem 8.1 promises us that the bundle words with level-2 ring holes of \mathcal{F}_i do not contain any level-0 ring bundles. Thus, we may apply the algorithm of Theorem 7.1 to each element of \mathcal{F}_i . If it outputs a solution for some sequence $(\mathbf{p}_i)_{i=1}^k$, we know that G_i is a YES-instance to k -DPP. Otherwise, by Theorem 7.1, there does not exist a solution to k -DPP on G_i consistent with an element of \mathcal{F}_i , and, by Theorem 8.1, G_i is a NO-instance. This finishes the proof of Theorem 1.2. \square

References

- [1] I. Adler, S. G. Kolliopoulos, P. K. Krause, D. Lokshtanov, S. Saurabh, and D. M. Thilikos. Tight bounds for linkages in planar graphs. In L. Aceto, M. Henzinger, and J. Sgall, editors, *ICALP (1)*, volume 6755 of *Lecture Notes in Computer Science*, pages 110–121. Springer, 2011.
- [2] D. Berwanger, A. Dawar, P. Hunter, S. Kreutzer, and J. Obdržálek. The dag-width of directed graphs. *J. Comb. Theory, Ser. B*, 102(4):900–923, 2012.
- [3] H. L. Bodlaender, M. R. Fellows, and M. T. Hallett. Beyond NP-completeness for problems of bounded width: hardness for the W hierarchy. In *STOC*, pages 449–458, 1994.
- [4] R. Diestel, T. R. Jensen, K. Y. Gorbunov, and C. Thomassen. Highly connected sets and the excluded grid theorem. *J. Comb. Theory, Ser. B*, 75(1):61–73, 1999.
- [5] G. Ding, A. Schrijver, and P. D. Seymour. Disjoint paths in a planar graph - a general theorem. *SIAM J. Discrete Math.*, 5(1):112–116, 1992.
- [6] G. Ding, A. Schrijver, and P. D. Seymour. Disjoint cycles in directed graphs on the torus and the klein bottle. *J. Comb. Theory, Ser. B*, 58(1):40–45, 1993.
- [7] R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Monographs in Computer Science. Springer, New York, 1999.
- [8] S. Fortune, J. Hopcroft, and J. Wyllie. The directed subgraph homeomorphism problem. *Theoret. Comput. Sci.*, 10(2):111–121, 1980.
- [9] R. Ganian, P. Hliněný, J. Kneis, A. Langer, J. Obdržálek, and P. Rossmanith. On digraph width measures in parameterized algorithmics. In *IWPEC*, pages 185–197, 2009.

- [10] R. Ganian, P. Hlinený, J. Kneis, D. Meister, J. Obdržálek, P. Rossmanith, and S. Sikdar. Are there any good digraph width measures? In *IPEC*, pages 135–146, 2010.
- [11] M. Grohe, K. Kawarabayashi, D. Marx, and P. Wollan. Finding topological subgraphs is fixed-parameter tractable. In *Proceedings of the 43rd ACM Symposium on Theory of Computing*, pages 479–488, 2011.
- [12] Q.-P. Gu and H. Tamaki. Improved bounds on the planar branchwidth with respect to the largest grid minor size. *Algorithmica*, 64(3):416–453, 2012.
- [13] T. Johnson, N. Robertson, P. D. Seymour, and R. Thomas. Unpublished.
- [14] T. Johnson, N. Robertson, P. D. Seymour, and R. Thomas. Directed tree-width. *J. Comb. Theory, Ser. B*, 82(1):138–154, 2001.
- [15] K. Kawarabayashi and S. Kreutzer. An excluded grid theorem for digraphs with forbidden minors, 2013. Manuscript.
- [16] K. Kawarabayashi, B. A. Reed, and P. Wollan. The graph minor algorithm with parity conditions. In *FOCS*, pages 27–36, 2011.
- [17] K. Kawarabayashi and P. Wollan. A shorter proof of the graph minor algorithm: the unique linkage theorem. In *STOC*, pages 687–694, 2010.
- [18] S. Kreutzer and S. Ordyniak. Digraph decompositions and monotonicity in digraph searching. *Theor. Comput. Sci.*, 412(35):4688–4703, 2011.
- [19] D. Meister, J. A. Telle, and M. Vatshelle. Characterization and recognition of digraphs of bounded kelly-width. In *WG*, pages 270–279, 2007.
- [20] M. Middendorf and F. Pfeiffer. On the complexity of the disjoint paths problem. *Combinatorica*, 13(1):97–107, 1993.
- [21] B. A. Reed, N. Robertson, A. Schrijver, and P. D. Seymour. Finding disjoint trees in planar graphs in linear time. In *Graph Structure Theory*, pages 295–302, 1991.
- [22] N. Robertson and P. D. Seymour. Graph minors. V. Excluding a planar graph. *J. Comb. Theory, Ser. B*, 41(1):92–114, 1986.
- [23] N. Robertson and P. D. Seymour. Graph minors. VII. Disjoint paths on a surface. *J. Comb. Theory, Ser. B*, 45(2):212–254, 1988.
- [24] N. Robertson and P. D. Seymour. Graph minors. XIII. The disjoint paths problem. *J. Combin. Theory Ser. B*, 63(1):65–110, 1995.
- [25] N. Robertson, P. D. Seymour, and R. Thomas. Quickly excluding a planar graph. *J. Comb. Theory, Ser. B*, 62(2):323–348, 1994.
- [26] M. A. Safari. D-width: A more natural measure for directed tree width. In *MFCS*, pages 745–756, 2005.
- [27] A. Schrijver. Free partially commutative groups, cohomology, and paths and circuits in directed graphs on surfaces. Preprint. Available at <http://homepages.cwi.nl/~lex/files/dppg.pdf>.

- [28] A. Schrijver. Finding k disjoint paths in a directed planar graph. *SIAM J. Comput.*, 23(4):780–788, 1994.
- [29] A. Slivkins. Parameterized tractability of edge-disjoint paths on directed acyclic graphs. *SIAM J. Discrete Math.*, 24(1):146–157, 2010.