Open access • Book Chapter • DOI:10.1007/978-3-662-53174-7_12

Finding Paths in Grids with Forbidden Transitions — Source link

Mamadou Moustapha Kanté, Fatima Zahra Moataz, Benjamin Momège, Nicolas Nisse

Institutions: Blaise Pascal University, University of Nice Sophia Antipolis

Published on: 17 Jun 2015 - Workshop on Graph-Theoretic Concepts in Computer Science

Topics: Longest path problem, Forbidden graph characterization, Treewidth, Planar graph and Path (graph theory)

Finding Paths in Grids with Forbidden Transitions

Mamadou Moustapha Kanté, Fatima Zahra Moataz, Benjamin Momège, Nicolas Nisse

Finding Paths in Grids with Forbidden Transitions

M. M. Kanté[1], F. Z. Moataz*[2,3], B. Momège[2,3], and N. Nisse[3,2]

[1]Univ. Blaise Pascal, LIMOS, CNRS, Clermont-Ferrand, France
[2]Univ. Nice Sophia Antipolis, CNRS, I3S, UMR 7271, 06900 Sophia Antipolis, France
[3]INRIA, France

Abstract

Une transition dans un graphe est une paire d'arêtes incidentes à un même sommet. Etant donnés un graphe $G = (V, E)$, deux sommets $s, t \in V$ et un ensemble associé de transitions interdites $\mathcal{F} \subseteq E \times E$, le problème de chemin évitant des transitions interdites consiste à décider s'il existe un chemin élémentaire de s à t qui n'utilise aucune des transitions de $\mathcal{F}$. C'est-à-dire qu'il est interdit d'emprunter consécutivement deux arêtes qui soient une paire de $\mathcal{F}$. Ce problème est motivé par le routage dans les réseaux routiers (où une transition interdite représente une interdiction de tourner) ainsi que dans les réseaux optiques avec des noeuds asymétriques. Nous prouvons que le problème est NP-difficile dans les graphes planaires et plus particulièrement dans les grilles. Nous montrons également que le problème peut être résolu en temps polynomial dans la classe des graphes de largeur arborescente bornée.

Introduction

Driving in New-York is not easy. Not because of the rush hours and the taxi drivers, but because of the no-left, no-right and no U-turn signs. Even in a "grid-like" city like New-York, prohibited turns might force you to cross several times the same intersection before eventually reaching your destination. In this report, we give hints explaining why it is difficult to deal with forbidden-turn signs when driving.

Let $G = (V, E)$ be a graph. A transition in G is a pair of two distinct edges incident to a same vertex. Let $\mathcal{F} \subseteq E \times E$ be a set of forbidden transitions in G. We say that a path $P = (v_0, \ldots, v_q)$ is $\mathcal{F}$-valid if it contains none of the transitions of $\mathcal{F}$, i.e., $\{(v_{i-1}, v_i), (v_i, v_{i+1})\} \notin \mathcal{F}$ for $i \in \{1, \ldots, q-1\}$. Given G, $\mathcal{F}$ and two vertices s and t, the Path Avoiding Forbidden Transitions (PAFT) problem is to find an $\mathcal{F}$-valid s-t-path. The PAFT problem arises in many contexts. In optical networks, nodes can have asymmetric switching capabilities mostly due to cost-relevant reasons [3, 6]. This means that a node has some restrictions on its internal connectivity: traffic on a certain ingress port can only reach a subset of the egress ports. In this setting, the optical nodes configured asymmetrically are vertices with forbidden transitions and the routing problem is an application of PAFT. The study of PAFT is also motivated by its relevance to vehicle routing. In road networks, it is possible that some roads are closed due to traffic jams, construction, etc. It is also frequent to encounter no-left, no-right and no U-turn signs at intersections. These prohibited roads and turns can be modeled by forbidden transitions.

When the PAFT problem is studied, a distinction has to be made according to whether the path to find is elementary (cannot repeat vertices) or non-elementary. Indeed, PAFT can be solved in polynomial time [8] for the non-elementary case (using a simple BFS from t) while finding an elementary path avoiding forbidden transitions has been proved NP-complete in [11]. In this report, we study the elementary version of the PAFT problem in planar graphs and more particularly in grids. Our interest for planar graphs is motivated by the fact that they are closely related to road networks. They are also an interesting special case

*This author is supported by a grant from the "Conseil régional Provence Alpes-Côte d'Azur".

to study while trying to capture the difficulty of the problem. Furthermore, to the best of our knowledge, this case has not been addressed before in the literature.

Related work. PAFT is a special case of the problem of finding a path avoiding forbidden paths (PFP) introduced in [14]. Given a graph G, two vertices s and t, and a set $\mathcal{S}$ of forbidden paths, PFP aims at finding an s-t-path which contains no path of $\mathcal{S}$ as a subpath. When the forbidden paths are composed of exactly two edges, PFP is equivalent to PAFT. Many papers address the non-elementary version of PFP, proposing exact polynomial solutions [14, 1]. The elementary counterpart has been recently studied in [9] where a mathematical formulation is given and two solution approaches are developed and tested. The computational complexity of the elementary PFP can be deduced from the complexity of PAFT which has been established in [11]. Szeider proved in [11] that finding an elementary path avoiding forbidden transitions is NP-complete and gave a complexity classification of the problem according to the types of the forbidden transitions. The NP-completeness proof in [11] does not extend to planar graphs.

PAFT is also a generalization of the problem of finding a properly colored path in an edge-colored graph (PEC). Given an edge-colored graph G^c and two vertices s and t, the PEC problem aims at finding an s-t-path such that any consecutive two edges have different colors. It is easy to see that PEC is equivalent to PAFT when the set of forbidden transitions consists of all pairs of adjacent edges that have the same color. The PEC problem is proved to be NP-complete in directed graphs [7] which directly implies that the PAFT problem is NP-complete in directed graphs[1].

Contribution. Our main contribution is the proof that the PAFT problem is NP-complete in grids. We also prove that the problem can be solved in time $O((3\Delta(k+1))^{2k+O(1)}n))$ in n-node graphs with treewidth at most k and maximum degree Δ. In other words, we prove that the PAFT problem is FPT in $k+\Delta$. Our NP-completeness result strengthens the one of Szeider [11] established in 2003 and extends to the problem of PFP. The FPT result, on the other hand, highlights some of the parameters that make PAFT tractable.

The document is organized as follows. The problem of PAFT is formally stated in Section 1. In Section 2, the problem is proven NP-complete in grids. Finally, a polynomial time algorithm for graphs with bounded treewidth is presented in Section 3.

1 Problem statement

Let $G = (V, E)$ be a (multi-)graph. We denote by Δ the maximum degree of the G. Given a subgraph H of G, a *transition* in H is a (not ordered) set of two distinct edges of H incident to a same node. Namely, $\{e, f\}$ is a transition if $e, f \in E(H)$, $e \neq f$ and $e \cap f \neq \emptyset$. Let $\mathcal{T}$ denote the set of all transitions in G. Let $\mathcal{F} \subseteq \mathcal{T}$ be a set of *forbidden transitions*. A transition in $\mathcal{A} = \mathcal{T} \setminus \mathcal{F}$ is said *allowed*.

A *path* is any sequence $(v_0, v_1, \cdots, v_r)$ of vertices such that $v_i \neq v_j$ for any $0 \leq i < j \leq r$ and $e_i = \{v_i, v_{i+1}\} \in E$ for any $0 \leq i < r$. Given two vertices s and t in G, a path $P = (v_0, v_1, \cdots, v_r)$ is called an *st-path* if $v_0 = s$ and $v_r = t$. Finally, a path $P = (v_0, v_1, \cdots, v_r)$ is *$\mathcal{F}$-valid* if any transition in P is allowed, i.e., $\{e_i, e_{i+1}\} \notin \mathcal{F}$ for any $0 \leq i < r$.

Problem 1 (Problem of Finding a Path Avoiding Forbidden Transitions, PAFT)**.** *Given a (multi)-graph $G = (V, E)$, a set $\mathcal{F}$ of forbidden transitions and two vertices $s, t \in V$. Is there an $\mathcal{F}$-valid st-path in G?*

2 NP-completeness results

We start by proving that the PAFT problem is NP-complete in grids. For this purpose, we first prove that it is NP-complete in planar graphs with maximum degree at most 8 by a reduction from 3-SAT. Then, we propose simple transformations to reduce the degree of the vertices and prove that the PAFT problem is NP-complete in planar graphs with degree at most 4. Finally, we prove it is NP-complete in grids.

[1]Note that, in [7], the authors state that their result can be extended to planar graph. However, there is a mistake in the proof of the corresponding Corollary 7: to make their graph planar, vertices are added when edges intersect. Unfortunately, this transformation does not preserve the fact that the path is elementary.

Lemma 1. *The Problem of finding a path avoiding forbidden transitions is NP-complete in planar multi-graphs with maximum degree* 8.

Proof. The problem is clearly in NP. We prove the completeness using a reduction from the 3-SAT problem. Let Φ be an instance of 3-SAT, i.e., Φ is a boolean formula with variables $\{X_1, \cdots, X_n\}$ and clauses $\{C_1, \cdots, C_m\}$.

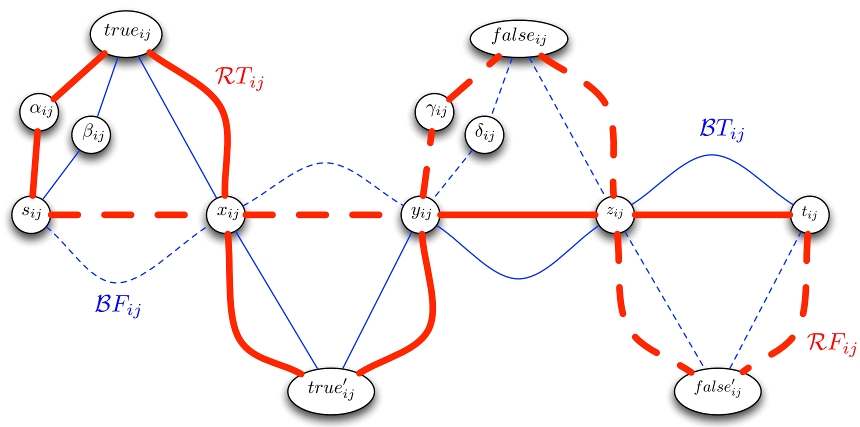

Figure 1: Mini-gadget H_{ij} which is the "superposition" of 4 edge-disjoint paths $\mathcal{R}T_{ij}$ (full bold and red), $\mathcal{B}T_{ij}$ (full thin and blue), $\mathcal{R}F_{ij}$ (dotted bold and red) and $\mathcal{B}F_{ij}$ (dotted thin and blue). The mini-gadget G_{ij} is obtained from H_{ij} by possibly adding the edge $\{\alpha_{ij}, \beta_{ij}\}$ (if variable X_i appears positively in Clause C_j) or the edge $\{\gamma_{ij}, \delta_{ij}\}$ (if variable X_i appears negatively in clause C_j).

Mini-gadget G_{ij}. For any $1 \leq i \leq n$ and $1 \leq j \leq m$, let us define the mini-gadget G_{ij} and the corresponding set $\mathcal{F}_{ij}$ of forbidden transitions. First, let us define the multi-graph H_{ij} which consists of the "superposition" of 4 edge-disjoint paths:

- $\mathcal{R}T_{ij} = (s_{ij}, \alpha_{ij}, true_{ij}, x_{ij}, true'_{ij}, y_{ij}, z_{ij}, t_{ij})$;
- $\mathcal{B}T_{ij} = (s_{ij}, \beta_{ij}, true_{ij}, x_{ij}, true'_{ij}, y_{ij}, z_{ij}, t_{ij})$;
- $\mathcal{R}F_{ij} = (s_{ij}, x_{ij}, y_{ij}, \gamma_{ij}, false_{ij}, z_{ij}, false'_{ij}, t_{ij})$;
- $\mathcal{B}F_{ij} = (s_{ij}, x_{ij}, y_{ij}, \delta_{ij}, false_{ij}, z_{ij}, false'_{ij}, t_{ij})$;

That is, H_{ij} has 13 vertices and 28 edges as illustrated in Figure 1. Note that $\mathcal{R}T_{ij}, \mathcal{B}T_{ij}, \mathcal{R}F_{ij}, \mathcal{B}F_{ij}$ are four paths from s_{ij} to t_{ij}. Note also that $\mathcal{R}T_{ij}$ and $\mathcal{B}T_{ij}$ (resp., $\mathcal{R}F_{ij}$ and $\mathcal{B}F_{ij}$) are somehow "parallel".

Finally, the graph G_{ij} is obtained from H_{ij} by adding an edge $\{\alpha_{ij}, \beta_{ij}\}$ if variable X_i appears positively in clause C_j, respectively, by adding an edge $\{\gamma_{ij}, \delta_{ij}\}$ if variable X_i appears negatively in Clause C_j. If X_i does not appear in C_j, then $G_{ij} = H_{ij}$.

To simplify the presentation, edges of G_{ij} are given "colors". Each edge of $\mathcal{R}T_{ij}$ or $\mathcal{R}F_{ij}$ is said *red* and each edge of $\mathcal{B}T_{ij}$ or $\mathcal{B}F_{ij}$ is said *blue*. Moreover, the edges $\{\alpha_{ij}, \beta_{ij}\}$ and $\{\gamma_{ij}, \delta_{ij}\}$ (if any) are said *green*. Finally, to make our notation clearer, let us note that, intuitively, using the path $\mathcal{R}T_{ij}$ or $\mathcal{B}T_{ij}$ (full in Figure 1) will correspond to setting variable X_i to $True$. Respectively, using the path $\mathcal{R}F_{ij}$ or $\mathcal{B}F_{ij}$ (dotted in Figure 1)will correspond to setting variable X_i to $False$.

The **allowed** transitions in G_{ij} are the transitions of the four paths $\mathcal{R}T_{ij}, \mathcal{B}T_{ij}, \mathcal{R}F_{ij}, \mathcal{B}F_{ij}$. Moreover, if the edge $e = \{\alpha_{ij}, \beta_{ij}\}$ exists, then the four transitions $\{\{s_{ij}, \alpha_{ij}\}, e\}, \{\{s_{ij}, \beta_{ij}\}, e\}, \{\{true_{ij}, \alpha_{ij}\}, e\}$,

$\{\{true_{ij}, \beta_{ij}\}, e\}$ are also allowed. Respectively, if the edge $f = \{\gamma_{ij}, \delta_{ij}\}$ exists, then the four transitions $\{\{y_{ij}, \gamma_{ij}\}, f\}, \{\{y_{ij}, \delta_{ij}\}, f\}$, $\{\{false_{ij}, \gamma_{ij}\}, f\}, \{\{false_{ij}, \delta_{ij}\}, f\}$ are also allowed. Then, the set $\mathcal{F}_{ij}$ of forbidden transitions consists of all the transitions that have not been allowed.

The key properties of gadget G_{ij} are described in the following claims:

Claim 1. *The $\mathcal{F}_{ij}$-valid $s_{ij}t_{ij}$-paths in G_{ij} are $\mathcal{RT}_{ij}, \mathcal{BT}_{ij}, \mathcal{RF}_{ij}, \mathcal{BF}_{ij}$ and*

- *if variable X_i appears positively in Clause C_j:*
  - *the path $\mathcal{RBT}_{ij}$ that starts with the first edge $\{s_{ij}, \alpha_{ij}\}$ of $\mathcal{RT}_{ij}$, then uses green edge $\{\alpha_{ij}, \beta_{ij}\}$ and ends with all edges of $\mathcal{BT}_{ij}$ but the first one;*
  - *the path $\mathcal{BRT}_{ij}$ that starts with the first edge $\{s_{ij}, \beta_{ij}\}$ of $\mathcal{BT}_{ij}$, then uses green edge $\{\alpha_{ij}, \beta_{ij}\}$ and ends with all edges of $\mathcal{RT}_{ij}$ but the first one;*
- *if variable X_i appears negatively in Clause C_j:*
  - *the path $\mathcal{RBF}_{ij}$ that starts with the subpath $(s_{ij}, x_{ij}, y_{ij}\ \gamma_{ij})$ of $\mathcal{RF}_{ij}$, then uses green edge $\{\gamma_{ij}, \delta_{ij}\}$ and ends with the subpath of $\mathcal{BF}_{ij}$ that starts at δ_{ij} and ends at t_{ij};*
  - *the path $\mathcal{BRF}_{ij}$ that starts with the subpath $(s_{ij}, x_{ij}, y_{ij}\ \delta_{ij})$ of $\mathcal{BF}_{ij}$, then uses green edge $\{\delta_{ij}, \gamma_{ij}\}$ and ends with the subpath of $\mathcal{RF}_{ij}$ that starts at γ_{ij} and ends at t_{ij};*

Claim 2. *Let P be a $\mathcal{F}_{ij}$-valid $s_{ij}t_{ij}$-paths in G_{ij}. Then, either*

- *P passes through $true_{ij}$ and $true'_{ij}$ and does not pass through $false_{ij}$ nor $false'_{ij}$, or*
- *P passes through $false_{ij}$ and $false'_{ij}$ and does not pass through $true_{ij}$ nor $true'_{ij}$.*

Claim 3. *Let P be a $\mathcal{F}_{ij}$-valid $s_{ij}t_{ij}$-paths in G_{ij}. Then the first and last edges of P have different colors if and only if P uses a green edge, i.e., $P \in \{\mathcal{RBT}_{ij}, \mathcal{BRT}_{ij}, \mathcal{RBF}_{ij}, \mathcal{BRF}_{ij}\}$.*

The main graph of our reduction will be obtained from a copy of each gadget G_{ij}, $1 \leq i \leq n$, $1 \leq j \leq m$, combined in a "grid-like" manner. Let us continue with the description of the Clause-gadget that will correspond to the "row" of the main "grid".

Clause-gadget Row_j**.** Let $1 \leq j \leq m$, the Clause-gadget Row_j corresponding to clause C_j is obtained from a copy of each G_{ij}, $1 \leq i \leq n$, and from two nodes s_j and t_j. For any $1 \leq i \leq n$, the allowed transitions of G_{ij} are still allowed. Then, there are two cases depending on the parity of j. Roughly, when j is odd, the mini-gadgets are combined "from left to right" (from $i = 1$ to n), when j is even they are combined "from right to left" (from $i = n$ to 1).

Case j **odd.** First, let us add the red edge $e = \{s_j, s_{1j}\}$ and allow the two transitions between e and any red edge incident to s_{1j} in G_{1j}. Namely, the transitions $\{e, \{s_{1j}, \alpha_{1j}\}\}$ and $\{e, h\}$ are allowed, where h is the red edge $\{s_{1j}, x_{1j}\}$.

Then, add the blue edge $f = \{t_{nj}, t_j\}$ and allow the two transitions between f and any blue edge incident to t_{nj} in G_{nj}. Namely, the transitions $\{h, f\}$ and $\{h', f\}$ are allowed, where h is the blue edge $\{t_{nj}, false'_{nj}\}$ and h' is the blue edge $\{t_{nj}, z_{nj}\}$.

Finally, for any $1 \leq i < n$, let us identify t_{ij} and $s_{i+1,j}$. For any pair (e, f) of edges incident to $t_{ij} = s_{i+1,j}$ with $e \in E(G_{ij})$ and $f \in E(G_{i+1,j})$, the transition $\{e, f\}$ is allowed if e and f have the same color. Namely, let $\{a, b\}_r$ (resp. $\{a, b\}_b$) be the red (resp. blue) edge between vertices a and b. The eight allowed transitions around $t_{ij} = s_{i+1,j}$ are $\{\{z_{ij}, t_{ij}\}_b, \{s_{i+1,j}, x_{i+1,j}\}_b\}$, $\{\{z_{ij}, t_{ij}\}_r, \{s_{i+1,j}, x_{i+1,j}\}_r\}$, $\{\{false'_{ij}, t_{ij}\}_b, \{s_{i+1,j}, x_{i+1,j}\}_b\}$, $\{\{false'_{ij}, t_{ij}\}_r, \{s_{i+1,j}, x_{i+1,j}\}_r\}$, $\{\{z_{ij}, t_{ij}\}_b, \{s_{i+1,j}, \beta_{i+1,j}\}\}$, $\{\{false'_{ij}, t_{ij}\}_b, \{s_{i+1,j}, \beta_{i+1,j}\}\}$, $\{\{z_{ij}, t_{ij}\}_r, \{s_{i+1,j}, \alpha_{i+1,j}\}\}$, and $\{\{false'_{ij}, t_{ij}\}_r, \{s_{i+1,j}, \alpha_{i+1,j}\}\}$.

Case j even. First, let us add the blue edge $e = \{s_j, s_{nj}\}$ and allow the two transitions between e and any blue edge incident to s_{nj} in G_{nj}. Namely, the transitions $\{e, \{s_{nj}, \alpha_{nj}\}\}$ and $\{e, h\}$ are allowed, where h is the red edge $\{s_{nj}, x_{nj}\}$.

Then, add the red edge $f = \{t_{1j}, t_j\}$ and allow the two transitions between f and any red edge incident to t_{1j} in G_{1j}. Namely, the transitions $\{h, f\}$ and $\{h', f\}$ are allowed, where h is the red edge $\{t_{1j}, false'_{1j}\}$ and h' is the red edge $\{t_{1j}, z_{1j}\}$.

Finally, for any $1 < i \leq n$, let us identify t_{ij} and $s_{i-1,j}$. For any pair (e, f) of edges incident to $t_{ij} = s_{i-1,j}$ with $e \in E(G_{ij})$ and $f \in E(G_{i-1,j})$, the transition $\{e, f\}$ is allowed if e and f have the same color. Namely, let $\{a, b\}_r$ (resp. $\{a, b\}_b$) be the red (resp. blue) edge between vertices a and b. The eight allowed transitions around $t_{ij} = s_{i-1,j}$ are $\{\{z_{ij}, t_{ij}\}_b, \{s_{i-1,j}, x_{i-1,j}\}_b\}$, $\{\{z_{ij}, t_{ij}\}_r, \{s_{i-1,j}, x_{i-1,j}\}_r\}$, $\{\{false'_{ij}, t_{ij}\}_b, \{s_{i-1,j}, x_{i-1,j}\}_b\}$, $\{\{false'_{ij}, t_{ij}\}_r, \{s_{i-1,j}, x_{i-1,j}\}_r\}$, $\{\{z_{ij}, t_{ij}\}_b, \{s_{i-1,j}, \beta_{i-1,j}\}\}$, $\{\{false'_{ij}, t_{ij}\}_b, \{s_{i-1,j}, \beta_{i-1,j}\}\}$, $\{\{z_{ij}, t_{ij}\}_r, \{s_{i-1,j}, \alpha_{i-1,j}\}\}$, and $\{\{false'_{ij}, t_{ij}\}_r, \{s_{i-1,j}, \alpha_{i-1,j}\}\}$.

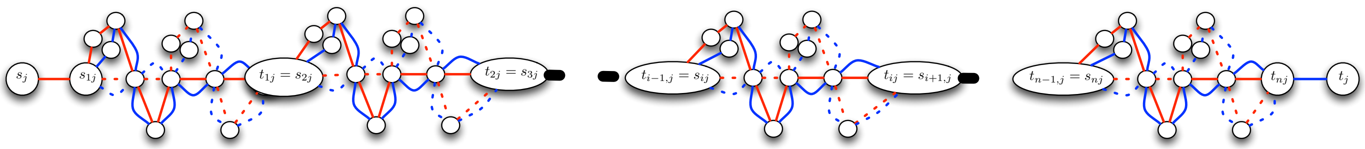

Figure 2: Clause Gadget Row_j when j odd.

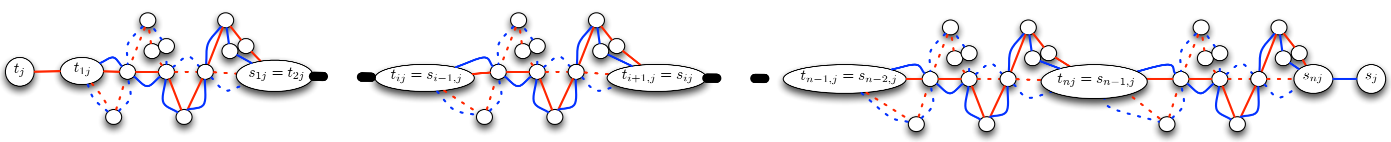

Figure 3: Clause Gadget Row_j when j even.

Let $\mathcal{F}_j$ be the set of forbidden transitions in Row_j. The subgraph Row_j is depicted in Figures 2 and 3. The key property of Row_j relates to the structure of $\mathcal{F}_j$-valid paths from s_j to t_j.

Claim 4. *Any $\mathcal{F}_j$-valid path P from s_j to t_j in Row_j consists of the concatenation of:*

Case j odd. *the red edge $\{s_j, s_{1j}\}$, then the concatenation of $\mathcal{F}_{ij}$-valid paths from s_{ij} to t_{ij} in G_{ij}, for $1 \leq i \leq n$ in this order (from $i = 1$ to n), and finally the blue edge $\{t_{nj}, t_j\}$;*

Case j even. *the blue edge $\{s_j, s_{nj}\}$, then the concatenation of $\mathcal{F}_{ij}$-valid paths from s_{ij} to t_{ij} in G_{ij}, for $1 \leq i \leq n$ in the reverse order (from $i = n$ to 1), and finally the red edge $\{t_{1j}, t_j\}$.*

By the previous claim, for any $\mathcal{F}_j$-valid path P from s_j to t_j, the colors of the first and last edges differ. Hence, by Claim 3 and the definition of the allowed transitions between two mini-gadgets:

Claim 5. *Any $\mathcal{F}_j$-valid path P from s_j to t_j must use a green edge in a mini-gadget G_{ij} for some $1 \leq i \leq n$.*

Main graph G. We are now ready to define the main graph G and associated set of forbidden transitions $\mathcal{F}$. As already said, G looks like a grid in which the subgraphs Row_j, $1 \leq j \leq m$, correspond to the rows.

More formally, let G be the graph obtained from a copy of Row_j for any $1 \leq j \leq m$ and the corresponding set of forbidden transitions $\mathcal{F}_j$, as follows.

For each $1 \leq j < m$, let us identify the vertices t_j and s_{j+1} and add the pair of the 2 edges incident to $t_j = s_{j+1}$ in the set of allowed transitions.

Afterwards, for each $1 \leq j < m$ and for each $1 \leq i \leq n$, let us identify the vertices $true_{i,j+1}$ and $false'_{ij}$ on the one hand, and the vertices $true'_{ij}$ and $false_{i,j+1}$ on the other hand. No new allowed transitions are added at this step.

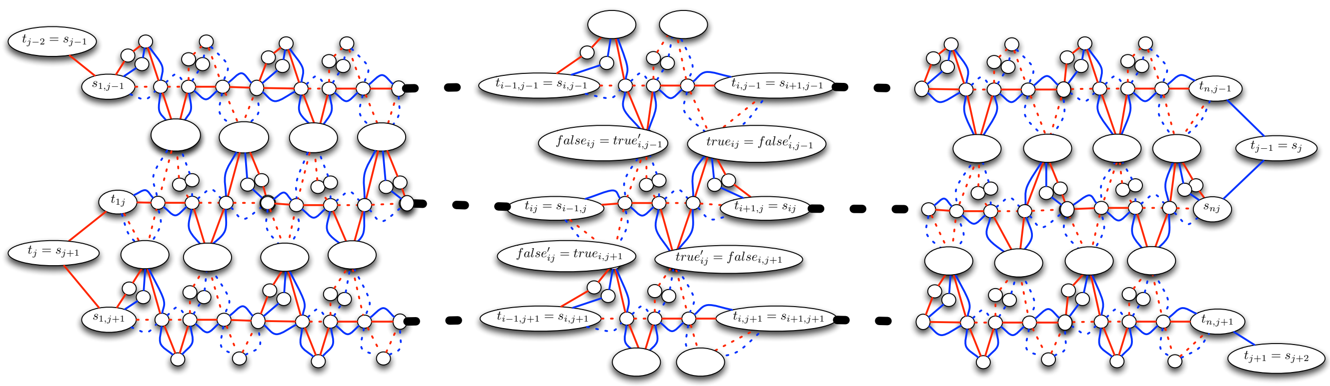

Figure 4: Part of graph G, combining Row_{j-1}, Row_j and Row_{j+1} for j even.

The graph G is depicted in Figure 4. Because no transitions are allowed from one subgraph Row_j to another except the one going through t_j, we have the following:

Claim 6. *Any $\mathcal{F}$-valid path P from s_1 to t_m in G consists of the concatenation of $\mathcal{F}_j$-valid paths from s_j to t_j in Row_j from $j = 1$ to m.*

The next claim is crucial since it implies that any $\mathcal{F}$-valid s_1t_m-path in G defines an assignment of the variables $X_1, \cdots, X_n$.

Claim 7. *Let P be an $\mathcal{F}$-valid s_1t_m-path in G. Then, for any $1 \leq i \leq n$, either*

- *for any $1 \leq j \leq m$, the subpath of P between s_{ij} and t_{ij} passes through $true_{ij}$ and $true'_{ij}$ and does not pass through $false_{ij}$ nor $false'_{ij}$, or*
- *for any $1 \leq j \leq m$, the subpath of P between s_{ij} and t_{ij} passes through $false_{ij}$ and $false'_{ij}$ and does not pass through $true_{ij}$ nor $true'_{ij}$.*

Proof. By Claims 4 and 6, for any $1 \leq i \leq n$ and any $1 \leq j \leq m$, there is a subpath P_{ij} of P that goes from s_{ij} to t_{ij}. Moreover, the paths P_{ij} are pairwise vertex-disjoint.

For $1 \leq i \leq n$, by Claim 2, P_{i1} either passes through $true_{i1}$ and $true'_{i1}$, or through $false_{i1}$ and $false'_{i1}$. Let us assume that we are in the first case (the second case can be handled symmetrically). We prove by induction on $j \leq m$ that P_{ij} passes through $true_{ij}$ and $true'_{ij}$ and does not pass through $false_{ij}$ nor $false'_{ij}$.

Indeed, if P passes through $true_{ij} = false'_{i,j+1}$ and $true'_{ij} = false_{i,j+1}$, then $P_{i,j+1}$ cannot use $false_{i,j+1}$ nor $false'_{i,j+1}$ since P_{ij} and $P_{i,j+1}$ are vertex-disjoint. By Claim 2, $P_{i,j+1}$ passes through $true_{i,j+1}$ and $true'_{i,j+1}$. $\diamond$

Note that $(G, \mathcal{F})$ can be constructed in polynomial-time. Moreover, G is clearly planar with maximum degree 8. Hence, the next claim allows to prove the lemma 1.

Claim 8. *Φ is satisfiable if and only if there is an $\mathcal{F}$-valid s_1t_m-path in G.*

Proof. Let φ be a truth assignment which satisfies Φ. We can build an $\mathcal{F}$-valid s_1t_m-path in G as follows. For each row $1 \leq j \leq m$, we build a path P_j from s_i to t_j by concatenating the paths P_{ij}, $1 \leq j \leq m$, which are built as follows. Among the variables that appear in C_j, let X_k be the variable with the smallest index, which satisfies the clause.

- For $1 \leq i < k$, if $\varphi(X_i) = true$, then $P_{ij} = \mathcal{R}T_{ij}$ ($P_{ij} = \mathcal{B}T_{ij}$) if j is odd (even), respectively. If $\varphi(X_i) = false$, then $P_{ij} = \mathcal{R}F_{ij}$ if j is odd, and $P_{ij} = \mathcal{B}F_{ij}$ if j is even.
- If $\varphi(X_k) = true$, then $P_{ij} = \mathcal{RB}T_{ij}$ if j is odd, and $P_{ij} = \mathcal{BR}T_{ij}$ if j is even. If $\varphi(X_k) = false$, then $P_{ij} = \mathcal{RB}F_{ij}$ if j is odd, and $P_ij = \mathcal{BR}F_{ij}$ if j is even.
- For $k < i \leq n$, if $\varphi(X_i) = true$, then $P_{ij} = \mathcal{B}T_{ij}$ if j is odd, and $P_ij = \mathcal{R}T_{ij}$ if j is even. If $\varphi(X_i) = false$, then $P_{ij} = \mathcal{B}F_{ij}$ if j is odd, and $P_{ij} = \mathcal{R}F_{ij}$ otherwise.

The path P obtained from the concatenation of paths P_j for $1 \leq j \leq m$ is an $\mathcal{F}$-valid path from s_1 to t_m.

Now let us suppose that there is an $\mathcal{F}$-valid path P from s_1 to t_m. According to Claim 7, for any $1 \leq i \leq n$, for any $1 \leq j \leq m$, P passes through $true_{ij}$ and $true'_{ij}$ or for any $1 \leq j \leq m$, P passes through $false_{ij}$ and $false'_{ij}$. Let us then consider the truth assignment φ of Φ such that for each $1 \leq i \leq n$:

- If P uses $true_{ij}$ and $true'_{ij}$ in all rows $1 \leq j \leq m$, then $\varphi(X_i) = true$.
- If P uses $false_{ij}$ and $false'_{ij}$ in all rows $1 \leq j \leq m$, then $\varphi(X_i) = false$.

Thanks to Claim 7, φ is a valid truth assignment. We need to prove that φ satisfies Φ. According to Claims 6, for each row $1 \leq j \leq m$, P contains an $\mathcal{F}_j$-valid path P_j from s_j to t_j. Each path P_j uses a green edge as stated by Claim 3. With respect to the possible ways to use a green edge which are stated in Claim 2, the use of a green edge in P_j forces P_j (and hence P) to use, for a variable X_i that appears in C_j, the vertices $true_{ij}$ and $true'_{ij}$ ($false_{ij}$ and $falce'_{ij}$) if X_i appears positively (negatively) in C_j, respectively. This means that for each clause C_j, for one of the variables that appear in C_j which we denote X_i, $\varphi(X_i) = true$ ($\varphi(X_i) = false$) if X_i appears positively (negatively) in C_j, respectively. Thus, the truth assignment φ satisfies Φ. ◇

□

Corollary 1. *The problem of finding a path avoiding forbidden transitions is NP-hard in planar graphs with maximum degree* 8.

Proof. Let $G_m = (E_m, V_m)$ be a multi-graph, $\mathcal{F}_m$ be a set of forbidden transitions and $s, t \in V_m$. Let $G = (V, E)$ be the graph obtained from G_m by subdividing once each edge. That is, $V = V \cup \{v_e \mid e \in E_m\}$ and $E = \{\{u, v_e\}, \{w, v_e\} \mid e = \{u, w\} \in E_m\}$. Moreover, let $\mathcal{F}$ be the set of forbidden transitions in G that corresponds to $\mathcal{F}_m$, i.e., $\mathcal{F} = \{\{\{v_e, u\}, \{u, v_f\}\} \mid \{e, f\} \in \mathcal{F}_m, u \in e \cap f\}$.

Clearly, there is a $\mathcal{F}_m$-valid st-path in G_m if and only if there is a $\mathcal{F}$-valid st-path in G.

Finally, planarity and maximum degree are preserved by our transformation. □

Applying the above transformation to the multi-graph used in the proof of lemma 1 will give a graph G that verifies the following properties:

- G is planar.
- Each vertex of G has either degree 8 or degree ≤ 4. In fact, there are also vertices of degree 5 which we transform to vertices of degree 3 as follows. For vertices s_{1j} (resp. t_{1j}) where j is odd we delete the 2 blue edges incident to s_{1j} (resp. t_{1j}). For vertices s_{nj} (resp. t_{1j}) where j is even, we delete the 2 red edges incident to s_{nj} (resp. t_{1j}). This transformation does not affect the reduction or the proof.
- According to its forbidden transitions and to its disposition in the planar embedding (depicted in Figure 4), a vertex v of G of degree 8 has one of three following types:

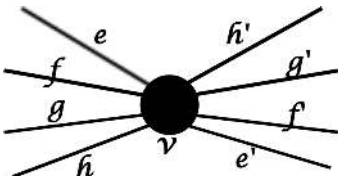

(a) A vertex of degree 8 and allowed transtions $A(v) = \{\{e, e'\}, \{f, f'\}, \{g, g'\}, \{h, h'\}\}$ (edges are ordered as in the planar embedding of G)

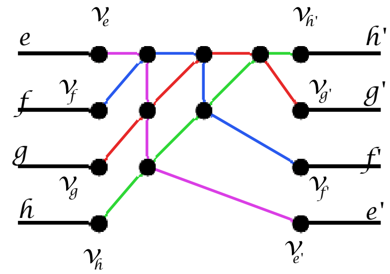

(b) Gadget D_v: the paths $P_{ee'}$, $P_{ff'}$,$P_{gg'}$, and $P_{hh'}$ are respectively the pink, blue, red and green paths. Transitions around vertices v_i and transitions of paths $P_{ee'}$, $P_{ff'}$,$P_{gg'}$, and $P_{hh'}$ are allowed

Figure 5: Type 1

Type 1: The edges incident to v are $\omega(v) = \{e, e', f, f', g, g', h, h'\}$ and the allowed transitions around v are $A(v) = \{\{e, e'\}, \{f, f'\}, \{g, g'\}, \{h, h'\}\}$. The edges of v in the planar embedding are as depicted in Figure 5a. (v is a vertex of type x_{ij}, y_{ij} or z_{ij} in the graph of Figure 4)

Type 2: The edges incident to v are $\omega(v) = \{e, e', f, f', g, g', h, h'\}$ and the allowed transitions around v are $A(v) = \{\{e, e'\}, \{f, f'\}, \{g, g'\}, \{h, h'\}\}$. The edges of v in the planar embedding are as depicted in Figure 6a. (v is a vertex of type $true_{ij}$, $true'_{ij}$, $false_{ij}$, or $false'_{ij}$ in the graph of Figure 4)

Type 3: The edges incident to v are $\omega(v) = \{e, e', f, f', g, g', h, h'\}$ and the allowed transitions around v are $A(v) = \{\{e, e'\}, \{e, f'\}, \{f, f'\}, \{f, e'\}, \{g, g'\}, \{g, h'\}, \{h, h'\}, \{h, g'\}\}$. The edges of v in the planar embedding are as depicted in Figure7a. (v is a vertex of type s_{ij} in the graph of Figure 4)

Let us denote by $\mathcal{G}_8$ the class of graphs verifying the three properties above. We can safely say that PAFT is NP-complete in this class of graph. We will use this fact to prove the following lemma.

Lemma 2. *The Problem of finding a path avoiding forbidden transitions is NP-hard in planar graphs with maximum degree* 4*.*

Proof. Let G be a graph of the class $\mathcal{G}_8$ and $\mathcal{F}$ its set of forbidden transitions. To prove the lemma, we are going to replace each vertex v of degree 8 in G with a gadget D_v. After replacing all vertices of degree 8, we will obtain a graph G' of maximum degree 4 and a new set of forbidden transitions $\mathcal{F}'$ such that finding an $\mathcal{F}-$valid path from s to t in G is equivalent to finding an $\mathcal{F}'$-valid path from s to t in G'. Let v be a node of degree 8 of G. D_v is constructed according to the type of v as follows:

Type 1 In this case, D_v is constructed as follows. For each $i \in \omega(v)$, a vertex v_i is created. For each $\{i, j\} \in A(v)$, vertices v_i and v_j are linked with a path P_{ij} of length four . The four paths P_{ij} are pairwise intersecting in distinct vertices as illustrated in Figure 5b. The allowed transitions in D_v are the transitions of the paths P_{ij}. Now to replace v with D_v in G, we do the following: each edge $i \in \omega(v)$ of G is linked to vertex v_i of D_v. The gadget D_v is planar, and edges $i \in \omega(v)$ are connected to it in the same "order" they were connected to v in the planar embedding of G as illustrated in Figure 5. Note the gadget D_v cannot be crossed twice with the same path, otherwise the path is not simple. Moreover, D_v can be crossed if and only if the edges used to enter and leave form an allowed transition around v.

Type 2 In this case, D_v is constructed as follows. For each $i \in \omega(v)$, a vertex v_i is created. For each $\{i, j\} \in A(v)$, vertices v_i and v_j are linked with a path P_{ij} of length 7 . Each two of the four paths P_{ij} intersect in two different vertices as illustrated in Figure 6b. The allowed transitions in D_v are the transitions of the paths P_{ij}. Now to replace v with D_v in G, we do the following: each edge $i \in \omega(v)$

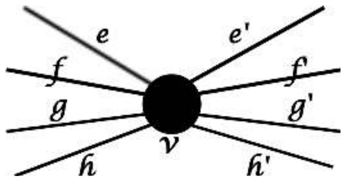

(a) A vertex of degree 8 and allowed transtions $A(v) = \{\{e,e'\},\{e,f'\},\{f,f'\},\{f,e'\},\{g,g'\},\{g,h'\},\{h,h'\},\{h,g'\}\}$ (edges are ordered as in the planar embedding of G)

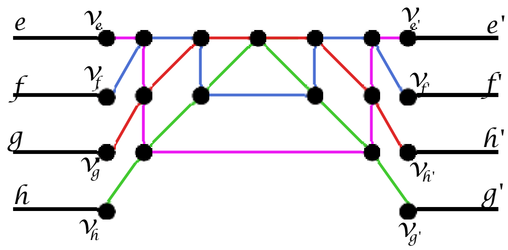

(b) Gadget D_v: the paths $P_{ee'}$, $P_{ff'}$,$P_{gg'}$, and $P_{hh'}$ are respectively the pink, blue, red and green paths. Transitions around vertices v_i and transitions of paths $P_{ee'}$, $P_{ff'}$,$P_{gg'}$, and $P_{hh'}$ are allowed

Figure 6: Type 2

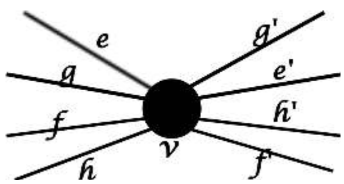

(a) A vertex of degree 8 and allowed transtions $A(v) = \{\{e,e'\},\{e,f'\},\{f,f'\},\{f,e'\},\{g,g'\},\{g,h'\},\{h,h'\},\{h,g'\}\}$ (edges are ordered as in the planar embedding of G)

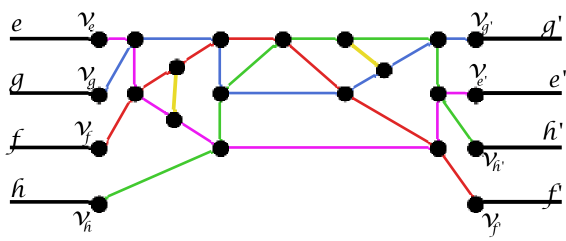

(b) Gadget D_v: the paths $P_{ee'}$, $P_{ff'}$,$P_{gg'}$, and $P_{hh'}$ are respectively the pink, red, blue and green paths. transitions around vertices v_i, transitions of the paths $P_{ee'}$, $P_{ff'}$,$P_{gg'}$, and $P_{hh'}$, and transitions containing the yellow edge are allowed

Figure 7: Type 3

of G is linked to vertex v_i of D_v. The gadget D_v is planar, and edges $i \in \omega(v)$ are connected to it in the same "order" they were connected to v in the planar embedding of G as illustrated in Figure 6. Note that the gadget D_v cannot be crossed twice with the same path, otherwise the path is not simple. Moreover, D_v can be crossed if and only if the edges used to enter and leave form an allowed transition around v.

Type 3 In this case, D_v is constructed as follows. For each $\{i,i'\} \in A(v)$, vertices v_i and v_j are linked with a path P_{ij} of length 7. Each two of the paths P_{ij} intersect twice in distinct vertices as illustrated in Figure 7b. Furthermore, we add two edges linking the paths $P_{ee'}$ and $P_{ff'}$, and $P_{gg'}$ and $P_{hh'}$, respectively. Now to replace v with D_v in G, we do the following: each edge $i \in \omega(v)$ of G is linked to vertex v_i of D_v. The gadget D_v is planar, and edges $i \in \omega(v)$ are connected to it in the same "order" they were connected to v in the planar embedding of G as illustrated in Figure 7. Note that the gadget D_v cannot be crossed twice with the same path, otherwise the path is not simple. Moreover, D_v can be crossed if and only if the edges used to enter and leave form an allowed transition around v.

The graph G' is the one obtained from G after replacing vertices of degree 8 with the gadgets described above. The set of forbidden transitions $\mathcal{F}'$ consists of the transitions of the set $\mathcal{F}$ and the forbidden transitions of the gadgets D_v as described above. The maximum degree of G' is 4 and G' is planar.

Let us now suppose that there is an $\mathcal{F}$-valid path P from s to t in G. Let P' be the st-path of G' constructed as follows: P' uses all edges used by P. Furthermore, if P uses a degree 8 vertex of type 1 or 2 with a transition $\{e,e'\}$ then P' uses e, subpath $P_{ee'}$, and e'. If P uses a degree 8 vertex of type 3 with transition $\{e,e'\}$ (or $\{e,f'\}$), then P' uses e, e' and the subpath $P_{ee'}$ (e,f', and the subpath $P_{ef'}$ which is the concatenation of a subpath of path $P_{ee'}$, a yellow edge and a subpath of path $P_{ee'}$), respectively. The path P' is $\mathcal{F}'$-valid.

Now, let us suppose that there is an $\mathcal{F}'$-valid path P' from s to t in G'. If P' only uses edges from G, then it can be considered as an $\mathcal{F}$-valid path from s to t in G. If P' uses an edge that is not in G, then P' crosses one of the gadgets D_v. As we have specified above, the gadgets D_v can only be crossed in specified ways that ensure that the edges used to enter and leave the gadget form an allowed transition. We can then remove the edges of P' that do not belong to G to obtain an $\mathcal{F}$-valid path P in G. For any v of degree 8, the path P does not pass twice through v since gadget D_v in G' cannot be crossed twice by the same path. □

Theorem 1. *The Problem of finding a path avoiding forbidden transitions is NP-complete in grids.*

Proof. To prove that PAFT is NP-complete in grids, we use the notion of planar grid embedding. As defined in [12], a planar grid embedding of a graph G is a mapping Q of G into a grid such that:

- Q maps each vertex of G into a distinct vertex of the grid.
- Q maps edge e of G into a path of the grid $Q(e)$ whose endpoints are mappings of vertices linked by e.
- For every pair $\{e, e'\}$ of edges of G, the corresponding paths $Q(e)$ and $Q(e')$ have no points in common, except, possibly, the endpoints.

The following theorem has been proven in [13]:

Theorem 2 ([13])**.** *Let $G = (V, E)$ be a planar graph such that $|V| = n$ and $\Delta \leq 4$, a planar grid embedding of G in a grid of size at most $9n^2$ can be found in polynomial-time.*

Now, to prove Theorem 1, let us consider an instance of the problem of PAFT in a planar graph $G = (V, E)$ of maximum degree at most 4 with a set of allowed transitions $\mathcal{A}$. Let Q be a grid planar embedding of G into a grid K of size at most $O(|V|^2)$. We assume that except $Q((s, t))$ no other path between s and t consists of only one edge in K, if it is not the case, we increase the size of K to make sure that the distance between s and t is at least 2.Finding a path avoiding forbidden transitions between two nodes s and t in G with the set $\mathcal{A}$ is equivalent to finding a path avoiding forbidden transitions between the nodes $Q(s)$ and $Q(t)$ in K with the set of allowed transitions $\mathcal{A}'$ defined as follows:

- For each $e \in E$, all the transitions in the path $Q(e)$ are allowed.
- For each $\{e, e'\} \in \mathcal{A}$, the pair of edges of $Q(e)$ and $Q(e')$, which share a vertex, is an allowed transition.

Since we have proved in Lemma 2 that CFT is NP-complete in planar graphs with maximum degree 4, we deduce that CFT is NP-complete in grids. □

3 An algorithm for graphs with bounded treewidth

A *tree-decomposition* of a graph [10] is a way to represent G by a family of subsets of its vertex-set organized in a tree-like manner and satisfying some connectivity property. The *treewidth* of G measures the proximity of G to a tree. More formally, a tree decomposition of $G = (V, E)$ is a pair $(T, \mathcal{X})$ where $\mathcal{X} = \{X_t | t \in V(T)\}$ is a family of subsets, called *bags*, of V, and T is a tree, such that:

- $\bigcup_{t \in V(T)} X_t = V$;
- for any edge $uv \in E$, there is a bag X_t (for some node $t \in V(T)$) containing both u and v;
- for any vertex $v \in V$, the set $\{t \in V(T) | v \in X_t\}$ induces a subtree of T.

The *width* of a tree-decomposition $(T, \mathcal{X})$ is $max_{t \in V(T)} |X_t| - 1$ and its *size* is order $|V(T)|$ of T. The treewidth of G, denoted by $tw(G)$, is the minimum width over all possible tree-decompositions of G.

Computing optimal tree-decomposition - i.e., with width $tw(G)$ - is NP-complete in the class of general graphs G [2]. For any fixed $k \geq 1$, Bodlaender designed an algorithm that computes, in time $O(k^{k^3} n)$, a tree-decomposition of width k of any n-node graph with treewidth at most k [4]. Very recently, a single-exponential (in k) algorithm has been proposed that computes a tree-decomposition with width at most $5k$ in the class of graphs with treewidth at most k [5].

Many NP-hard problems can be solved in polynomial time in the class of graphs of bounded treewidth using dynamic programming algorithms. For instance, the Maximum Independent Set, the 3-Coloring, the Vertex Cover, and the Hamiltonian cycle are all FPT when parametrized by the treewidth of the graph.

In this section, we adapt the algorithm that solves the Hamiltonian cycle [we need references] on graphs with bounded treewidth to find a path avoiding forbidden transitions. We start by introducing more definitions.

Definition 1. *A rooted tree decomposition $((T, \mathcal{X}), r)$ of G is nice if for every node $u \in V(T)$:*

- *u has no children and $|X_u| = 1$ (u is called a leaf node), or*
- *u has one child v with $X_u \subset X_v$ and $|X_u| = |X_v| - 1$ (u is called a forget node),or*
- *u has one child v with $X_v \subset X_u$ and $|X_u| = |X_v| + 1$ (u is called an introduce node), or*
- *u has two children v and w with $X_u = X_v = X_w$ (u is called a join node.).*

Lemma 3. *When given a tree decomposition of width w of G, in polynomial time we can construct a nice tree decomposition $(T, \mathcal{X})$ of G of width k, with $|V(T)| \in O(kn)$, where $n = |V(G)|$.*

We use the notion of nice tree decomposition and adapt the dynamic programming algorithm for finding a Hamiltonian cycle in a graph to prove Theorem 3.

Theorem 3. *The problem of finding a path avoiding forbidden transition can be solved in polynomial time in graphs with bounded treewidth. In particular, there exists an algorithm that finds the shortest path avoiding forbidden transitions between two nodes in a graph of treewidth k in time $O((3\Delta(k+1))^{2k+O(1)} n))$*

Proof. Let $G = (V, E)$ be a graph with bounded treewidth k, $\mathcal{F} \subseteq E \times E$ a set of forbidden transitions (and $\mathcal{A} \subseteq E \times E$ the set of allowed transitions), and s and t two vertices of V. We would like to find the shortest path P from s to t avoiding the forbidden transitions $\mathcal{F}$.

Let $G_{e,f}$ such that e and f are edges incident to s and t, respectively, be the graph obtained from G, by deleting all edges incident to s and t except for e and f. Finding the shortest path avoiding forbidden transitions $\mathcal{F}$ from s to t in G is equivalent to finding the shortest path among all shortest paths avoiding forbidden transitions $\mathcal{F}$ from s to t in $G_{e,f}$, for each possible pair e, f. In the following we will present how to solve the CFT problem in $G_{e,f}$. To obtain the solution in G, we will need to repeat the algorithm at most Δ^2 times.

Let $(T, \mathcal{X})$ be a nice tree-decomposition of width k of $G_{e,f}$. We assume that s appears in one introduce bag and t in two bags, a leaf and its introduce parent. We root T at the node containing s. Let $G[A]$ be the subgraph of $G_{e,f}$ induced by the set of vertices A. . For each $u \in V(T)$ we denote by X_u,T_u and V_u the vertices of the bag corresponding to u, the subtree of T rooted at u, and the vertices of the bags corresponding to the nodes of T_u , respectively.

If there exists an $\mathcal{F}$-valid path P from s to t, then the intersection of this path with $G[V_u]$ for a node $u \in T$ is a set of paths (avoiding forbidden transitions) each having both endpoints in X_u. If $t \in V_u$, then one of the paths has only one endpoint in X_u.

With respect to the parts of path P that are in $G[V_u]$, vertices in X_u can be partitioned into three subsets: X_u^0, X_u^1, and X_u^2 which are the vertices of degree 0, 1 and 2 in $P \cap G[V_u]$, respectively. Furthermore, a matching M of X_u^1 decides which vertices are endpoints of the same subpath and a set of edges S defines which edges incident to X_u^1 are in P. For each node $u \in T$ and each subproblem $(X_u^0, X_u^1, X_u^2, M, S)$ where

(X_u^0, X_u^1, X_u^2) is a partition of X_u, M is a matching of X_u^1 and S is a set of edges incident to the vertices of X_u^1, we need to see if there exists a set of paths avoiding forbidden transitions in V_u such that their endpoints are exactly X_u^1 according to the matching M, they contain the edges of S and the vertices of X_u^2 and they do not contain any vertex of X_u^0. For the case where $t \in V_u$, we will need to check the possible matchings of each subset of X_u^1 of size $|X_u^1| - 1$. For each node, we will need to solve at most $3^{k+1}(k+1)^{k+1}\Delta^{k+1}$ subproblems; there are at most 3^{k+1} possible partitions of the vertices of X_u into the 3 different sets, $(k+1)^{k+1}$ possible matchings for a set of $k+1$ elements and Δ possible edges for each element of X_u^1.

Let us see how to solve a problem $(X_u^0, X_u^1, X_u^2, M, S)$ at a node u supposing that all the problems at its descendants have been solved:

- If u is a leaf, then $X_u = \{a\}$. The only problem that has a solution is $(X_u^0 = \{a\}, X_u^1 = \emptyset, X_u^2 = \emptyset, M = \emptyset, S = \emptyset)$.
- If u is a forget node, let v be the child of u. We have $X_u = X_v \setminus a$. We can distinguish two cases:
  - If $a \neq t$, then the problem $(X_u^0, X_u^1, X_u^2, M, S)$ has a solution if and only if one of the problems $(X_u^0 \cup \{a\}, X_u^1, X_u^2, M, S)$ and $(X_u^0, X_u^1, X_u^2 \cup \{a\}, M, S)$ at node v has a solution.
  - If $a = t$, then the problem $(X_u^0, X_u^1, X_u^2, M, S)$ has a solution if and only if the problem $(X_u^0, X_u^1 \cup \{a\}, X_u^2, M, S)$ at v has a solution.
- If u is an introduce node, let v be the child of u. We have $X_u = X_v \cup a$ (all neighbors of a in V_u are in X_u). Note that $a \neq t$ since t appears in a forget node and its introduce parent. In this case we proceed as follows.
  - If $a \in X_u^0$, then solving $(X_u^0, X_u^1, X_u^2, M, S)$ at u is equivalent to solving $(X_u^0 \setminus \{a\}, X_u^1, X_u^2, M, S)$ at v.
  - if $a \in X_u^1$, let ab be the edge incident to a in S. Since all neighbors of a in V_u are in X_u, then $b \in X_u \cap X_v$. Let us consider the following cases: (still need to treat the case where $b = t$)
    - If $b = t$, then the only problem that has a solution at u is $(X_u \setminus \{a, t\}, \{a, t\}, \emptyset, \{(a, t)\}, \{(a, t)\})$. To solve it, we need to check at v the solution of the problem $(X_u \setminus \{a\}, \emptyset, \emptyset, \emptyset, \emptyset)$.
    - If $b \in X_u^1$ $(b \neq t)$, (the problem has a solution only if $(a, b) \in M$ and the edge incident to b in S is ab) then check at v the solution of the problem $(X_u^0 \cup \{b\}, X_u^1 \setminus \{a, b\}, X_u^2, M', S')$ where $M' = M \setminus (a, b)$ and $S' = S \setminus ab$.
    - If $b \in X_u^2$ $(b \neq t)$, then check at v the solution of the problem $(X_u^0, X_u^1 \setminus \{a\} \cup \{b\}, X_u^2 \setminus \{b\}, M', S')$ where $M' = M \setminus (a, h) \cup (b, h)$ and S' contains the set S minus the edge ab plus an edge incident to b that forms an allowed transition with edge ba (there are at most Δ such problems).
  - If $a \in X_u^2$, then for every two neighbors b and c of a in X_u such that (ba, ac) is an allowed transition do the following.
    - If $b \in X_u^1$ and $c \in X_u^1$, then check the solution at v of the problem $(X_u^0 \cup \{b, c\}, X_u^1 \setminus \{b, c\}, X_u^2 \setminus \{a\}, M', S')$ where $M' = M \setminus (b, c)$ and remove ab and bc from S to obtain S'.
    - If $b \in X_u^2$ and $c \in X_u^2$, then check the solution at v of the problem $(X_u^0, X_u^1 \cup \{b, c\}, X_u^2 \setminus \{a, b, c\}, M', S')$ where $M' = M \cup \{bh, ch'\} \setminus hh'$ (bc should not be in the matching) and to obtain S', add to S two edges incident to b and c and forming allowed transitions with ab and ac, respectively (there are $\frac{k+1}{2}$ possible choices for hh' and Δ^2 possible choices for the two edges to add to S).
    - If $b \in X_u^1$ and $c \in X_u^2$, then check the solution at v of the problem $(X_u^0 \cup \{b\}, X_u^1 \setminus \{b\} \cup \{c\}, X_u^2 \setminus \{a, c\}, M', S')$ where $M' = M \setminus bh \cup ch$ and to obtain S' remove ab from S and add an edge incident to c that forms an allowed transition with ca. (There are Δ possibilities).

  Note that the number of pairs of neighbors of a to consider are of order of k^2.

- If u is a join node, let v and w be its children. For any two subproblems at v and w we check if the union of the two solutions is a solution for $(X_u^0, X_u^1, X_u^2, M, S)$ of node u. (At most $(3^{k+1}k + 1^{k+1}\Delta^{k+1})^2$ possibilities).

At the node containing s, we only need to solve subproblems where s and t are of degree 1 and all other vertices have either degree 2 or 0.

To find the shortest path, one has to choose, whenever having a choice between different solutions for a subproblem at a node, the solution with the minimum number of edges.

□

References

[1] Ahmed, M., and Lubiw, A. Shortest paths avoiding forbidden subpaths. *Networks 61*, 4 (2013), 322–334.

[2] Arnborg, S., Corneil, D. G., and Proskurowski, A. Complexity of finding embeddings in a k-tree. *SIAM J. Algebraic Discrete Methods 8*, 2 (Apr. 1987), 277–284.

[3] Bernstein, G., Lee, Y., Gavler, A., and Martensson, J. Modeling wdm wavelength switching systems for use in gmpls and automated path computation. *Optical Communications and Networking, IEEE/OSA Journal of 1*, 1 (June 2009), 187–195.

[4] Bodlaender, H. L. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput. 25*, 6 (1996), 1305–1317.

[5] Bodlaender, H. L., Drange, P. G., Dregi, M. S., Fomin, F. V., Lokshtanov, D., and Pilipczuk, M. A $o(c^k n)$ 5-approximation algorithm for treewidth. *CoRR abs/1304.6321* (2013).

[6] Chen, Y., Hua, N., Wan, X., Zhang, H., and Zheng, X. Dynamic lightpath provisioning in optical wdm mesh networks with asymmetric nodes. *Photonic Network Communications 25*, 3 (2013), 166–177.

[7] Gourvès, L., Lyra, A., Martinhon, C. A., and Monnot, J. Complexity of trails, paths and circuits in arc-colored digraphs. *Discrete Applied Mathematics 161*, 6 (2013), 819 – 828.

[8] Gutiérrez, E., and Medaglia, A. Labeling algorithm for the shortest path problem with turn prohibitions with application to large-scale road networks. *Annals of Operations Research 157*, 1 (2008), 169–182.

[9] Pugliese, L. D. P., and Guerriero, F. Shortest path problem with forbidden paths: The elementary version. *European Journal of Operational Research 227*, 2 (2013), 254 – 267.

[10] Robertson, N., and Seymour, P. D. Graph minors. II. algorithmic aspects of tree-width. *J. Algorithms 7*, 3 (1986), 309–322.

[11] Szeider, S. Finding paths in graphs avoiding forbidden transitions. *Discrete Applied Mathematics 126*, 2 - 3 (2003), 261 – 273.

[12] Tamassia, R. On embedding a graph in the grid with the minimum number of bends. *SIAM J. Comput. 16*, 3 (June 1987), 421–444.

[13] Valiant, L. G. Universality considerations in VLSI circuits. *IEEE Transactions on Computers 30*, 2 (1981), 135–140.

[14] Villeneuve, D., and Desaulniers, G. The shortest path problem with forbidden paths. *European Journal of Operational Research 165*, 1 (2005), 97 – 107.