

Planar Disjoint-Paths Completion*

Isolde Adler[†] Stavros G. Kolliopoulos[‡] Dimitrios M. Thilikos[§]

January 17, 2015

Abstract

We introduce PLANAR DISJOINT PATHS COMPLETION, a completion counterpart of the Disjoint Paths problem, and study its parameterized complexity. The problem can be stated as follows: given a, not necessarily connected, plane graph G , k pairs of terminals, and a face F of G , find a minimum-size set of edges, if one exists, to be added inside F so that the embedding remains planar and the pairs become connected by k disjoint paths in the augmented network. Our results are twofold: first, we give an upper bound on the number of necessary additional edges when a solution exists. This bound is a function of k , independent of the size of G . Second, we show that the problem is fixed-parameter tractable, in particular, it can be solved in time $f(k) \cdot n^2$.

Keywords: Completion Problems, Disjoint Paths, Planar Graphs.

1 Introduction

Suppose we are given a planar road network with n cities and a set of k pairs of them. An empty area of the network is specified and we wish to add a minimum-size set of intercity roads in that area so that the augmented network remains planar and the pairs are connected by k internally disjoint roads. In graph-theoretic terms, we are looking for a minimum-size edge-completion of a plane graph so that an infeasible instance of the DISJOINT PATHS problem becomes feasible while maintaining planarity. In this paper we give an algorithm that solves this problem in $f(k) \cdot n^2$ steps. Our algorithm uses a combinatorial lemma stating that, whenever such a solution exists, its size depends exclusively on k .

The renowned DISJOINT PATHS PROBLEM (DP) is defined as follows.

*The second and the third author were co-financed by the E.U. (European Social Fund - ESF) and Greek national funds through the Operational Program “Education and Lifelong Learning” of the National Strategic Reference Framework (NSRF) - Research Funding Program: “Thales. Investing in knowledge society through the European Social Fund”.

[†]Institut für Informatik, Goethe-Universität, Frankfurt, Germany. Email: iadler@informatik.uni-frankfurt.de

[‡]Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Athens, Greece. Email: sgk@di.uoa.gr

[§]Department of Mathematics, University of Athens, Athens, Greece, and AIGCo project team, CNRS, LIRMM, France. Email: sedthilk@thilikos.info

$\text{DP}(G, s_1, t_1, \dots, s_k, t_k)$

Input: An undirected graph G and k pairs of *terminals* $s_1, t_1, \dots, s_k, t_k \in V(G)$.

Question: are there k pairwise internally vertex-disjoint paths Q_1, \dots, Q_k in G such that path Q_i connects s_i to t_i ?

(By *pairwise internally vertex-disjoint* we mean that two paths can only intersect at a vertex which is a terminal for both.)

DP is NP-complete even on planar graphs [14] but, when parameterized by k , the problem belongs to the parameterized complexity class FPT, i.e., it can be solved in time $f(k) \cdot n^{O(1)}$, for some function f . More precisely, it can be solved in $f(k) \cdot n^3$ time by the celebrated algorithm of Robertson and Seymour [19] from the Graph Minors project. For planar graphs, the same problem can be solved in $f(k) \cdot n$ [16].

We introduce a completion counterpart of this problem, PLANAR DISJOINT PATHS COMPLETION (PDPC), which is of interest on infeasible instances of DP, and we study its parameterized complexity, when parameterized by k . We are given an embedding of a, possibly disconnected, planar graph G in the sphere, k pairs of terminals $s_1, t_1, \dots, s_k, t_k \in V(G)$, a positive integer ℓ , and an open connected subset \mathbf{F} of the surface of the sphere, such that \mathbf{F} and G do not intersect (we stress that the boundary of \mathbf{F} is not necessarily a cycle). We want to determine whether there is a set of at most ℓ edges to add, the so-called *patch*, so that

- (i) the new edges lie inside \mathbf{F} and are incident only to vertices of G on the boundary of \mathbf{F} ,
- (ii) the new edges do not cross with each other or with G , and
- (iii) in the resulting graph, which consists of G plus the patch, DP has a solution.

PDPC is NP-complete even when ℓ is not a part of the input and G is planar by the following simple reduction from DP: add a triangle T to G and let \mathbf{F} be the interior of T . That way, we force the set of additional edges to be empty and obtain DP as a special case.

Notice that our problem is polynomially equivalent to the minimization problem where we ask for a minimum-size patch: simply solve the problem for all possible values of ℓ . Requiring the size of the patch to be at most ℓ is the primary source of difficulty. In case there is no restriction on the size of the patch and we simply ask whether one exists, the problem is in FPT by a reduction to DP, which is summarized as follows. For simplicity, let \mathbf{F} be an open disk. Let G' be the graph obtained by “sewing” along the boundary of \mathbf{F} an $O(n) \times O(n)$ -grid. By standard arguments, PDPC has a solution on G if and only if DP has a solution on G' . A similar, but more involved, construction applies when \mathbf{F} is not an open disk.

Parameterizing completion problems. Completion problems are natural to define: take any graph property, represented by a collection of graphs \mathcal{P} , and ask whether it is possible to add edges to a graph so that the new graph is in \mathcal{P} . Such problems have been studied for a long time and some of the most prominent are the following: HAMILTONIAN COMPLETION [8, GT34], PATH GRAPH COMPLETION [8, GT36] PROPER INTERVAL GRAPH COMPLETION [9] MINIMUM FILL-IN [20] INTERVAL GRAPH COMPLETION [8, GT35].

Kaplan et al. in their seminal paper [12] initiated the study of the parameterized complexity of completion problems and showed that MINIMUM FILL-IN, PROPER INTERVAL GRAPH COMPLETION and STRONGLY CHORDAL GRAPH COMPLETION are in FPT when parameterized by the number of edges to add. Recently, the problem left open by [12], namely INTERVAL GRAPH COMPLETION was also shown to be in FPT [11]. Certainly, for all these problems the testing of the corresponding property is in P, while for problems such as HAMILTONIAN COMPLETION, where \mathcal{P} is the class of Hamiltonian graphs, there is no FPT algorithm, unless P=NP. For the same reason, one cannot expect an FPT-algorithm when \mathcal{P} contains all YES-instances of DP, even on planar graphs. We consider an alternative way to parameterize completion problems, which is appropriate for the hard case, i.e., when testing \mathcal{P} is intractable: we parameterize the property itself. In this paper, we initiate this line of research, by considering the parameterized property \mathcal{P}_k that contains all YES-instances of DP on planar graphs with k pairs of terminals.

Basic concepts. As open sets are not discrete structures, we introduce some formalism that will allow us to move seamlessly from topological to combinatorial arguments. The definitions may look involved at first reading, but this is warranted if one considers, as we do, the problem in its full generality where the input graph is not necessarily connected.

Let G be a graph embedded in the sphere Σ_0 . Embeddings are always without crossings and we view the graph G as a subset of Σ_0 . Given a set $\mathbf{X} \subseteq \Sigma_0$, let $\mathbf{clos}(\mathbf{X})$, $\mathbf{int}(\mathbf{X})$, and $\partial\mathbf{X}$ denote the closure, the interior, and the boundary of X , respectively. We define $V(\mathbf{X}) = V(G) \cap \partial\mathbf{X}$. A *noose* is a Jordan curve of Σ_0 that meets G only on vertices. Let \mathcal{D} be a finite collection of mutually non-intersecting open disks of Σ_0 whose boundaries are nooses and such that each point that belongs to at least two such nooses is a vertex of G . We define $\mathbf{I}_{\mathcal{D}} = \bigcup_{D \in \mathcal{D}} D$ and define $\Gamma_{\mathcal{D}}$ as the Σ_0 -embedded graph whose vertex set is $V(\mathbf{I}_{\mathcal{D}})$ and whose edge set consists of the connected components of the set $\partial\mathbf{I}_{\mathcal{D}} \setminus V(\mathbf{I}_{\mathcal{D}})$. Notice that, in the definition of $\Gamma_{\mathcal{D}}$, we permit multiple edges, loops, or vertex-less edges.

Let \mathbf{J} be an open subset of Σ_0 . \mathbf{J} is a *cactus set* of G if there is a collection \mathcal{D} as above such that $\mathbf{J} = \mathbf{I}_{\mathcal{D}}$, all biconnected components of the graph $\Gamma_{\mathcal{D}}$ are cycles, $V(G) \subseteq \mathbf{clos}(\mathbf{J})$, and each edge of $E(G)$ is a subset of \mathbf{J} . Given such a \mathbf{J} , we define $\Gamma_{\mathbf{J}} = \Gamma_{\mathcal{D}}$. Two cactus sets \mathbf{J} and \mathbf{J}' of G are *isomorphic* if $\Gamma_{\mathbf{J}}$ and $\Gamma_{\mathbf{J}'}$ are topologically isomorphic. Throughout this paper, we use the standard notion of topological isomorphism between planar embeddings, see Section 2.

Given a cactus set \mathbf{J} , we define for each vertex $v \in V(\mathbf{J})$ its *multiplicity* $\mu(v)$ to be equal to the number of connected components of the set $\Gamma_{\mathbf{J}} \setminus \{v\}$, minus the number of connected components of $\Gamma_{\mathbf{J}}$, plus one. We also define $\mu(\mathbf{J}) = \sum_{v \in V(\mathbf{J})} \mu(v)$. Observe that, given a cactus set \mathbf{J} of G , the edges of G lie entirely within the interior of \mathbf{J} . See Figure 1. The boundary of \mathbf{J} corresponds to a collection of simple closed curves such that

- (i) no two of them intersect at more than one point and
- (ii) they intersect with G only at (some of) the vertices in $V(G)$.




Figure 1: A cactus set \mathbf{J} and the vertices of $V(\mathbf{J})$. Next to each vertex v we give its multiplicity $\mu(v)$.

Cactus sets are useful throughout our paper as “capsule” structures that surround G and thus they abstract the interface of a graph embedding with the rest of the sphere surface.

We say that an open set \mathbf{F} of Σ_0 is an *outer-cactus set* of G if $\Sigma_0 \setminus \mathbf{clos}(\mathbf{F})$ is a cactus set of G . See Fig. 2.(ii). For example, if G is planar, any face F of G can be used to define an outer-cactus set, whose boundary meets G only at the vertices incident to F . Our definition of an outer-cactus set is more general: it can be a subset of a face F , meeting the boundary of F only at some of its vertices.

Let G be an input graph to DP, see Figure 2.(i). Given an outer-cactus set \mathbf{F} of G , an \mathbf{F} -patch of G is a pair (P, \mathbf{J}) where (i) \mathbf{J} is a cactus set of G , where $\Sigma_0 \setminus \mathbf{clos}(\mathbf{J}) \subseteq \mathbf{F}$ and (ii) P is a graph embedded in Σ_0 *without crossings* such that $E(P) \subseteq \Sigma_0 \setminus \mathbf{clos}(\mathbf{J})$, $V(P) = V(\mathbf{J})$ (see Figures 2.(iii) and 2.(iv)). Observe that the edges of P do not cross any edge in $E(G)$. In the definition of the \mathbf{F} -patch, the graph P contains the new edges we add. The vertices in $V(\mathbf{F})$ define the vertices of G which we are allowed to include in P . $V(\mathbf{J})$ is meant to contain those vertices of $V(\mathbf{F})$ that become vertices (possibly isolated) of P . In terms of data structures, we assume that a cactus set \mathbf{J} is represented by the (embedded) graph $\Gamma_{\mathbf{J}}$. Similarly an outer-cactus set \mathbf{F} is represented by the (embedded) graph $\Gamma_{\Sigma_0 \setminus \mathbf{clos}(\mathbf{F})}$.

We restate now the definition of the PLANAR DISJOINT PATHS COMPLETION problem as follows:

PDPC($G, s_1, t_1, \dots, s_k, t_k, \ell, \mathbf{F}$)

Input: A graph G embedded in Σ_0 without crossings, terminals $s_1, t_1, \dots, s_k, t_k \in V(G)$, a positive integer ℓ , and an outer-cactus set \mathbf{F} of G .

Parameter: k

Question: Is there an \mathbf{F} -patch (P, \mathbf{J}) of G , such that $|P| \leq \ell$ and $\text{DP}(G \cup P, s_1, t_1, \dots, s_k, t_k)$ has a solution? Compute such an \mathbf{F} -patch if it exists.




Figure 2: An example input of the PDPC problem and a solution to it when $\ell = 8$: (i) The graph embedding in the input and the terminals $s_1, t_1, s_2, t_2, s_3, t_3$. The closure of the grey area contains the graph G and the big vertices are the terminals. The white area is a face of G . (ii) The input of the problem, consisting of G , the terminals and the outer-cactus set \mathbf{F} . The solid black vertices are the vertices of G that are also vertices of $V(\mathbf{F})$. (iii) The solution of the problem consists of the \mathbf{F} -patch (P, \mathbf{J}) where the edges of P are the dashed lines and $\mathbf{J} = \mathbf{J}_1 \cup \mathbf{J}_2 \cup \mathbf{J}_3$. (iv) The input and the solution together where the validity of the patch is certified by 3 disjoint paths.

If such an \mathbf{F} -patch exists, we call it a *solution* for PDPC. In the corresponding optimization problem, denoted by MIN-PDPC, one asks for the minimum ℓ for which PDPC has a solution, if one exists. See Fig. 2 for an example input of PDPC and a solution to it.

Our results. Notice that in the definition of PDPC the size of the patch does not depend on the parameter k . Thus, it is not even obvious that PDPC belongs to the parameterized complexity class \mathbf{XP} , i.e., it has an algorithm of time $n^{f(k)}$ for some function f . Our first contribution, Theorem 2, is a combinatorial one: we prove that if a patch exists, then its size is bounded by k^{2^k} . Therefore, we can always assume that ℓ is bounded by a function of k . This bound is the departure point for the proof of the main algorithmic result of this paper:

Theorem 1 PDPC \in FPT. In particular, PDPC can be solved in $f(k) \cdot n^2$ steps, where f is a function that depends only on k . Therefore, MIN-PDPC can be solved in $g(k) \cdot n^2$ steps.

We present now the proof strategy and the ideas underlying our results.

1.1 Proof strategy

Combinatorial Theorem. In Theorem 2, we prove that every patch whose size is larger than k^{2^k} , can be replaced by another one of strictly smaller size. In particular, we identify a region \mathbf{B} of \mathbf{F} that is traversed by a large number of segments of different paths of the DP solution. Within that region, we apply a global topological transformation that replaces the old patch by a new, strictly smaller one, while preserving its embeddability in \mathbf{F} . The planarity of the new patch is based on the fact that the new segments are reflections in \mathbf{B} of a set of segments of the feasible DP solution that previously lied outside \mathbf{B} . This combinatorial result allows us to reduce the search space of the problem to one whose size is bounded by $\min\{\ell, k^{2^k}\}$. Therefore, the construction of the corresponding collection of “candidate solutions” can be done in advance, for each given k , without requiring *any a priori knowledge* of the input graph G .

We note that the proof of our combinatorial theorem could be of independent interest. In fact, it is one of the ingredients of the proof of the main result of [1].

The algorithm for PDPC. As the number of patches is bounded by a function of k , we need to determine whether there is a correct way to glue one of them on vertices of the boundary of the open set \mathbf{F} so that the resulting graph is a YES-instance of the DP problem. For each candidate patch \tilde{P} , together with its corresponding candidate cactus set $\tilde{\mathbf{J}}$, we define the *set of compatible* graphs embedded in $\tilde{\mathbf{J}}$. Each compatible graph \tilde{H} consists of unit-length paths and has the property that $\tilde{P} \cup \tilde{H}$ contains k disjoint paths. Intuitively, each \tilde{H} is a certificate of the part of the DP solution that lies within G when the patch in \mathbf{F} is isomorphic to \tilde{P} . It therefore remains to check for each \tilde{H} whether it can be realized by a collection of actual paths within G . For this, we set up a collection \mathcal{H} of all such certificates. Checking for a suitable realization of a member of \mathcal{H} in G is still a topological problem that depends on the embedding of G : graphs that are isomorphic, but not *topologically isomorphic*, may certify different completions. For this reason, our next step is to *enhance* the structure of the members of \mathcal{H} so that their realization in G reduces to a purely combinatorial check. (Cf. Section 4.1 for the definition of the enhancement operation). We show in Lemma 6 that for the enhanced certificates, this check can be implemented by rooted topological minor testing. For this check, we can apply the recent algorithm of [10] that runs in $h_1(k) \cdot n^3$ steps and obtain an algorithm of overall complexity $h_2(k) \cdot n^3$.

We note that the use of the complicated machinery of the algorithm in [10] can be bypassed towards obtaining a simpler and faster $f(k) \cdot n^2$ algorithm. This is possible because the generated instances of the rooted topological minor problem satisfy certain structural properties. This allows the direct application of the *Irrelevant Vertex Technique* introduced in [19] for solving, among others, the DISJOINT PATHS Problem. The

details of this improvement are in Section 5.

2 Preliminaries

We consider finite graphs. For a graph G we denote the vertex set by $V(G)$ and the edge set by $E(G)$. If G is embedded in the sphere Σ_0 , the edges of G and the graph G refer also to the corresponding sets of points in Σ_0 . Clearly the edges of G correspond to open sets and G itself is a closed set. We denote by $F(G)$ the set of all the faces of G , i.e., all connected components of $\Sigma_0 \setminus G$. Given a set $S \subseteq V(G)$, we say that the pair (G, S) is a graph *rooted* at S . We also denote as $\mathcal{P}(G)$ the set of all paths in G with at least one edge. Given a path $P \in \mathcal{P}(G)$, we denote by $I(P)$ the set of internal vertices of P . Given a vertex $v \in V(G)$ and a positive integer r , we denote by $N_G^r(v)$ the set of all vertices in G that are within distance at most r from v . Given a vertex v of a graph G with exactly two neighbors x and y , the result of the *dissolution* of v in G is the graph obtained if we remove v from G and add, if it does not already exist, the edge $\{x, y\}$.

Rooted topological minors. Let H and G be graphs, S_H be a subset of vertices in $V(H)$, S_G be a subset of vertices in $V(G)$, and ρ be a bijection from S_H to S_G . We say that (H, S_H) is a ρ -*rooted topological minor* of (G, S_G) , if there exist injections $\psi_0: V(H) \rightarrow V(G)$ and $\psi_1: E(H) \rightarrow \mathcal{P}(G)$ such that

1. $\rho \subseteq \psi_0$,
2. for every $e = \{x, y\} \in E(H)$, $\psi_1(e)$ is a $(\psi_0(x), \psi_0(y))$ -path in $\mathcal{P}(G)$, and
3. all $e_1, e_2 \in E(H)$ with $e_1 \neq e_2$ satisfy $I(\psi_1(e_1)) \cap V(\psi_1(e_2)) = \emptyset$.

In words, when H is a topological minor of G , G contains a subgraph which is isomorphic to a subdivision of H . In addition, when (H, S_H) is a ρ -*rooted topological minor* of (G, S_G) then this isomorphism respects the bijection ρ between the vertex sets S_H and S_G .

Contractions. Let G and H be graphs and let $\sigma: V(G) \rightarrow V(H)$ be a surjective mapping such that

1. for every vertex $v \in V(H)$, the graph $G[\sigma^{-1}(v)]$ is connected;
2. for every edge $\{v, u\} \in E(H)$, the graph $G[\sigma^{-1}(v) \cup \sigma^{-1}(u)]$ is connected;
3. for every $\{v, u\} \in E(G)$, either $\sigma(v) = \sigma(u)$, or $\{\sigma(v), \sigma(u)\} \in E(H)$.

We say that H is a σ -*contraction* of G or simply that H is a *contraction* of G if such a σ exists.

Observation 1 *Let H and G be graphs such that H is a σ -contraction of G . If $x, y \in V(G)$, then the distance in G between x and y is at least the distance in H between $\sigma(x)$ and $\sigma(y)$.*

We also need the following topological lemma.

Lemma 1 *Let G be a Σ_0 -embedded graph and let \mathbf{J} be a cactus set of it. Let also M be a Σ_0 -embedded graph such that $M \cap \mathbf{J} = \emptyset$ and $V(M) \subseteq V(\mathbf{J})$. Then there is a closed curve K in $\Sigma_0 \setminus \mathbf{clos}(\mathbf{J})$ meeting each edge of M twice.*

Proof. We consider the dual graph of $\Gamma_{\mathbf{J}} \cup M$ and we remove from it all vertices lying inside \mathbf{J} . We denote by Q the resulting graph and notice that Q is connected because the set $\Sigma_0 \setminus \mathbf{clos}(\mathbf{J})$ is connected. Next, construct the graph Q' by subdividing once every edge of Q such that the subdivision vertex is the intersection point of the edge and its dual. Let T be a spanning tree of Q' and let K be a closed curve such that T is inside one of the connected components of the set $\Sigma_0 \setminus K$ and \mathbf{J} in the other. If we further require K to intersect M a minimum number of times, we obtain the claimed curve. \square

Topological isomorphism. Given a graph G embedded in Σ_0 , let \mathbf{f} be a face in $F(G)$ whose boundary has ξ connected components A_1, \dots, A_ξ . We define the set $\pi(\mathbf{f}) = \{\pi_1, \dots, \pi_\xi\}$ such that each π_i is the cyclic ordering of $V(A_i)$, possible with repetitions, defined by the way vertices are met while walking along A_i in a way that the face \mathbf{f} is always on our left side. Clearly, repeated vertices in this walk are cut-vertices of G .

Let G and H be graphs embedded in Σ_0 . We say that G and H are *topologically isomorphic* if there exist bijections $\phi: V(G) \rightarrow V(H)$ and $\theta: F(G) \rightarrow F(H)$ such that

1. ϕ is an isomorphism from G to H , i.e. for every pair $\{x, y\}$ of distinct vertices in $V(G)$, $\{x, y\} \in E(G)$ iff $\{\phi(x), \phi(y)\} \in E(H)$.
2. For every face $\mathbf{f} \in F(G)$, $\phi(\pi(\mathbf{f})) = \pi(\theta(\mathbf{f}))$.

In the definition above, by $\phi(\pi(\mathbf{f}))$ we mean $\{\phi(\pi_1), \dots, \phi(\pi_\xi)\}$, where, if $\pi_i = (x_1, \dots, x_{\zeta_i}, x_1)$, then by $\phi(\pi_i)$, we mean $(\phi(x_1), \dots, \phi(x_{\zeta_i}), \phi(x_1))$. Notice that it is possible for two isomorphic planar graphs to have embeddings that are not topologically isomorphic (see [6, page 93] for such an example and further discussion on this topic).

Treewidth. A *tree decomposition* of a graph G is a pair (\mathcal{X}, T) where T is a tree with nodes $\{1, \dots, m\}$ and $\mathcal{X} = \{X_i \mid i \in V(T)\}$ is a collection of subsets of $V(G)$ (called *bags*) such that:

1. $\bigcup_{i \in V(T)} X_i = V(G)$,
2. for each edge $\{x, y\} \in E(G)$, $\{x, y\} \subseteq X_i$ for some $i \in V(T)$, and
3. for each $x \in V(G)$ the set $\{i \mid x \in X_i\}$ induces a connected subtree of T .

The *width* of a tree decomposition $(\{X_i \mid i \in V(T)\}, T)$ is $\max_{i \in V(T)} \{|X_i| - 1\}$. The *treewidth* of a graph G denoted $\mathbf{tw}(G)$ is the minimum width over all tree decompositions of G .

3 Bounding the size of the completion

In this section we show the following.

Theorem 2 *If there is a solution for PDPC($G, s_1, t_1, \dots, s_k, t_k, \ell, \mathbf{F}$), then there is a solution (P, \mathbf{J}) with $|E(P)| \leq k^{2^k}$.*

For the proof, we need the following combinatorial lemma.

Lemma 2 *Let Σ be an alphabet of size $|\Sigma| = k$. Let $w \in \Sigma^*$ be a string over Σ . If $|w| > 2^k$, then w contains an infix y with $|y| \geq 2$, such that every letter occurring in y , occurs an even number of times in y .*

Proof. Let $\Sigma = \{a_1, \dots, a_k\}$, and let $w = w_1 \cdots w_n$ with $n > 2^k$. Define vectors $z_i \in \{0, 1\}^k$ for $i \in \{1, \dots, n\}$, and we let the j th entry of vector z_i be 0 if and only if letter a_j occurs an even number of times in the prefix $w_1 \cdots w_i$ of w and 1 otherwise. Since $n > 2^k$, there exist $i, i' \in \{1, \dots, n\}$ with $i \neq i'$, such that $z_i = z_{i'}$. Then $y = w_{i+1} \cdots w_{i'}$ proves the lemma. \square

We also need the following easy topological lemma.

Lemma 3 *Let G be a connected outerplanar graph that may have loops but no multiple edges (or loops) and let Γ be an embedding of G in Σ_0 such that all its vertices belong to the boundary of a face f . Let also E_{in} be the set of all the edges or loops of G that do not belong to the boundary of f . Then there is a unique (up to topological isomorphism) embedding Γ' of G such that $\Gamma \cap \Gamma'$ is the boundary of f and that all edges of E_{in} are embedded inside f .*

Proof. Let $\tilde{\Gamma}$ be the boundary of f . Let \mathcal{F} be the set of all connected components of $\Sigma_0 \setminus \tilde{\Gamma}$, except from f .

For each $e \in E_{\text{in}}$ we denote by S_e the set of endpoints of e . For each edge $e \in E_{\text{in}}$ we define a set \bar{e} that will be the image of e in the new embedding Γ' . We distinguish the following two cases.

Case 1. $|S_e| = 2$. Notice that S_e is a separator of G . Notice also that $G \setminus S_e$ contains two connected components. Also there exists a (unique) cycle C_e containing the edge e and no other points of Γ such that each of these connected components belong each to a distinct disk among the two disks that are bounded by C_e . We define $\bar{e} = C_e \setminus \text{clos}(e)$.

Case 2. $|S_e| = 1$. Notice that there exists a unique cycle C_e defining two disks D and D' such that (i) $C_e \cap \Gamma = S_e$ and (ii) $\Gamma_e \subseteq D$ where Γ_e is the (unique) member of \mathcal{F} that contains the loop e . We define $\bar{e} = C_e \setminus S_e$.

Notice that it is possible to define all \bar{e} 's, one by one, such that they do not intersect. Given the above definitions, it follows that $\tilde{\Gamma} \cup \bigcup_{e \in E_{\text{in}}} \bar{e}$ is the required embedding. \square

Proof of Theorem 2. Let (P, \mathbf{J}) be a solution for PDPC($G, s_1, t_1, \dots, s_k, t_k, \ell, \mathbf{F}$) with $|E(P)|$ minimal. Consider the embedding of $G \cup P$ in the sphere Σ_0 , and let Q_1, \dots, Q_k be the paths of a DP solution in $G \cup P$. By the minimality of $|E(P)|$ we can assume that the edges of P are exactly the edges of $\bigcup_{i \in \{1, \dots, k\}} Q_i$ that are not in G . For the same reason, two edges in P have a common endpoint x that is not a terminal *only if* x is a cut-vertex of $\Gamma_{\mathbf{J}}$.

Let P^* denote the graph obtained from the dual of $P \cup \Gamma_{\mathbf{J}}$, after removing the vertices corresponding to the faces of $\Gamma_{\mathbf{J}}$ that are disjoint from \mathbf{F} . We first show that the maximum degree of P^* is bounded by k . Assume, to the contrary, that P^* has a vertex incident to two edges e_1^* and e_2^* such that e_1 and e_2 belong to the same path, say

Q_x , in $\{Q_1, \dots, Q_k\}$. Then it is possible to choose an endpoint p_1 of e_1 and an endpoint p_2 of e_2 such that, (P', \mathbf{J}') is also a solution such that $|E(P')| < |E(P)|$. Indeed, p_1 and p_2 are the vertices incident to e_1 and e_2 belonging to the two connected components of $Q_x \setminus \{e_1, e_2\}$ that contain respectively s_x and t_x . Therefore we established that for every $i \in [k]$, no vertex of P^* is incident to two distinct edges of Q_i , thus all vertices of P^* have degree at most k .

Here $P' = P \setminus \{e_1, e_2\} \cup \{p_1, p_2\}$ and \mathbf{J}' is defined such that $\Gamma_{\mathbf{J}'}$ is obtained from $\Gamma_{\mathbf{J}}$ after dissolving the vertices that became isolated during the construction of P' .

Our next aim is to prove that the diameter of P^* is bounded by 2^k . Then $|E(P^*)| = |E(P)| \leq k^{2^k}$ and we are done. Note that every edge in $E(P^*)$ corresponds to an edge in exactly one path of Q_1, \dots, Q_k . Hence, given a path $R = (r_0, \dots, r_\zeta)$ in P^* , it corresponds to a string $w \in \{Q_1, \dots, Q_k\}^*$ in a natural way. It is enough to prove the following claim.

Claim. *The string w contains no infix y with $|y| \geq 2$, such that every letter occurring in y occurs an even number of times in y .*

Proof of the Claim. Towards a contradiction, suppose that w contains such an infix y . We may assume that $w = y$. Let $E_R \subseteq E(P)$ be the set of edges corresponding to the edges of the path $R \subseteq P^*$. Then $|E_R| \geq 2$ because w (and hence R) has length at least 2. Let $\mathbf{B} \subseteq \Sigma_0$ be the open set defined by the union of all edges in E_R and all faces of the graph $P \cup \Gamma_{\mathbf{J}}$ that are incident to them. Clearly, \mathbf{B} is an open connected subset of \mathbf{F} with the following properties:

- (a) \mathbf{B} contains all edges in E_R and no other edges of P ,
- (b) the ends of every edge in E_R lie on the boundary $\partial\mathbf{B}$, and
- (c) every edge in $E(P) \setminus E_R$ has empty intersection with \mathbf{B} .

We consider an ‘up-and-down’ partition $(U = \{u_1, \dots, u_r\}, D = \{d_1, \dots, d_r\})$ of the endpoints of the edges in E_R as follows: traverse the path R in P^* in some arbitrary direction and when the i th edge $e_i \in E_R$ is met, the endpoint u_i of e_i on the left of this direction is added to U and the right endpoint d_i is added to D . Notice that U and D may be multisets because it is not necessary that all vertices in P have degree 1. For each $i \in \{1, \dots, r\}$ we say that u_i is the *counterpart* of d_i and vice versa.

Because the paths Q_1, \dots, Q_k are vertex-disjoint, the following holds.

Observation: If $x \in V(P)$ has degree larger than one, then either x is a terminal and has degree at most k or x is a cutpoint of $\Gamma_{\mathbf{J}}$ and has degree exactly 2.

By assumption, every path Q_i crosses R an even, say $2n_i$, number of times. Now for every path Q_i satisfying $E(Q_i) \cap E_R \neq \emptyset$, we number the edges in $E(Q_i) \cap E_R$ by $e_1^i, \dots, e_{2n_i}^i$ in the order of their appearance when traversing Q_i from s_i to t_i and we orient them from s_i to t_i . We introduce shortcuts for Q_i as follows: for every odd number $j \in \{1, \dots, 2n_i\}$, we replace the subpath of Q_i from $\text{tail}(e_j^i)$ to $\text{head}(e_{j+1}^i)$ by a new edge f_j^i in D (see Figure 3).




Figure 3: Example of the transformation in the proof of the Claim in the proof of Theorem 2; P is on the left and P' is shown on the right. The dashed lines represent the edges of C .

After having done this for all odd numbers $j \in \{1, \dots, 2n_i\}$, we obtain a new path Q'_i from s_i to t_i that uses strictly less edges in \mathbf{B} than Q_i . Having replaced all paths Q_i with $E(Q_i) \cap E_R \neq \emptyset$ in this way by a new path Q'_i , we obtain from P a new graph P' by replacing every pair of edges $e_j^i, e_{j+1}^i \in E(P)$ by f_j^i for all $i \in \{1, \dots, k\}$ with $E(Q_i) \cap E_R \neq \emptyset$, and for all $j \in \{1, \dots, 2n_i\}$, j odd. We denote by E'_R the set of all replacement edges f_j^i . We also remove every vertex that becomes isolated in P' during this operation.

Then it is easy to verify that:

- None of the edges of E_R survives in $E(P')$.
- $|E(P')| < |E(P)|$.
- $\text{DP}(G \cup P', s_1, t_1, \dots, s_k, t_k)$ has a solution.

If we show that, for some suitable cactus set \mathbf{J}' of G , (P', \mathbf{J}') is an \mathbf{F} -patch, then we are done, because $|E(P')| < |E(P)|$. In what follows, we prove that P' can also be embedded without crossings in $\mathbf{clos}(\mathbf{F})$ such that $E(P') \subseteq \Sigma_0 \setminus \partial\mathbf{F}$. For this it suffices to prove that the edges in E'_R can be embedded in \mathbf{B} without crossings.

For every path Q_i with $E(Q_i) \cap E_R \neq \emptyset$ let F_j^i denote the subpath of Q_i from $\text{head}(e_j^i)$ to $\text{tail}(e_{j+1}^i)$, for $j \in \{1, \dots, 2n_i\}$, j odd (this path may be edgeless only in the case where $\text{head}(e_j^i) = \text{tail}(e_{j+1}^i)$ is a cut-vertex of $\Gamma_{\mathbf{J}}$). We replace F_j^i by a single edge c_j^i (when the corresponding path is edgeless, the edge c_j^i is a loop outside B). We consider the graph C with vertex set $V(P)$ and edge set $\{c_j^i \mid i \in \{1, \dots, k\}, E(Q_i) \cap E_R \neq \emptyset, j \in \{1, \dots, 2n_i\}, j \text{ odd}\}$.

Our strategy consists of a two-step transformation of this embedding. The first step creates an embedding of C inside $\mathbf{clos}(\mathbf{B})$ without moving the vertices. Indeed, notice that C is embedded in $\Sigma_0 \setminus \mathbf{B}$ without crossings such that all the endpoints of the edges in C lie on the boundary of \mathbf{B} . As for every odd j in $[2n_i]$, $j+1$ also belongs to $[2n_i]$, it follows that none of the endpoints of F_j^i can be a terminal. This, together with the observation above, implies that no two edges of C have a common endpoint.

By applying Lemma 3 on the Σ_0 -embedded outerplanar graph $\Gamma = \Gamma_{\mathbf{J}} \cup C$ where \mathbf{B} plays the role of the face f , we obtain a new non-crossing embedding Γ' . Notice that $\Gamma' \setminus \Gamma_{\mathbf{J}}$ is a new non-crossing embedding of C where all of its edges lie inside \mathbf{B} . (Recall that none of the edges of E_R is present in this embedding.) This transformation maps every edge c_j^i to a new edge inside \mathbf{B} with the same endpoints.

The second step “reflects” the resulting embedding along the axis defined by the path R such that each vertex is exchanged with its counterpart. Now define $(c_j^i)'$ so that it connects $\text{tail}(e_j^i)$ and $\text{head}(e_{j+1}^i)$ – these are exactly the counterparts of $\text{head}(e_j^i)$ and $\text{tail}(e_{j+1}^i)$. Due to symmetry, the $(c_j^i)'$ are pairwise non-crossing, and none of them crosses a drawing of an edge in $E(P) \setminus E_R$. Hence the $(c_j^i)'$ together with the drawing of edges in $E(P) \setminus E_R$ provide a planar drawing of P' (where $(c_j^i)'$ is the drawing of f_j^i). We finally define (up to isomorphism) the cactus set \mathbf{J}' of G such that $\Gamma_{\mathbf{J}'}$ is obtained from $\Gamma_{\mathbf{J}}$ after dissolving the vertices of $P \cap \Gamma_{\mathbf{J}}$ that are isolated in P' . It is easy to verify that (P', \mathbf{J}') is an \mathbf{F} -patch of G . This concludes the proof of the Claim.

Using the Claim above and from Lemma 2, it follows that $|w| \leq 2^k$, and hence the diameter of P^* is bounded by 2^k . The proof of Theorem 2 is now complete. \square

Let \mathcal{L} be a list of all simple plane graphs with at most $\min\{\ell, k^{2^k}\}$ edges and no isolated vertices. We call a graph in \mathcal{L} a *completion*. As a first step, our algorithm for PDPC computes the list \mathcal{L} . Obviously, the running time of this process is bounded by a function depending only on k .

4 The algorithm for PLANAR-DPC

The fact that the size of \mathcal{L} is bounded by a function of k implies that PDPC is in XP. Indeed, given the list \mathcal{L} , for each completion $\tilde{P} \in \mathcal{L}$ we define the graph $Q_{\tilde{P}} = (V(\tilde{P}), \emptyset)$ and we consider all cactus sets $\tilde{\mathbf{J}}$ of $Q_{\tilde{P}}$ where $(\tilde{P}, \tilde{\mathbf{J}})$ is a $(\Sigma_0 \setminus \mathbf{clos}(\tilde{\mathbf{J}}))$ -patch of $Q_{\tilde{P}}$ and $V(\tilde{\mathbf{J}}) = V(\tilde{P})$. We denote the set of all such pairs $(\tilde{P}, \tilde{\mathbf{J}})$ by \mathcal{J} and observe that the number of its elements (up to topological isomorphism of the graph $\tilde{P} \cup \Gamma_{\tilde{\mathbf{J}}}$) is bounded by a function of k .

For each pair $(\tilde{P}, \tilde{\mathbf{J}}) \in \mathcal{J}$, we check whether there exists an \mathbf{F} -patch (P, \mathbf{J}) of G such that $\tilde{P} \cup \Gamma_{\tilde{\mathbf{J}}}$ and $P \cup \Gamma_{\mathbf{J}}$ are topologically isomorphic and DP has a solution in the graph $G \cup P$. As there are $n^{z(k)}$ ways to choose (P, \mathbf{J}) and each check can be done in $O(z_1(k) \cdot n^3)$ steps, we conclude that PDPC can be solved in $n^{z_2(k)}$ steps. In the remainder of the paper, we will prove that the problem is actually in FPT.

The main bottleneck is that there are too many ways to identify $V(\tilde{\mathbf{J}})$ with vertices of $V(\mathbf{F})$, because we cannot bound $|V(\mathbf{F})|$ by a function of k . To overcome this, we associate with a positive instance of PDPC a rooted topological minor (\tilde{H}, \tilde{T}) of the original graph G , that witnesses the fact that $(\tilde{P}, \tilde{\mathbf{J}})$ corresponds to the desired \mathbf{F} -patch of G . For convenience, we assume from now on that $\Gamma_{\mathbf{int}(\Sigma_0 \setminus \mathbf{F})}$ does not contain any loops or multiple edges. In other words, none of the open disks corresponding to the connected components of $\Sigma_0 \setminus \mathbf{F}$ contains less than 3 vertices in its boundary. If the input instance violates this property we may enforce it by adding isolated vertices in G




Figure 4: In the leftmost image the dotted lines are the edges of \tilde{P} . Together with the interior of the grey areas they form a pair $(\tilde{P}, \tilde{\mathbf{J}}) \in \mathcal{J}$. $\tilde{\mathbf{J}}$ has two weakly connected components. $V(\tilde{\mathbf{J}})$ consists of the 12 vertices on the boundary of the grey areas. The solid lines that intersect the open set $\tilde{\mathbf{J}}$ are the edges of the graph \tilde{H} , which is compatible with $(\tilde{P}, \tilde{\mathbf{J}})$. Let $\tilde{T} = \{s_1, t_1, \dots, s_4, t_4\}$. The triple $\tau = (\tilde{\mathbf{J}}, \tilde{H}, \tilde{T}) \in \mathcal{H}$ has two parts. The middle and the leftmost pictures show how each of these parts is enhanced in order to construct the graphs R_τ^1 and R_τ^2 .

and modify \mathbf{F} such that the new vertices belong to its boundary. Similarly, we restrict \mathcal{J} to contain only pairs $(\tilde{P}, \tilde{\mathbf{J}})$ where $\Gamma_{\tilde{\mathbf{J}}}$ does not contain any loops or multiple edges.

Given a pair $(\tilde{P}, \tilde{\mathbf{J}}) \in \mathcal{J}$, we say that a rooted simple graph $(\tilde{H}, \tilde{T} = \{a_1, b_1, \dots, a_k, b_k\})$ embedded in Σ_0 , is *compatible* with $(\tilde{P}, \tilde{\mathbf{J}})$ when

1. for every $e \in E(\tilde{H})$, $e \subseteq \tilde{\mathbf{J}}$,
2. the vertices of \tilde{H} that are also vertices of $\Gamma_{\tilde{\mathbf{J}}}$ are $|V(\tilde{P})|$, while the vertices of \tilde{H} that are not vertices of $\Gamma_{\tilde{\mathbf{J}}}$ are all vertices of \tilde{T} ,
3. $V(\tilde{H}) \setminus \tilde{T} \subseteq V(\tilde{\mathbf{J}}) \subseteq V(\tilde{H})$,
4. $\text{DP}(\tilde{P} \cup \tilde{H}, a_1, b_1, \dots, a_k, b_k)$ has a solution.

We define

$$\mathcal{H} = \{(\tilde{\mathbf{J}}, \tilde{H}, \tilde{T}) \mid \text{there exists a } (\tilde{P}, \tilde{\mathbf{J}}) \in \mathcal{J} \text{ such that } (\tilde{H}, \tilde{T}) \text{ is compatible with } (\tilde{P}, \tilde{\mathbf{J}})\}$$

and notice that $|\mathcal{H}|$ is bounded by some function of k . See the leftmost part of Fig. 4 for an example of a triple in \mathcal{H} .

Assuming that (P, \mathbf{J}) is a solution for $\text{PDPC}(G, s_1, t_1, \dots, s_k, t_k, \ell, \mathbf{F})$, consider the parts of the corresponding disjoint paths that lie within G . The intuition behind the definition above is that \tilde{H} is a certificate of these “partial paths” in G . Clearly, the number of these certificates is bounded by $|\mathcal{H}|$ and they can be enumerated in $f_0(k)$ steps, for some suitable function f_0 . For example, for the solution depicted in Fig. 4, \tilde{H} consists of 7 disjoint edges, one for each subpath within G . Our task is to find an

FPT-algorithm that for every such certificate checks whether the corresponding partial paths exist in G .

Given an open set \mathbf{O} , a *weakly connected component* of \mathbf{O} is the interior of some connected component of the set $\mathbf{clos}(\mathbf{O})$. Notice that a weakly connected component is not necessarily a connected set.

Let $\bar{\mathbf{F}}^1, \dots, \bar{\mathbf{F}}^\lambda$ be the weakly connected components of the set $\Sigma_0 \setminus \mathbf{clos}(\mathbf{F})$. We call such a component $\bar{\mathbf{F}}^i$ *active* if $\mathbf{clos}(\bar{\mathbf{F}}^i) \cap T \neq \emptyset$. We denote the collection of all active components by $\mathcal{F}_{\mathbf{F}}$. A crucial observation is that if an \mathbf{F} -patch exists we can always replace it by one that bypasses the inactive components.

Lemma 4 *Let $(G, s_1, t_1, \dots, s_k, t_k, \mathbf{F})$ be an instance for the PDPC problem and let $G' = G[\bigcup_{\bar{\mathbf{F}}^i \in \mathcal{F}_{\mathbf{F}}} \mathbf{clos}(\bar{\mathbf{F}}^i) \cap V(G)]$ and $\mathbf{F}' = \Sigma_0 \setminus \bigcup_{\bar{\mathbf{F}}^i \in \mathcal{F}_{\mathbf{F}}} \mathbf{clos}(\bar{\mathbf{F}}^i)$. Then $(G', s_1, t_1, \dots, s_k, t_k, \mathbf{F}')$ is an equivalent instance.*

Proof. For the non-trivial direction, we assume that G has an \mathbf{F} -patch (P, \mathbf{J}) such that $G \cup P$ contains a collection $\{P_1, \dots, P_k\}$ paths that certify the feasibility of $\text{DP}(G \cup P, s_1, t_1, \dots, s_k, t_k)$. Let P' be the graph obtained if in $(\bigcup_{i=1, \dots, k} P_i) \cap \mathbf{clos}(\mathbf{F}')$ we dissolve all vertices that are in \mathbf{F}' . Let also \mathbf{J}' be the union of all weakly connected components of the set \mathbf{J} that contain some open set from $\mathcal{F}_{\mathbf{F}}$. Observe that (P', \mathbf{J}') is an \mathbf{F}' -patch of G' . \square

By Lemma 4, we can assume from now on that the number λ of the weakly connected components of the set $\Sigma_0 \setminus \mathbf{clos}(\mathbf{F})$ is at most $2k$. Also we *restrict* \mathcal{H} so that it contains only triples $(\tilde{\mathbf{J}}, \tilde{H}, \tilde{T})$ such that the weakly connected components of the set $\tilde{\mathbf{J}}$ are exactly λ .

4.1 The enhancement operation

Consider the triple $\tau = (\tilde{\mathbf{J}}, \tilde{H}, \tilde{T}) \in \mathcal{H}$. Let $\tilde{\mathbf{J}}^1, \dots, \tilde{\mathbf{J}}^\lambda$ be the weakly connected components of the set $\tilde{\mathbf{J}}$. Then we define $\tilde{C}^i = \Gamma_{\tilde{\mathbf{J}}^i} \cup (\mathbf{clos}(\tilde{\mathbf{J}}^i) \cap \tilde{H})$ for $i \in \{1, \dots, \lambda\}$ and we call them *parts* of τ . Also we set $\tilde{T}^i = \tilde{T} \cap V(\tilde{C}^i)$, $1 \leq i \leq \lambda$. We now apply the following *enhancement* operation on each part of τ : For $i = 1, \dots, \lambda$, we consider the sequence $\mathcal{R}_\tau = (R_\tau^1, \dots, R_\tau^\lambda)$ where R_τ^i is the rooted graph $(R_\tau^i, \tilde{T}^i \cup \{x_{\text{new}}^i\})$ such that R_τ^i is defined as follows. Take the disjoint union of the graph \tilde{C}^i and a copy of the wheel $W_{\mu(\tilde{\mathbf{J}}^i)}$ with center x_{new}^i and add $\mu(\tilde{\mathbf{J}}^i)$ edges, called *i-external*, between the vertices of $V(\tilde{\mathbf{J}}^i)$ and the peripheral vertices of $W_{\mu(\tilde{\mathbf{J}}^i)}$ such that the resulting graph remains Σ_0 -embedded and each vertex $v \in V(\tilde{\mathbf{J}}^i)$ is incident to $\mu(v)$ non-homotopic edges not in $\tilde{\mathbf{J}}$. As the graph $\Gamma_{\tilde{\mathbf{J}}^i}$ is connected and planar, the construction of R_τ^i is possible. Observe also that $R_\tau^i \setminus \tilde{\mathbf{J}}^i$ is *unique up to topological isomorphism*. To see this, it is enough to verify that for every two vertices in $R_\tau^i \setminus \tilde{\mathbf{J}}^i$ of degree ≥ 3 there are always 3 disjoint paths connecting them (here we use our assumption that $\Gamma_{\tilde{\mathbf{J}}}$ does not contain any loops or multiple edges).

We define $\mathcal{R} = \{\mathcal{R}_\tau \mid \tau \in \mathcal{H}\}$ and observe that $|\mathcal{R}|$ is bounded by a function of k because the same holds for $|\mathcal{H}|$. (For an example of the construction of \mathcal{R}_τ , see Fig. 4)

We now define (C^1, \dots, C^λ) such that $C^i = \Gamma_{\bar{\mathbf{F}}^i} \cup (\mathbf{clos}(\bar{\mathbf{F}}^i) \cap G)$, $i \in \{1, \dots, \lambda\}$. We call the graphs in (C^1, \dots, C^λ) *parts* of G and let $T^i = T \cap V(C^i)$, $1 \leq i \leq \lambda$, where $T = \{s_1, t_1, \dots, s_k, t_k\}$. As above we define the *enhancement* of the parts of G as follows. For each $i = 1, \dots, \lambda$ we define the rooted graph $G^{*i} = (G^i, T^i \cup \{x_{\text{new}}^{*i}\})$ where G^i is defined as follows: take the disjoint union of C^i and the wheel $W_{\mu(\bar{\mathbf{F}}^i)}^*$ with center x_{new}^{*i} and add $\mu(\bar{\mathbf{F}}^i)$ edges, called **i-external*, between the vertices of $V(\bar{\mathbf{F}}^i)$ and the peripheral vertices of $W_{\mu(\bar{\mathbf{F}}^i)}^*$ such that the resulting graph remains Σ_0 -embedded and each vertex $v \in V(\bar{\mathbf{F}}^i)$ is incident to $\mu(v)$ non-homotopic edges. As above, each G^i is possible to construct and $G^i \setminus \bar{\mathbf{F}}^i$ is unique up to topological isomorphism.

The purpose of the above definitions is twofold. First, they help us to treat separately each of the parts of G and try to match them with the correct parts of τ . Second, the addition of the wheels to each part gives rise to a single, uniquely embeddable interface, between the part and its “exterior” and this helps us to treat embeddings as abstract graphs. Therefore, to check whether a part of τ is realizable within the corresponding part of G , we can use the rooted version of the topological minor relation on graphs as defined in Section 2.

4.2 The stretching lemma

A bijection ρ from $\tilde{T} = \{a_1, b_1, \dots, a_k, b_k\}$ to $T = \{s_1, t_1, \dots, s_k, t_k\}$ is *legal* if for every $i \in \{1, \dots, k\}$, there exists some $j \in \{1, \dots, k\}$ such that $\rho((a_i, b_i)) = (s_j, t_j)$.

Let $\tau \in \mathcal{H}$ and let ρ be a legal bijection from \tilde{T} to T and let ρ_i be the restriction of ρ in \tilde{T}_i . We say that $\mathcal{R}_\tau = (R_\tau^1, \dots, R_\tau^\lambda)$ is ρ -*realizable* in G if there exists a bijection $\phi: \{1, \dots, \lambda\} \rightarrow \{1, \dots, \lambda\}$ such that for $i = 1, \dots, \lambda$, R_τ^i is a $\hat{\rho}_i$ -rooted topological minor of $G^{*\phi(i)}$ where $\hat{\rho}_i = \rho_i \cup \{(x_{\text{new}}^i, x_{\text{new}}^{*\phi(i)})\}$.

By enumerating all possible bijections ϕ , we enumerate all possible correspondences between the parts of G and the parts of τ . In order to simplify notation, we assume in the remainder of this section that ϕ is the identity function.

The following lemma is crucial. It shows that when R_τ^i is a topological minor of G^{*i} we can always assume that all vertices and edges of \tilde{C}^i are mapped via ψ_0 and ψ_1 to vertices and paths in $\mathbf{clos}(\bar{\mathbf{F}}^i)$; the wheel $W_{\mu(\tilde{\mathbf{J}}^i)}$ is mapped to a “sub-wheel” of $W_{\mu(\bar{\mathbf{F}}^i)}$ while *i-external* edges are mapped to **i-external* edges. This will be useful in the proof of Lemma 6, as the *i-external* edges represent the interface of the completion \tilde{P} with \tilde{C}^i . The topological minor relation certifies that the same interface is feasible between the corresponding part C^i of G and its “exterior”. Lemma 5 establishes also that the image of $\Gamma_{\tilde{\mathbf{J}}^i}$ can be “stretched” so that it falls on $\Gamma_{\bar{\mathbf{F}}^i}$. As all the vertices in $V(\bar{\mathbf{F}}^i)$ are within distance 2 from the artificial terminal x_{new}^{*i} in G^{*i} , this allows us later in the proof of Lemma 9 to locate within G^{*i} the possible images of $V(R_\tau^i)$ in a neighborhood of the terminals. It is then safe to look for an irrelevant vertex “far away” from this neighborhood.

Lemma 5 *Let R_τ^i be a $\hat{\rho}_i$ -rooted topological minor of G^{*i} where $\hat{\rho}_i = \rho_i \cup \{(x_{\text{new}}^i, x_{\text{new}}^{*i})\}$, for $i = 1, \dots, \lambda$. Let also ψ_0^i and ψ_1^i be the functions (cf. Section 2) certifying this topological minor relation. Then ψ_0^i and ψ_1^i can be modified so that the following properties*

are satisfied.

1. if \tilde{e} is an edge of the wheel $W_{\mu(\tilde{\mathbf{J}}^i)}$ incident to x_{new}^i , then $\psi_1^i(\tilde{e})$ is an edge incident to x_{new}^{*i} .
2. if \tilde{e} is an edge of the wheel $W_{\mu(\tilde{\mathbf{J}}^i)}$ not incident to x_{new}^i , then $\psi_1^i(\tilde{e})$ is an x_{new}^{*i} -avoiding path of $W_{\mu(\tilde{\mathbf{F}}^i)}^*$. Moreover, each edge in $W_{\mu(\tilde{\mathbf{F}}^i)}^*$ is contained in a path $\psi_1^i(\tilde{f})$ for some edge \tilde{f} of the cycle $W_{\mu(\tilde{\mathbf{J}}^i)}^* \setminus x_{\text{new}}^i$.
3. if \tilde{e} is an i -external edge between $V(\tilde{\mathbf{J}}^i)$ and $V(W_{\mu(\tilde{\mathbf{J}}^i)} \setminus \{x_{\text{new}}^i\})$, then $\psi_1(e)$ is a path consisting of an $*i$ -external edge between $V(\tilde{\mathbf{F}}^i)$ and $V(W_{\mu(\tilde{\mathbf{F}}^i)} \setminus \{x_{\text{new}}^{*i}\})$.
4. $\psi_0^i(V(\tilde{\mathbf{J}}^i)) \subseteq V(\tilde{\mathbf{F}}^i)$.

Proof. Our first step is to enforce Properties 1 and 2. For each edge $\tilde{e} = \{x_{\text{new}}^i, x\}$ of R_τ^i , let $\tilde{e}' = \{x, x'\}$ be the unique edge of R_τ^i that is not an edge of $W_{\mu(\tilde{\mathbf{J}}^i)}$. Consider the path $P_{\tilde{e}}$ that is formed by the concatenation of $\psi_1^i(\tilde{e}')$ and $\psi_1^i(\tilde{e})$ and let \hat{x} be the neighbor of $\psi_0^i(x_{\text{new}}^i) = x_{\text{new}}^{*i}$ in the path $\psi_1^i(P_{\tilde{e}})$. For each such \tilde{e} and its incident vertex x , we simultaneously update ψ_0^i such that $\psi_0^i(x) = \hat{x}$. Accordingly, for every $\tilde{e} = \{x_{\text{new}}^i, x\}$ we simultaneously update ψ_1^i so that $\psi_1^i(\{x_{\text{new}}^i, x\})$ is the path consisting of the edge $\{x_{\text{new}}^{*i}, \hat{x}\}$ (enforcing Property 1) and $\psi_1^i(x, x')$ is the subpath of $P_{\tilde{e}}$ between \hat{x} and $\psi_0^i(x')$. We also update ψ_1^i such that for each $\tilde{e}_1 = \{x_{\text{new}}^i, x_1\}$ and $\tilde{e}_2 = \{x_{\text{new}}^i, x_2\}$ where x_1 and x_2 are adjacent, we set $\psi_1^i(\{x_1, x_2\})$ to be the (\hat{x}_1, \hat{x}_2) -path of $W_{\mu(\tilde{\mathbf{F}}^i)}^*$ that does not meet any other vertex of $\psi_0^i(V(W_{\mu(\tilde{\mathbf{J}}^i)}))$, enforcing Property 2.

So far, we have enforced that for each $i \in \{1, \dots, \lambda\}$, the wheel $W_{\mu(\tilde{\mathbf{F}}^i)}$ is mapped via ψ_0^i and ψ_1^i to the wheel $W_{\mu(\tilde{\mathbf{J}}^i)}$ (by slightly abusing the notation, we can say that $\psi_1(W_{\mu(\tilde{\mathbf{J}}^i)}) = W_{\mu(\tilde{\mathbf{F}}^i)}$). Notice that Properties 1 and 2 imply that all vertices and edges of \tilde{C}^i are mapped via ψ_0 and ψ_1 to vertices and paths of $\mathbf{clos}(\tilde{\mathbf{F}}^i)$.

Our second step is to enforce Properties 3 and 4. The transformation that we describe next, essentially “stretches” the image of $\tilde{\mathbf{J}}^i$ until it hits from within the boundary of $\tilde{\mathbf{F}}^i$. Let $\tilde{e} = \{x, x'\}$ be an edge of R_τ^i where $x \in V(\tilde{\mathbf{J}}^i)$ and $x' \in V(W_{\mu(\tilde{\mathbf{J}}^i)} \setminus \{x_{\text{new}}^i\})$. We also set $P_{\tilde{e}} = \psi_1^i(\tilde{e})$. By definition, $P_{\tilde{e}}$ is a path in G^{*i} which starts from $\psi_0^i(x)$ and ends at $\psi_0^i(x')$. As $\psi_0^i(x)$ is a point of $\mathbf{clos}(\tilde{\mathbf{F}}^i)$ and $\psi_0^i(x')$ is not, we can define \hat{x} as the first vertex of $P_{\tilde{e}}$ that is a vertex of $V(\tilde{\mathbf{F}}^i)$. For each such \tilde{e} and its incident vertex x , we simultaneously update ψ_0^i such that $\psi_0^i(x) = \hat{x}$. Accordingly, for every $\tilde{e} = \{x, x'\}$ we simultaneously update ψ_1^i so that $\psi_1^i(\{x, x'\})$ is the path consisting of the single edge $\{\hat{x}, \psi_0^i(x')\}$ – thus enforcing Property 3 – and for each edge $\{x_1, x_2\}$ of $\Gamma_{\tilde{\mathbf{J}}^i}$ we update ψ_1^i so that $\psi_1^i(\{x_1, x_2\})$ is a (\hat{x}_1, \hat{x}_2) -path avoiding any other vertex of $\psi_0^i(V(\tilde{\mathbf{J}}^i))$. As now all images of the vertices and the edges of $\Gamma_{\tilde{\mathbf{J}}^i}$ lie on $\Gamma_{\tilde{\mathbf{F}}^i}$, Property 4 holds. \square

4.3 Reducing PDPC to topological minor testing

Lemma 6 PDPC($G, t_1, s_1, \dots, t_k, s_k, \ell, \mathbf{F}$) has a solution if and only if there exists a $\tau = (\tilde{\mathbf{J}}, \tilde{H}, \tilde{T}) \in \mathcal{H}$ and a legal bijection $\rho: \tilde{T} \rightarrow T$ such that \mathcal{R}_τ is ρ -realizable in G .

Proof. Suppose that (P, \mathbf{J}) is a solution for $\text{PDPC}(G, t_1, s_1, \dots, t_k, s_k, \ell, \mathbf{F})$ giving rise to a collection $\mathcal{P} = \{P_1, \dots, P_k\}$ of disjoint paths in $G \cup P$ where P_i connects s_i with t_i .

From Theorem 2, we can assume that $|V(\mathbf{J})|$ is bounded by a suitable function of k and therefore, (P, \mathbf{J}) is isomorphic to a member $(\tilde{P}, \tilde{\mathbf{J}})$ of \mathcal{J} , in the sense that $\tilde{P} \cup \Gamma_{\tilde{\mathbf{J}}}$ and $P \cup \Gamma_{\mathbf{J}}$ are topologically isomorphic. Let (H, T) be the rooted subgraph of G formed by the edges of the paths in \mathcal{P} that are also edges of G . We now define the rooted graph (H', T) by dissolving all non-terminal vertices in the interior of \mathbf{J} and observe that (H', T) is isomorphic, with respect to some bijection ω , to some (\tilde{H}, \tilde{T}) that is compatible with $(\tilde{P}, \tilde{\mathbf{J}})$. Therefore $\tau = (\tilde{\mathbf{J}}, \tilde{H}, \tilde{T}) \in \mathcal{H}$ and let ρ be the correspondence between \tilde{T} and T , induced by ω . By partitioning ω with respect to the weakly connected components of the set $\tilde{\mathbf{J}}$, we generate a bijection ϕ between the parts $\{\tilde{C}^1, \dots, \tilde{C}^\lambda\}$ of τ and the parts $\{C^1, \dots, C^\lambda\}$ of G such that a subdivision of \tilde{C}^i is topologically isomorphic to a subgraph of $C^{\phi(i)}$, $i \in \{1, \dots, \lambda\}$. This topological isomorphism can be extended in the obvious manner to the graphs obtained after the enhancement operation applied to the parts of τ and the parts of G . This translates to the fact that each of the resulting R_τ^i is a ρ_i -rooted topological minor of $G^{*\phi(i)}$ where each ρ_i is obtained by the restriction of ρ to the vertices of $V(R_\tau^i)$ plus the pair $(x_{\text{new}}^i, x_{\text{new}}^{*i})$. We conclude that \mathcal{R}_τ is ρ -realizable in G .

For the opposite direction, let $\tau \in \mathcal{H}$. By renumbering if necessary the parts of G , let R_τ^i be a $\hat{\rho}_i$ -rooted topological minor of G^{*i} were $\hat{\rho}_i = \rho_i \cup \{(x_{\text{new}}^i, x_{\text{new}}^{*i})\}$, for $i = 1, \dots, \lambda$. Let also ψ_0^i and ψ_1^i be the functions satisfying the four properties of Lemma 5.




Figure 5: The graph \tilde{H}_K^+ when \tilde{H}^+ is the leftmost graph of Figure 4.

For the given $\tau = (\tilde{\mathbf{J}}, \tilde{H}, \tilde{T}) \in \mathcal{H}$ we assume w.l.o.g. that $|V(\tilde{\mathbf{J}})|$ is minimal. Recall that there exists a \tilde{P} such that (\tilde{H}, \tilde{T}) is compatible with $(\tilde{P}, \tilde{\mathbf{J}})$. Let $\tilde{H}^+ = \tilde{H} \cup \Gamma_{\tilde{\mathbf{J}}} \cup \tilde{P}$. From Lemma 1, there exists a curve K where $K \subseteq \Sigma_0 \setminus \text{clos}(\tilde{\mathbf{J}})$ and such that K intersects each edge of \tilde{P} twice. Let Δ_{in} and Δ_{out} be the two connected components of the set $\Sigma_0 \setminus K$ and w.l.o.g. we may assume that $\text{clos}(\tilde{\mathbf{J}}) \subseteq \Delta_{\text{out}}$. If we consider all

connected components of the set $K \setminus \tilde{H}^+$ as edges and take their union with the graph obtained by \tilde{H}^+ after subdividing the edges of \tilde{P} at the points of $K \cap \tilde{H}^+$, we create a new graph \tilde{H}_K^+ (see Figure 5 for an example). We also define $\tilde{H}_{K,\text{in}}^+ = \tilde{H}_K^+ \cap \Delta_{\text{in}}$ and $\tilde{H}_{K,\text{out}}^+ = \tilde{H}_K^+ \cap \text{clos}(\Delta_{\text{out}})$. For $i \in \{1, \dots, \lambda\}$, we define $\tilde{H}_{K,\text{out}}^{+(i)}$ as follows: first take the graph obtained by $\tilde{H}_{K,\text{out}}^+$ if we remove all edges except those that have at least one endpoint in \tilde{C}^i and those that belong in K and, second, dissolve in this graph all vertices that have degree 2. If we now add a new vertex $x_{\text{new}}^i \in \Delta_{\text{in}}$ and make it adjacent to all remaining vertices in K , it follows, by the minimality of the choice of τ , that the resulting graph is isomorphic to R_τ^i . In the above construction, we considered the vertex set $Q = K \cap \tilde{H}^+ \subseteq V(H_K^+)$ and we slightly abuse notation by denoting by Q also its counterpart in R_τ^i . We also denote by π_Q the cyclic ordering of Q defined by the order of appearance of the vertices in Q along K . Similarly, for $i \in \{1, \dots, \lambda\}$, we set $Q_i = Q \cap V(\tilde{H}_{K,\text{out}}^{+(i)})$ and we denote by π_{Q_i} the induced sub-ordering of π_Q . We denote by $\tilde{E}^i = \{\tilde{e}_1^i, \dots, \tilde{e}_{\mu(\tilde{\mathbf{J}}^i)}^i, \tilde{e}_1^i\}$ the edges of R_τ^i , not in $W_{\mu(\tilde{\mathbf{J}}^i)}$, that are incident to the vertices of Q_i , cyclically ordered according to π_{Q_i} .

We now take the graph G^{*i} and remove from it all edges outside $\text{clos}(\bar{\mathbf{F}}^i)$ except from the images of the edges of \tilde{E}^i (observe that $\text{clos}(\bar{\mathbf{F}}^i)$ is always a connected set while this is not necessarily the case for $\bar{\mathbf{F}}^i$). From Property 3 in Lemma 5, these images are just edges between $V(\bar{\mathbf{F}}^i)$ and $V(W_{\mu(\bar{\mathbf{F}}^i)} \setminus \{x_{\text{new}}^{*i}\})$. We denote the resulting graph by G^{*i-} and observe that it can be drawn in Σ_0 in a way that the vertices in $\psi_0(Q_i)$ lie on a virtual closed curve in accordance with the cyclic ordering π_{Q_i} of their ψ_0^i -preimages. We consider the disjoint union G_{out}^{*} of all G^{*i-} and we observe that it is possible to embed it in Σ_0 such that all vertices in $Q^* = \bigcup_{i \in \{1, \dots, \lambda\}} \psi_0^i(Q_i)$ lie on the same closed curve K^* . By the way each π_{Q_i} is defined from π_Q and the bijection between Q and Q^* we directly obtain that the following properties are satisfied: (i) the new embedding is planar, (ii) G_{out}^{*} is a subset of one of the connected components of the set $\Sigma_0 \cap K^*$, (iii) for all $i \in \{1, \dots, \lambda\}$, the cyclic ordering of $\psi_0^i(Q_i)$ is an induced sub-ordering of the cyclic ordering defined by the way the vertices of Q^* are being arranged along K^* . Note that $\psi_0 = \bigcup_{i \in \{1, \dots, \lambda\}} \psi_0^i$ defines a bijection from the vertices of Q to the vertices of Q^* . Our next step is to obtain the graph G_{in}^{*} by taking $\tilde{H}_{K,\text{in}}^+$ and renaming each vertex $x \in Q$ to $\psi_0(x)$. Now the graph $G^+ = G_{\text{in}}^{*} \cup G_{\text{out}}^{*}$ consists of the original graph G and a collection \mathcal{Z} of paths of length 3.

We are now ready to define the desired \mathbf{F} -patch (P, \mathbf{J}) of G . The edges of P are created by dissolving all internal vertices of each path in \mathcal{Z} . Also, \mathbf{J} can be any cactus set for which the embedding of $\Gamma_{\mathbf{J}}$ results from the embedding of $\Gamma_{\mathbf{F}}$ after dissolving all vertices of $V(\mathbf{F})$ that are not in $\psi_0(V(\tilde{\mathbf{J}}))$. \square

In [13], Grohe, Kawarabayashi, Marx, and Wollan gave an $h_1(k) \cdot n^3$ algorithm for checking rooted topological minor testing, where h_1 is some computable function. Combining their algorithm with Lemma 6, we obtain an $h_2(k) \cdot n^3$ algorithm for PDPC (here, again, h_2 is some computable function). Therefore, $\text{PDPC} \in \text{FPT}$.

In the next section, we show how to obtain the improved running time claimed in Theorem 1.

5 Applying the irrelevant vertex technique

Lemma 6 established that for every candidate patch \tilde{P} , certifying its feasibility for PDPC on G reduces to finding a rooted topological minor. By the results of [10], PDPC is in FPT by an $h_2(k) \cdot n^3$ algorithm. In this section we show that this running time can be improved. We first explain the intuition behind this improvement.

5.1 Proof strategy

At a high level, we have reduced the validity of \tilde{P} to the problem of determining whether a q -vertex rooted graph H is a rooted topological minor of an n -vertex graph G , where q is bounded by some function of k . By Lemma 5, we can assume that the images of the vertices of H are either terminals of G or lie on the boundary of the same face of G . This observation makes it possible to directly employ, in Lemma 9, the *irrelevant vertex technique* [19].

In particular, if the treewidth of G is big enough, one can detect a sufficiently large set of concentric cycles that are away from the images of the vertices of H in G ; this is possible due to Lemma 5. Using the “vital linkage” theorem of Robertson and Seymour [17, 18] (see also [13]), we obtain that the topological minor mapping can be updated so that the realization of H avoids the inner cycle of this collection. Therefore, the removal of any of the vertices of this cycle creates an equivalent instance of the problem with a smaller number of vertices. By repeating this vertex-removal operation, we end up with a graph whose treewidth is bounded by some function of q . In this case, since the rooted variant of the topological minor checking problem is definable in Monadic Second Order Logic (MSOL) (see Lemma 10 in the Appendix), the problem can be solved in a linear number of steps according to Courcelle’s Theorem.

For the running time of our algorithm, we use the fact that the detection of an irrelevant vertex in planar graphs requires to find a vertex that is “far enough” from all the terminals. As this can be done by standard BFS in $O(n)$ steps and at most n such vertices are deleted, the overall complexity of the algorithm is $f(k) \cdot n^2$.

5.2 Treewidth and linkages

Let G be a graph. A *linkage* in G is a set of pairwise disjoint paths of it. The *endpoints* of a linkage \mathcal{L} are the endpoints of the paths in \mathcal{L} . The *pattern* of \mathcal{L} is defined as

$$\pi(\mathcal{L}) = \{\{s, t\} \mid \mathcal{L} \text{ contains a path from } s \text{ to } t\}$$

Consider now the rooted graph $G = (G', T)$ where G' is a Σ_0 -embedded graph. We call a cycle of G' *T-respectful* if all the vertices of T are inside one of the two connected components of the set $\Sigma_0 \setminus C$. Given a T -respectful cycle C of G' we denote by $\Delta_{\text{ext}}(C)$ the connected component of the set $\Sigma_0 \setminus C$ that contains T and by $\Delta_{\text{int}}(C)$ the other. A sequence C_1, \dots, C_k of T -respectful cycles in G is *T-concentric* if $\Delta_{\text{ext}}(C_1) \subseteq \dots \subseteq \Delta_{\text{ext}}(C_k)$.

Let Γ_k ($k \geq 2$) be the graph obtained from the $(k \times k)$ -grid by triangulating internal faces of the $(k \times k)$ -grid such that all internal vertices become of degree 6, all non-corner




Figure 6: The leftmost graph is Γ_6 . The insulation of x in the rightmost graph is equal to 2 (the terminals are the blue vertices).

external vertices are of degree 4, and then one corner of degree two is joined by edges – we call them *external* – with all vertices of the external face (the *corners* are the vertices that in the underlying grid have degree two). Graph Γ_6 is shown in Fig. 6.(i). We also define the graph Γ_k^* as the graph obtained from Γ_k if we remove all its external edges. In Γ_k^* we call all vertices incident to its unique non-triangle face *perimetric*.

Let $G = (G', T)$ be a rooted Σ_0 -embedded graph. Given $x \in V(G')$ we define the *insulation between x and T* , denoted $\mathbf{ins}_T(x)$ as the maximum length of a sequence of T -concentric cycles C_1, \dots, C_l in G such that (i) $\{x\} \cap \bigcup_{i \in \{1, \dots, l\}} C_i = \emptyset$ and (ii) every line of Σ_0 connecting x and some vertex of T meets every $C_i, i \in \{1, \dots, l\}$. For an example, see Figure 6.(ii). We define the *planar thickness* of a rooted graph $G = (G', T)$ as follows:

$$\mathbf{pth}_T(G) = \max\{\mathbf{ins}_T(x) \mid x \in V(G')\}.$$

It is easy to verify that \mathbf{pth} is closed under contractions. In other words, the following holds.

Lemma 7 *Let $G_1 = (G'_1, T_1)$ and $G_2 = (G'_2, T_2)$ be two rooted graphs where G'_1 is a σ -contraction of G'_2 and $T_1 = \sigma(T_2)$. Then $\mathbf{pth}_{T_1}(G_1) \leq \mathbf{pth}_{T_2}(G_2)$.*

According to the following lemma, if the treewidth of G is big enough then every constant radius ball around T is sufficiently insulated from some vertex of G .

Lemma 8 *Let $G = (G', T)$ be a rooted plane graph where $\mathbf{tw}(G') \geq 24 \cdot (2l + 2r + 2)(\lceil \sqrt{|T|} + 1 \rceil) + 49$. Let also $T' = \bigcup_{t \in T} N_{G'}^r(t)$. Then $\mathbf{pth}_{T'}(G) \geq l$.*

Proof. By using [7, Lemma 6], the graph G' contains as a σ -contraction the graph $H' = \Gamma_{(2l+2r+2)(\lceil \sqrt{|T|} + 1 \rceil)}$. Define the rooted graph $H = (H', T_H)$ where $T_H = \sigma(T)$. We

consider a vertex packing of H' in $|T| + 1$ copies of $\Gamma_{2l+2r+2}^*$. By the pigeonhole principle and the fact that $|T_H| \leq |T|$, one, say Z , of these copies does not contain any vertex in T_H . By Observation 1, contractions do not increase distances, hence $\sigma(T') \subseteq N_{H'}^r(T_H)$. Let Z' be the subgraph of Z that is isomorphic to Γ_{2l+2}^* whose vertices have distance at least r from the perimeter vertices of Z and therefore are also at distance strictly greater than r from T_H . We conclude that $V(Z') \cap \sigma(T') = \emptyset$. Notice that Z' contains $l + 1$ $\sigma(T')$ -concentric cycles C_1, \dots, C_l, C_{l+1} that are also $\sigma(T')$ -concentric cycles of H .

Let $x \in C_{l+1}$. Then the cycles C_1, \dots, C_l certify that $\mathbf{ins}_{\sigma(T')}(x) \geq l$. Therefore $\mathbf{pth}_{\sigma(T')}(H') \geq l$. By Lemma 7, we conclude that $\mathbf{pth}_{T'}(G) \geq l$. \square

We need the following theorem from [17]. We present it using the terminology of [5].

Proposition 1 *There is a computable function g such that the following holds: Let Γ be a Σ_0 -embedded plane graph, \mathcal{L} be a linkage of Γ and let T be the set of vertices in the pairs of $\pi(\mathcal{L})$. Let also $C_1, \dots, C_{g(|\pi(\mathcal{L})|)}$ be T -concentric cycles of Γ . Then there is a linkage \mathcal{L}' with the same pattern as \mathcal{L} such that all paths in \mathcal{L} are contained in $\Delta_{\text{ext}}(C_{g(|\pi(\mathcal{L})|)})$.*

5.3 The algorithm

Using Lemmata 5, 7, 8, and Proposition 1, we can prove the following result.

Lemma 9 *There is an FPT algorithm, running in $f_2(k) \cdot n^2$, for some function f_2 , that given a $\tau = (\tilde{\mathbf{J}}, \tilde{H}, \tilde{T}) \in \mathcal{H}$ and a legal bijection $\rho: \tilde{T} \rightarrow T$, checks whether \mathcal{R}_τ is ρ -realizable in G .*

Proof. As the number of bijections $\phi: \{1, \dots, \lambda\} \rightarrow \{1, \dots, \lambda\}$ is bounded by a function of k it is enough to show how to check in FPT time whether, for $i = 1, \dots, \lambda$, R_τ^i is a $\hat{\rho}_i$ -rooted topological minor of $G^{*\phi(i)}$ where $\hat{\rho}_i = \rho_i \cup \{(x_{\text{new}}^i, x_{\text{new}}^{*i})\}$. To simplify notation, we drop indices and we denote $R_\tau^i = (R_\tau^i, \tilde{T}^i \cup \{x_{\text{new}}^i\})$ and $G^{*\phi(i)} = (G'^{\phi(i)}, T^{\phi(i)} \cup \{x_{\text{new}}^{*\phi(i)}\})$ by $R = (R', \tilde{T} \cup \{x_{\text{new}}\})$ and $G^* = (G', T \cup \{x_{\text{new}}^*\})$ respectively. We also use $\hat{\rho}$ instead of $\hat{\rho}_i$ and x_{new} instead of x_{new}^i .

We now apply the *irrelevant vertex technique*, introduced in [19] as follows. Using the algorithm from [2] (or, alternatively, the one from [15]) one can either compute a tree-decomposition of G' of width at most $q = 4 \cdot (24 \cdot (2 \cdot g(|E(R')|) + 8)(\sqrt{|T| + 2}) + 49)$ or prove that no tree decomposition exists with width less than $q/4$. In the case where $\mathbf{tw}(G') \leq q$, we recall that $|E(R')|$ is a function of k and $|T| \leq 2k$. Consequently, there exists a function f_3 such that $\mathbf{tw}(G') \leq f_3(k)$ and the result follows from Lemma 10 in the Appendix.

Suppose now that $\mathbf{tw}(G') \geq q/4 = 24 \cdot (2 \cdot g(|E(R')|) + 8)(\sqrt{|T| + 2}) + 49$. Applying Lemma 8 for $r = 3$ we have that $\mathbf{pth}_{T'}(G^*) \geq g(|E(R')|)$ where $T' = N_{G'}^3(T \cup \{x_{\text{new}}^*\})$.

We now prove that there is a vertex $x \in V(G^*)$ such that if R is a $\hat{\rho}$ -rooted topological minor of G^* then R is a $\hat{\rho}$ -rooted topological minor of $G^* \setminus \{x\}$. Let ψ_0 and ψ_1 be the functions certifying that R is a $\hat{\rho}$ -rooted topological minor of G^* . We apply on ψ_0 and ψ_1 the modifications of the Lemma 5 so that they satisfy Properties 1–4. An important consequence is that the images of all vertices of R under ψ_0 are close to terminals in T . Indeed, from Properties 1 and 2, all neighbors of x_{new} are mapped via ψ_0 to

vertices that belong in $N_{G'}^1(x_{\text{new}}^*)$. Moreover, from Properties 3 and 4 it follows that $\psi_0(V(\tilde{\mathbf{J}})) \subseteq V(\mathbf{F}) \subseteq N_{G'}^2(x_{\text{new}}^*)$. So far we have proved that all vertices of R , except those inside $\tilde{\mathbf{J}}$, are mapped via ψ_0 to vertices in $N_{G'}^2(x_{\text{new}}^*)$. As all vertices of R that are inside $\tilde{\mathbf{J}}$ belong to \tilde{T} , it follows that $\psi_0(V(R)) \subseteq T'$.

The set $L = \{\psi_1(e) \mid e \in E(R)\}$ is a set of paths in G^* . Let $L_1 \subseteq L$ be the set of those paths that have length at most 2 and define $L_2 = L \setminus L_1$. For each path $Q \in L_2$ define its *interior*, denoted $\text{int}(Q)$, as the subpath of Q consisting of all vertices of $I(Q)$. Clearly, $\mathcal{L} = \{\text{int}(Q) \mid Q \in L_2\}$ is a linkage in G^* and $\pi(\mathcal{L}) \subseteq T'$. Consider the collection \mathcal{C} of T' -concentric cycles $C_1, \dots, C_{g(|E(R')|)}$ certifying the fact that $\mathbf{pth}_{T'}(G^*) \geq g(|E(R')|)$. Define the graph $\Gamma = (\mathbf{clos}(\Delta_{\text{in}}(C_1)) \cap G^*) \cup \left(\bigcup_{Q \in L_2} \text{int}(Q)\right)$. We have that (i) \mathcal{L} is a linkage of Γ , and (ii) $|E(R')| \geq |\pi(\mathcal{L})|$. Let x be a vertex of $C_{g(|\pi(\mathcal{L})|)}$. By Proposition 1 there is another linkage \mathcal{L}' with the same pattern as \mathcal{L} such that all paths in \mathcal{L}' avoid x .

The vertices of all paths in L_1 belong to T' since they are at distance at most 3 from $T \cup x_{\text{new}}^*$. The endpoints of all paths in L_2 also belong T' since they are at distance at most 2 from $T \cup x_{\text{new}}^*$. Therefore all paths in L_1 and the endpoints of all paths in L_2 avoid Γ altogether. We now show that all paths in L can be rerouted so that they avoid x , while they remain internally vertex-disjoint. The paths in L_1 stay the same. For the paths in L_2 , we only have to reroute their interiors within Γ . This is achieved by connecting the pairs in $\pi(\mathcal{L})$ via the linkage \mathcal{L}' . After this substitution all paths in the updated set L avoid x . By updating ψ_1 to reflect the new interiors of the paths in L_2 , we obtain that R is a $\hat{\rho}$ -rooted topological minor of $G^* \setminus \{x\}$.

Notice that x can be found in linear time applying BFS starting from T' . After deleting x , we create an equivalent instance of the problem with smaller size. By recursively applying the same reduction to the new instance at most $|V(G')|$ times, at some point the treewidth will drop below q and then we solve the problem by applying Lemma 10 in the Appendix as above. \square

From Lemmata 6 and 9 we obtain Theorem 1.

6 Further extensions and open problems

We chose to tackle the disjoint-paths completion problem with the topological restriction of having non-crossing patch edges. A natural extension of this problem is to allow a fixed number $\xi > 0$ of crossings in the patch. We believe that, using the same techniques, one may devise an $f(k) \cdot n^2$ algorithm for this problem as well. The only substantial difference is a generalization of our combinatorial result (Theorem 2) under the presence of crossings.

An interesting topic for future work is to define and solve the disjoint-paths completion problem for graphs embedded in surfaces of higher genus. A necessary step in this direction is to extend Theorem 2 for the case where the face to be patched contains handles.

Another issue is to extend the whole approach for the case where the faces to be patched are more than one. This aim can be achieved without significant deviation from our methodology, in case the number of these faces is bounded. However, when this

restriction does not apply, the problem seems challenging and, in our opinion, it is not even clear whether it belongs to FPT.

Acknowledgement. We wish to thank the anonymous reviewers of an earlier version of this paper for valuable comments and suggestions.

References

- [1] I. ADLER, S. G. KOLLIPOULOS, P. K. KRAUSE, D. LOKSHTANOV, S. SAURABH, AND D. M. THILIKOS, *Tight bounds for linkages in planar graphs*, in Proceedings of the 38th International Colloquium on Automata, Languages and Programming (ICALP 2011), 2011.
- [2] E. AMIR, *Efficient approximation for triangulation of minimum treewidth*, in Uncertainty in Artificial Intelligence: Proc. of the Seventeenth Conference (UAI-2001), San Francisco, CA, Morgan Kaufmann Publishers, 2001, pp. 7–15.
- [3] B. COURCELLE, *The monadic second-order logic of graphs. I. Recognizable sets of finite graphs*, Information and Computation, 85 (1990), pp. 12–75.
- [4] B. COURCELLE AND M. MOSBAH, *Monadic second-order evaluations on tree-decomposable graphs*, Theoretical Computer Science, 109 (1993), pp. 49–82.
- [5] A. DAWAR, M. GROHE, AND S. KREUTZER, *Locally excluding a minor*, in Proc. of the 21st IEEE Symposium on Logic in Computer Science (LICS'07), New York, 2007, IEEE, pp. 270–279.
- [6] R. DIESTEL, *Graph theory*, vol. 173 of Graduate Texts in Mathematics, Springer-Verlag, Berlin, third ed., 2005.
- [7] F. V. FOMIN, P. GOLOVACH, AND D. M. THILIKOS, *Contraction bidimensionality: the accurate picture*, in 17th Annual European Symposium on Algorithms (ESA '09), Lecture Notes in Computer Science, Springer Verlag, 2009, pp. 706–717.
- [8] M. R. GAREY AND D. S. JOHNSON, *Computers and Intractability, A Guide to the Theory of NP-Completeness*, W.H. Freeman and Company, New York, 1979.
- [9] M. C. GOLUMBIC, H. KAPLAN, AND R. SHAMIR, *On the complexity of DNA physical mapping*, Adv. in Appl. Math., 15 (1994), pp. 251–261.
- [10] M. GROHE, K. KAWARABAYASHI, D. MARX, AND P. WOLLAN, *Finding topological subgraphs is fixed-parameter tractable*, in 43rd ACM Symposium on Theory of Computing (STOC 2011), San Jose, California, June 6-8 2011.
- [11] P. HEGGERNES, C. PAUL, J. A. TELLE, AND Y. VILLANGER, *Interval completion with few edges*, in Proceedings of the 39th Annual ACM Symposium on the Theory of Computing (STOC 2007), San Diego, California, American Mathematical Society, 2007, pp. 374–381.

- [12] H. KAPLAN, R. SHAMIR, AND R. E. TARJAN, *Tractability of parameterized completion problems on chordal, strongly chordal, and proper interval graphs*, SIAM Journal on Computing, 28 (1999), pp. 1906–1922.
- [13] K. KAWARABAYASHI AND P. WOLLAN, *A shorter proof of the graph minor algorithm - the unique linkage theorem*, in Proc. of the 42nd annual ACM Symposium on Theory of Computing (STOC 2010), 2010.
- [14] M. R. KRAMER AND J. VAN LEEUVEN, *The complexity of wire-routing and finding minimum area layouts for arbitrary VLSI circuits*, Advances in Comp. Research, 2 (1984), pp. 129–146.
- [15] B. A. REED, *Finding approximate separators and computing tree width quickly*, in Proc. of the Twenty-fourth annual ACM Symposium on Theory of Computing, ACM Press, 1992, pp. 221–228.
- [16] B. A. REED, N. ROBERTSON, A. SCHRIJVER, AND P. D. SEYMOUR, *Finding disjoint trees in planar graphs in linear time*, in Graph structure theory (Seattle, WA, 1991), vol. 147 of Contemp. Math., Amer. Math. Soc., Providence, RI, 1993, pp. 295–301.
- [17] N. ROBERTSON AND P. SEYMOUR, *Graph minors. XXII. Irrelevant vertices in linkage problems, preprint*, 1992.
- [18] N. ROBERTSON AND P. SEYMOUR, *Graph minors. XXI. Graphs with unique linkages*, J. Combin. Theory Ser. B, 99 (2009), pp. 583–616.
- [19] N. ROBERTSON AND P. D. SEYMOUR, *Graph minors. XIII. The disjoint paths problem*, J. Combin. Theory Ser. B, 63 (1995), pp. 65–110.
- [20] M. YANNAKAKIS, *Computing the minimum fill-in is NP-complete*, SIAM J. Algebraic Discrete Methods, 2 (1981), pp. 77–79.

Appendix

A.1 MSOL and rooted topological minors

In this section we define a parameterized version of Rooted Topological Minor Testing that uses treewidth as a parameter and we show, using Monadic Second Order Logic (MSOL), that it is in FPT. The corresponding algorithm will be used as a subroutine in the proof of Theorem 1.

p-BOUNDED TREEWIDTH ROOTED TOPOLOGICAL MINOR TESTING

Input: a rooted graph (H, S_H) , a rooted graph (H, S_G) , and a bijection $\rho: S_H \rightarrow S_G$.

Parameter: $k = \text{tw}(G) + |V(H)|$

Question: is (H, S_H) a ρ -rooted topological minor of the rooted graph (G, S_G) ?

Lemma 10 *p*-BOUNDED TREEWIDTH ROOTED TOPOLOGICAL MINOR TESTING can be solved in $f(k) \cdot |V(G)|$ steps.

Proof. Let $S_H = \{a_1, \dots, a_s\}$, let $V(H) \setminus S_H = \{u_1, \dots, u_t\}$, and let $E(H) = \{e_1, \dots, e_\epsilon\}$.

Let $\tau_s = \{\text{VERT}, \text{EDGE}, I, c_1, \dots, c_s\}$ be the vocabulary of graphs (as incidence structures) with s constants. We give an MSOL $[\tau_s]$ formula ϕ_{H, S_H} such that for every graph G with $S_G = \{b_1, \dots, b_n\}$ and $\rho(a_i) = b_i$ we have

$$(G, b_1, \dots, b_n) \models \phi_{H, S_H} \iff (H, S_H) \text{ is a } \rho\text{-rooted topological minor of } (G, S_G).$$

Let $\text{path}(x, y, Z)$ be the MSOL $[\tau_s]$ formula stating that Z is a path from x to y . (This can be easily done by saying that Z is a set of edges with the property that for every vertex v incident to an edge in Z , the vertex v is either incident to precisely two edges in Z , or v is incident to one edge in Z and $v = x$ or $v = y$. Finally, we can express that the path Z is connected.)

Let

$$\begin{aligned} \phi_{H, S_H} := & \exists Z_1 \dots Z_\epsilon \exists x_1 \dots x_t \left(\bigwedge_{i \neq j; i, j \leq t} x_i \neq x_j \wedge \bigwedge_{e_i = \{u_k, u_\ell\} \subseteq V(H) \setminus S_H} \text{path}(x_k, x_\ell, Z_i) \wedge \right. \\ & \bigwedge_{\substack{e_i = \{u_k, a_\ell\} \\ x_k \in V(H) \setminus S_H \\ a_\ell \in S_H}} \text{path}(x_k, c_\ell, Z_i) \wedge \bigwedge_{e_i = \{a_k, a_\ell\} \subseteq S_H} \text{path}(c_k, c_\ell, Z_i) \wedge \\ & \bigwedge_{i \neq j; i, j \leq \epsilon} \exists y \left((\text{VERT } y \wedge \text{'}y \text{ is incident to an edge in } Z_i \text{ and to an edge in } Z_j\text{'}) \right. \\ & \left. \rightarrow \left(\bigvee_{i=1}^s y = c_i \vee \bigvee_{i=1}^t y = x_i \right) \right). \end{aligned}$$

The constants c_i are interpreted by the b_i , hence they make sure that a_i is mapped to b_i , and Condition 1 of rooted topological minors is satisfied. In addition we make sure

that every edge of H is mapped to a path in G . Finally we make sure that Condition 3 is satisfied. (The statement ‘ x is incident to an edge in Z_i and to an edge in Z_j ’ can be easily expressed in MSOL.)

Observe that the length of the formula ϕ_{H,S_H} only depends on H . Hence by Courcelle’s Theorem [3,4] there is a computable function f_1 such that p -BOUNDED TREEWIDTH ROOTED TOPOLOGICAL MINOR TESTING can be solved in time $f_1(\mathbf{tw}(G) + |V(H)|) \cdot |V(G)|$. \square