Discrete Applied Mathematics 312 (2022) 86–105

Contents lists available at ScienceDirect

Discrete Applied Mathematics

journal homepage: www.elsevier.com/locate/dam

Computing the zig-zag number of directed graphs

Mitre C. Dourado [a], Celina M.H. de Figueiredo [a], Alexsander A. de Melo [a,*], Mateus de Oliveira Oliveira [c], Uéverton S. Souza [b]

[a] Federal University of Rio de Janeiro, Rio de Janeiro, Brazil
[b] Fluminense Federal University, Niterói, Brazil
[c] University of Bergen, Bergen, Norway

ARTICLE INFO

Article history:
Received 29 June 2020
Received in revised form 12 June 2021
Accepted 15 September 2021
Available online 2 October 2021

Keywords:
Zig-zag number
Directed graphs
Directed width measure
NP-completeness
Computational complexity

ABSTRACT

The notion of zig-zag number was introduced as an attempt to provide a unified algorithmic framework for directed graphs. Nevertheless, little was known about the complexity of computing this directed graph invariant. We prove that deciding whether a directed graph has zig-zag number at most k is in NP for each fixed $k \geq 0$. Although for most of the natural decision problems this is an almost trivial result, settling k-ZIG-ZAG NUMBER in NP is surprisingly difficult. In addition, we prove that 2-ZIG-ZAG NUMBER is already an NP-hard problem.

1. Introduction

Structural graph parameters, such as treewidth, cutwidth and cliquewidth, have been crucial in the development of parameterized complexity theory. Indeed, many problems that are hard on general graphs become tractable when parameterized by such parameters [5,6]. However, one of the limitations of these parameters is the fact that they do not take the direction of edges into account. For instance, directed acyclic graphs (DAGs) in general have unbounded width with respect to any of the parameters mentioned above. Nevertheless, certain problems can be solved efficiently on DAGs by using straightforward algorithms. For instance, DIRECTED HAMILTONIAN PATH can be solved in linear time on DAGs with a depth-first search algorithm.

Building on this observation, Johnson, Robertson, Seymour and Thomas [12] initiated a quest for the development of width measures that explicitly take the direction of edges into consideration. In particular, they defined in [12] the notion of directed treewidth and showed that some linkage problems that are NP-hard on general directed graphs can be solved in polynomial time on directed graphs of constant directed treewidth. Additionally, a directed analog of the notion of pathwidth was also defined by Reed, Thomas, and Seymour around the same time (see for instance *cf.* [1]).

The introduction of directed treewidth and directed pathwidth motivated the development of many other width measures for directed graphs that focus on distinct algorithmic or structural properties [2,3,7,9,10,15,16]. A general algorithmic framework for directed width measures was developed in [13] with the introduction of the notion of *zig-zag* number of a directed graph, and subsequently generalized in [14] with the definition of the notion of *tree-zig-zag* number of a directed graph.

* Corresponding author.
E-mail addresses: mitre@dcc.ufrj.br (M.C. Dourado), celina@cos.ufrj.br (C.M.H. de Figueiredo), aamelo@cos.ufrj.br (A.A. de Melo), mateus.oliveira@uib.no (M. de Oliveira Oliveira), ueverton@ic.uff.br (U.S. Souza).

https://doi.org/10.1016/j.dam.2021.09.013

It was shown in [13] that if $\mathcal{G}$ is a class of directed graphs expressible by a monadic-second order logic formula φ and there is a positive integer p such that each directed graph in $\mathcal{G}$ can be cast as a union of p directed paths, then, given a decomposition of a directed graph G of zig-zag number at most k, one can count in time $f(\varphi, p, k)\cdot|G|^{\mathcal{O}(p\cdot k)}$ the number of subgraphs of G isomorphic to some member of $\mathcal{G}$, for some computable function f. Since directed path decompositions of width d can be efficiently converted into decompositions of zig-zag number $\mathcal{O}(d)$, the counting problem described above can also be solved in time $f(\varphi, p, d)\cdot|G|^{\mathcal{O}(p\cdot d)}$ on directed graphs of directed pathwidth at most d. These results were subsequently generalized in [14] to their respective counterparts for directed graphs of tree-zig-zag number at most k and of directed treewidth at most d. The results in [13] and in [14] were the first algorithmic metatheorems relating the monadic-second order logic of graphs to directed pathwidth and directed treewidth, respectively.

In a seminal paper analyzing the algorithmic potential of directed width measures, Ganian et al. [8] defined a width measure to be *algorithmically useful* if it satisfies the following properties: (1) every graph problem expressible in MSO_1 logic admits an XP-time algorithm when parameterized by the measure and (2) for each constant k, the class of graphs of width at most k is closed under taking directed topological minors. Interestingly, it was shown in [8] that, under standard complexity theoretic assumptions, any width measure satisfying properties (1) and (2) behaves essentially in the same way as the usual notion of undirected treewidth (see Theorems 6.6 and 6.7 of [8] for precise statements). We note that any directed width measure that is constant on DAGs, including zig-zag number, tree-zig-zag number and most of the width measures defined so far, fail to satisfy property (1) since 3-COLORING is MSO_1 definable and NP-complete on DAGs. Despite of this fact, zig-zag number and tree-zig-zag number have been proved to be algorithmically relevant, by establishing through the metatheorems presented in [13,14] a unified algorithmic framework to solve problems on directed graphs of low directed pathwidth and of low directed treewidth, respectively. Another interesting aspect of zig-zag number and tree-zig-zag number is the fact that they can be regarded as graph invariants with challenging theoretical open problems, from the perspectives of computational complexity and graph theory.

In fact, several complexity questions with respect to computing zig-zag number and tree-zig-zag number of a directed graph remain open. In particular, the computational complexity of the problem of determining whether a directed graph has zig-zag number at most k, even for constant k, has remained open since the introduction of this notion in [13].

We show in Section 3 that determining whether a directed graph G has zig-zag number at most k can be solved *non-deterministically* in time $|G|^{\mathcal{O}(k)}$, implying that this problem lies in NP for each fixed k. While the respective statement is almost trivial with respect to other directed width measures, such as directed pathwidth, which is known to be in P [18], our proof settling k-ZIG-ZAG NUMBER in NP turned out to be an interesting quest. This is due to the fact that the definition of zig-zag number, which we formally present in Section 2, involves the alternation of an existential and a universal quantifiers. Thus, a naive application of the definition would only lead to a Σ_2^{P}-upper bound for the problem. To circumvent this, and settle the problem in NP, our proof may be regarded as a way of redefining the property of a directed graph having zig-zag number at most k in a purely existential fashion.

On the other hand, through a polynomial-time reduction from POSITIVE NOT ALL EQUAL 3SAT, we prove in Section 4 that deciding whether a directed graph has zig-zag number at most 2 is an NP-hard problem. It is worth noting that this intractability result does not affect the applicability of the algorithmic metatheorem presented in [13], since for each $k \in \mathbb{N}$, directed path decompositions of width k can be converted efficiently into linear orderings of zig-zag number $\mathcal{O}(k)$ [13], and directed path-decompositions of width k, whenever exist, can be constructed in time $n^{\mathcal{O}(k)}$ [18].

Besides these proposed results, we analyze in Section 5 how zig-zag number and directed treewidth are related to each other. We prove that there are directed graphs of constant directed treewidth but unbounded zig-zag number. As a consequence, with the results of [14], we obtain that the family of directed graphs of constant tree-zig-zag number is strictly richer than the family of directed graphs of constant zig-zag number.

2. Preliminaries

A *directed simple graph* (or, simply *directed graph*) is a pair $G = (V, E)$ comprising a non-empty *vertex set V* and an *edge set* $E \subseteq \{(u, v) : (u, v) \in V \times V, u \neq v\}$. In what follows, we may write n or $|G|$ to denote the number of vertices of G, and we may write $V(G)$ and $E(G)$ to refer to the vertex set and to the edge set of G, respectively.

For each positive integer n, we let $[n] \doteq \{1, 2, \dots, n\}$. Let G be a directed graph on n vertices. For each bijection $\pi : V(G) \to [n]$, we let $<_\pi \subseteq V(G)\times V(G)$ be the linear order associated with π such that, for each $u, v \in V(G)$, $u <_\pi v$ if and only if $\pi(u) < \pi(v)$. Analogously, we let $>_\pi \subseteq V(G)\times V(G)$ be the linear order such that, for each $u, v \in V(G)$, $u >_\pi v$ if and only if $\pi(u) > \pi(v)$. Let $X, Y \subseteq V(G)$ be two non-empty sets. We write $X <_\pi Y$ to denote that $u <_\pi v$ for each $u \in X$ and each $v \in Y$. We define $X >_\pi Y$ similarly. Moreover, for any non-empty set $X \subseteq V(G)$, we write $\min_\pi X$ to denote the unique vertex $u \in X$ such that $\{u\} <_\pi X \setminus \{u\}$. We define $\max_\pi X$ similarly.

The zig-zag number of a directed graph. Let n be a positive integer, G be a directed graph on n vertices and $\pi : V(G) \to [n]$ be a bijection. For simplicity, assume that $V(G) = \{u_1, \dots, u_n\}$ and, for each $u_i, u_j \in V(G)$, $i < j$ if and only if $u_i <_\pi u_j$.

For a proper subset $X \subset V(G)$, we let $E_G(X)$ denote the subset of edges of G with one endpoint in X and another endpoint in $V(G) \setminus X$. An *edge cut* (or simply *cut*) of G is defined as a subset $S \subseteq E(G)$ such that, for some $X \subset V(G)$, $S = E_G(X)$. For each $i \in [n-1]$, we let $S_G(\pi, i) \doteq E_G(\{u_1, \dots, u_i\})$ be the *ith cut of G with respect to* π. Then, the *cutwidth of G with respect to* π is defined as $\mathrm{cw}(G, \pi) \doteq \max_{i\in[n-1]}|S_G(\pi, i)|$, and the *cutwidth of G* is defined as the minimum $\mathrm{cw}(G, \pi)$ over all bijections $\pi : V(G) \to [n]$.

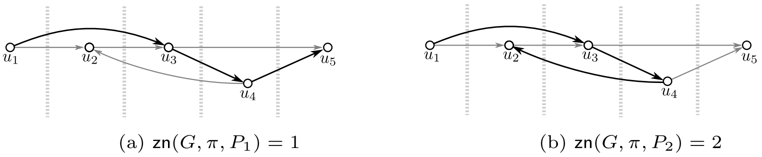

Fig. 1. Directed graph G, bijection $\pi : V(G) \to [|G|]$, where $i < j$ iff $u_i <_\pi u_j$, and directed paths P_1 and P_2 (in bold), such that $\mathsf{zn}(G, \pi, P_1) = 1$ and $\mathsf{zn}(G, \pi, P_2) = 2$, respectively.

Let P be a directed path of G. We let $\mathsf{zn}(G, \pi, P)$ be the maximum number of edges of P that are part of the cut $S_G(\pi, i)$, where the maximum is taken over all $i \in [n-1]$. More formally,

$$\mathsf{zn}(G, \pi, P) \doteq \max_{i \in [n-1]} |E(P) \cap S_G(\pi, i)|.$$

Then, we let $\mathsf{zn}(G, \pi)$ be the maximum $\mathsf{zn}(G, \pi, P)$ over all directed paths P of G. Finally, we define the *zig-zag number of* G, denoted by $\mathsf{zn}(G)$, as the minimum $\mathsf{zn}(G, \pi)$ over all bijections $\pi : V(G) \to [n]$.

Fig. 1 exemplifies a directed graph G and a bijection $\pi : V(G) \to [|G|]$ such that $\mathsf{zn}(G, \pi) = 2$. In fact, one can verify that $\mathsf{zn}(G) = \mathsf{zn}(G, \pi) = 2$.

It is straightforward from the definition of zig-zag number that a directed graph has zig-zag number 0 if and only if it does not contain any edge. Moreover, one can verify that every directed acyclic graph with at least one edge has zig-zag number 1. Indeed, it is known that a directed graph G is directed acyclic if and only if it admits a *topological ordering*, *i.e.* a linear order $<_\pi$ such that $u <_\pi v$ for each $(u, v) \in E(G)$. Thus, one can verify that, if G is a directed acyclic graph and π corresponds to a topological ordering of G, then $\mathsf{zn}(G, \pi) = 1$. In other words, graphs of zig-zag number at least 2 must contain directed cycles. On the other hand, every directed graph G with a directed cycle of length at least 3 necessarily has zig-zag number greater than or equal to 2. Indeed, in this case, for each bijection $\pi : V(G) \to [|G|]$, there always exist three distinct vertices $a, b, c \in V(G)$ such that $\langle a, b, c \rangle$ is a directed path of G, where $a <_\pi b$ and $c <_\pi b$. Intuitively, the zig-zag number of a directed graph measures how much its directed cycles are nested.

Next, we formally define the Zig-zag number problem.

Zig-zag number

Input: A directed graph G and a non-negative integer k.

Question: Is $\mathsf{zn}(G) \leq k$? In other words, does there exist a bijection $\pi : V(G) \to [|G|]$ such that, for every directed path P of G,

$$\mathsf{zn}(G, \pi, P) = \max_{i \in [|G|-1]} |E(P) \cap S_G(\pi, i)| \leq k?$$

In particular, for each fixed non-negative integer k, we define k-zig-zag number as the decision problem that, given a directed graph G, asks whether $\mathsf{zn}(G) \leq k$. More formally:

k-zig-zag number

Input: A directed graph G.

Question: Is $\mathsf{zn}(G) \leq k$? In other words, does there exist a bijection $\pi : V(G) \to [|G|]$ such that, for every directed path P of G,

$$\mathsf{zn}(G, \pi, P) = \max_{i \in [|G|-1]} |E(P) \cap S_G(\pi, i)| \leq k?$$

3. NP-membership for fixed k

In this section, we prove that k-zig-zag number is in NP for each fixed k. We remark that a naive application of the definition of zig-zag number of a directed graph naturally leads to a Σ_2^{P}-upper bound. To circumvent this and settle k-zig-zag number in NP, we show how to replace the inner universal quantifier, which iterates over all directed paths, with an XP-time deterministic computation corresponding to a guessed linear order of the vertices of the input graph and the integer k. More specifically, we prove the following theorem.

Theorem 1. *Let G be a directed graph and k be a non-negative integer. One can non-deterministically decide in time $|G|^{\mathcal{O}(k)}$ whether $zn(G) \leq k$.*

In order to prove Theorem 1, we reduce the problem of deciding whether $\mathsf{zn}(G, \pi) \geq k + 1$, for a guessed bijection $\pi : V(G) \to [|G|]$, to the Reachability problem in a suitably defined directed acyclic graph, denoted by $D_G(\pi, k)$, which we call *compatibility graph* of the triple (G, π, k). The formal definition of such a graph is properly given later on. Next, we describe how this section is organized.

In Section 3.1, we define the concept of *compatible subcut sequence* of a directed graph G with respect to a bijection $\pi : V(G) \rightarrow [|G|]$. Based on this concept, we provide a necessary and sufficient condition for $\mathsf{zn}(G, \pi) \geq k + 1$. Considering such a condition, we formally define in Section 3.2 the notion of *compatibility graph* and, then, we introduce a characterization relating the existence of the compatible subcut sequences of interest to the existence of directed paths with $|G| - 1$ vertices in the compatibility graph of (G, π, k). The proof of this characterization is presented in Section 3.3.

3.1. Compatible subcut sequence

Let G be a directed graph on n vertices and $\pi : V(G) \rightarrow [n]$ be a bijection. For simplicity, assume throughout this section that $V(G) = \{u_1, \ldots, u_n\}$ and, for each $u_i, u_j \in V(G)$, $i < j$ if and only if $u_i <_\pi u_j$.

The *cut sequence of G with respect to π* is defined as the sequence

$$\gamma_{G,\pi} \doteq \langle S_G(\pi, 1), \ldots, S_G(\pi, n-1)\rangle.$$

For each $i, j \in [n-1]$, with $i < j$, and each two *subcuts* $S'_i \subseteq S_G(\pi, i)$ and $S'_j \subseteq S_G(\pi, j)$, we say that S'_i is *compatible* with S'_j, and we denote this fact by $S'_i \prec_{G,\pi} S'_j$, if for each $e = (u, v) \in E(G)$ the two conditions below are observed.

(1) If $e \in S'_i$, and either $\pi(u) > j$ or $\pi(v) > j$, then $e \in S'_j$.
(2) If $e \in S_G(\pi, i) \setminus S'_i$, then $e \notin S'_j$.

Intuitively, condition (1) says that, if e belongs to the subcut S'_i and either the order of u or the order of v, with respect to π, is greater than j, then e must belong to the subcut S'_j. On the other hand, condition (2) says that if e belongs to the cut $S_G(\pi, i)$ but does not belong to the subcut S'_i, then e cannot belong to the subcut S'_j.

A *compatible subcut sequence* of $\gamma_{G,\pi}$ is a sequence of subcuts

$$\gamma' = \langle S'_1, \ldots, S'_{n-1}\rangle$$

such that $S'_i \subseteq S_G(\pi, i)$ for each $i \in [n-1]$, and $S'_j \prec_{G,\pi} S'_{j+1}$ for each $j \in [n-2]$. A neat idea behind the definition of compatible subcut sequence is to focus on neighboring subcuts S'_j and S'_{j+1} for each $j \in [n-2]$. We let $\boldsymbol{\Gamma}(\gamma_{G,\pi})$ be the set of all compatible subcut sequences of $\gamma_{G,\pi}$.

The next proposition establishes that the compatibility conditions (1) and (2) described above are sufficient to ensure that, if S'_i and S'_j are two subcuts in a same compatible subcut sequence, then there do not exist any inconsistency with respect to the edges that belong to S'_i and to S'_j. More specifically, provided that S'_i and S'_j are two subcuts in a same compatible subcut sequence, for any edge e belonging simultaneously to the cuts $S_G(\pi, i)$ and $S_G(\pi, j)$, we have that e belongs to S'_i if and only if it belongs to S'_j.

Proposition 1. *Let G be a directed graph, $\pi : V(G) \rightarrow [|G|]$ be a bijection and $\gamma' = \langle S'_1, \ldots, S'_{|G|-1}\rangle \in \boldsymbol{\Gamma}(\gamma_{G,\pi})$. For each edge $e \in E(G)$ and each $i \in [|G|-1]$, if $e \in S_G(\pi, i) \setminus S'_i$, then $e \notin S'_j$ for any $j \in [|G|-1]$.*

Proof. Let $i, j \in [|G|-1]$, with $i \neq j$, and $e = (u, v)$ be an edge in $S_G(\pi, i)$. The proof is split into two cases.

First, assume that $i < j$. Note that $i < |G| - 1$, otherwise $j < i$. Moreover, we have by hypothesis that $S'_l \prec_{G,\pi} S'_{l+1}$ for each $l \in \{i, \ldots, |G|-2\}$. Thus, if $e \in S_G(\pi, l) \setminus S'_l$, then $e \notin S'_{l+1}$ for each $l \in \{i, \ldots, |G|-2\}$. This inductively implies that, if $e \in S_G(\pi, i) \setminus S'_i$, then either $e \notin S_G(\pi, l)$ or $e \in S_G(\pi, l) \setminus S'_l$ for each $l \in \{i+1, \ldots, |G|-1\}$. In particular, if $e \in S_G(\pi, i) \setminus S'_i$, then $e \notin S'_j$.

Now, assume that $i > j$. We prove that, if $e \in S'_j$, then $e \in S'_i$. Thus, additionally assume that $e \in S'_j$. Since $e \in S_G(\pi, j)$ and $e \in S_G(\pi, i)$, either $\pi(u) > l$ or $\pi(v) > l$ for each $l \in \{j, \ldots, i\}$. This and the hypothesis that $S'_l \prec_{G,\pi} S'_{l+1}$ imply that, if $e \in S'_l$, then $e \in S'_{l+1}$ for each $l \in \{j, \ldots, i-1\}$. Therefore, since $e \in S'_j$, we inductively obtain that $e \in S'_l$ for each $l \in \{j, \ldots, i\}$. In particular, $e \in S'_i$. □

Let $\gamma' = \langle S'_1, \ldots, S'_{n-1}\rangle$ be a compatible subcut sequence in $\boldsymbol{\Gamma}(\gamma_{G,\pi})$. The *width* of γ' is defined as

$$\omega(\gamma') \doteq \max_{i \in [n-1]} |S'_i|.$$

If $E' = \bigcup_{i \in [n-1]} S'_i \neq \emptyset$, then we define $G[\gamma'] \doteq G[E']$ as the directed graph induced by E'. In particular, we remark that $\gamma_{G,\pi}$ is a compatible subcut sequence itself of width $\mathsf{cw}(G, \pi)$ and that $G[\gamma_{G,\pi}]$ consists in the directed graph obtained from G by removing all of its isolated vertices.

The next lemma states that deciding whether $\mathsf{zn}(G, \pi) \geq k + 1$ is equivalent to deciding whether there is a compatible subcut sequence of $\gamma_{G,\pi}$ of width at least $k + 1$, whose associated directed graph is a directed path.

Lemma 1. *Let G be a directed graph, $\pi : V(G) \rightarrow [|G|]$ be a bijection and k be a non-negative integer. Then, $zn(G, \pi) \geq k + 1$ if and only if there is a compatible subcut sequence $\gamma' = \langle S'_1, \ldots, S'_{|G|-1}\rangle \in \boldsymbol{\Gamma}(\gamma_{G,\pi})$ such that $\omega(\gamma') \geq k + 1$ and $G[\gamma']$ is a directed path.*

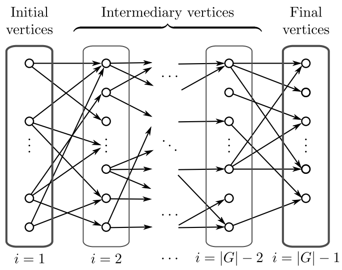

Fig. 2. A compatibility graph.

Proof. First, suppose that $\mathsf{zn}(G, \pi) \geq k + 1$, and let P be a directed path of G such that $\mathsf{zn}(G, \pi, P) \geq k + 1$. Consider the sequence $\gamma' = \langle S'_1, \ldots, S'_{|G|-1}\rangle$ of subcuts such that $S'_i = E(P) \cap S_G(\pi, i)$ for each $i \in [|G| - 1]$. We prove that γ' is a compatible subcut sequence of $\gamma_{G,\pi}$. In other words, we prove that $S'_i \prec_{G,\pi} S'_{i+1}$ for each $i \in [|G| - 2]$. Note that, if for some $i \in [|G| - 2]$ there exists an edge $e \in S_G(\pi, i) \setminus S'_i$, then $e \in E(G) \setminus E(P)$ and, consequently, $e \notin S'_{i+1}$. Now, suppose that for some $i \in [|G| - 2]$ there exists an edge $e = (u, v) \in S'_i$ such that either $\pi(u) > i + 1$ or $\pi(v) > i + 1$. Clearly, $e \in E(P)$. Moreover, note that $e \in S_G(\pi, i + 1)$. Thus, $e \in S'_{i+1}$, otherwise $e \in S_G(\pi, i + 1) \setminus S'_{i+1}$, which would imply that $e \in E(G) \setminus E(P)$. Therefore, γ' is indeed a compatible subcut sequence of $\gamma_{G,\pi}$. Additionally, one can straightforwardly verify that $\omega(\gamma') \geq k + 1$ and $G[\gamma']$ is a directed path.

Conversely, suppose that there exists a compatible subcut sequence $\gamma' = \langle S'_1, \ldots, S'_{|G|-1}\rangle$ of $\gamma_{G,\pi}$ such that $\omega(\gamma') \geq k+1$ and $G[\gamma']$ is a directed path. Thus, there exists $i \in [|G| - 1]$ such that $|S'_i| \geq k + 1$. As a result, if $P = G[\gamma']$, then $\mathsf{zn}(G, \pi, P) \geq |S'_i| \geq k + 1$. Therefore, $\mathsf{zn}(G, \pi) \geq \mathsf{zn}(G, \pi, P) \geq k + 1$. □

3.2. Compatibility graph

In this section, we define the notion of *compatibility graph*. Intuitively, each directed path with $|G| - 1$ vertices of the *compatibility graph* $D_G(\pi, k)$ corresponds to a compatible subcut sequence γ' of $\gamma_{G,\pi}$ satisfying the conditions described in Lemma 1. More specifically, the vertices of $D_G(\pi, k)$ consist in special tuples which, along with the directed edges between them, define a dynamic programming table. This table stores all the information needed to guarantee that, if there is a directed path in $D_G(\pi, k)$ with $|G| - 1$ vertices, then there exists a compatible subcut sequence γ' of $\gamma_{G,\pi}$ such that $\omega(\gamma') \geq k + 1$.

In order to capture the above property, we partition the vertex set of $D_G(\pi, k)$ into $|G| - 1$ distinct levels, such that each level $i \in [|G| - 1]$ is associated with the cut $S_G(\pi, i)$ and there is a *directed* edge in $D_G(\pi, k)$ from a vertex $\mathfrak{u}$ to a vertex $\mathfrak{v}$ only if $\mathfrak{u}$ belongs to a level i and $\mathfrak{v}$ belongs to the level $i + 1$, and some additional constraints (described in Section 3.2.2) are satisfied. The vertices in the level $i = 1$ are called *initial*, the vertices in a level $i \in [|G| - 1] \setminus \{1, |G| - 1\}$ are called *intermediary*, and the vertices in the level $i = |G| - 1$ are called *final*. We note that, by definition, the initial vertices of $D_G(\pi, k)$ have in-degree 0 and the final vertices of $D_G(\pi, k)$ have out-degree 0. Fig. 2 illustrates the partitioning of the vertex set of the compatibility graph $D_G(\pi, k)$ into these $|G| - 1$ distinct levels.

One can alternatively regard $D_G(\pi, k)$ as an acyclic finite automaton — with transition set defined by the adjacency relation described in Section 3.2.2. From this perspective, the initial vertices represent the initial states of the automaton and the final vertices represent the final states of the automaton.

The following immediate observation provides the basis for the definition of compatibility graph.

Observation 1. *A directed graph P is a directed path if and only if it satisfies the following four conditions:*

(DP1) P has exactly one vertex, called source vertex, with in-degree 0 *and out-degree* 1*;*
(DP2) P has exactly one vertex, called target vertex, with in-degree 1 *and out-degree* 0*;*
(DP3) All the other vertices of P have in-degree 1 *and out-degree* 1*;*
(DP4) P is weakly connected.

In particular, for a compatible subcut sequence γ' of $\gamma_{G,\pi}$, we have that $G[\gamma']$ is a directed path if and only if it satisfies Conditions (DP1)–(DP4). Based on that, and aiming at devising a dynamic programming method that determines whether there exists a compatible subcut sequence γ' such that $G[\gamma']$ is a directed path and that, more generally, satisfies the conditions described in Lemma 1, we define the set $\boldsymbol{B}_G(\pi, k)$. This set consists of all tuples of the form

$$\mathfrak{u}_i = \left(i, S'_i, \phi_i, \varphi_i, \mathcal{S}_i, \tau_i, \psi_i\right),$$

where $i \in [n - 1]$, S'_i is a subcut of $S_G(\pi, i)$ with cardinality at most $k + 1$, $\mathcal{S}_i$ is either an empty set or a partition of S'_i, $\phi_i, \varphi_i, \tau_i \in \{0, 1, 2\}$ are ternary flags and $\psi_i \in \{0, 1\}$ is a boolean flag. We remark that, for each $i \in [n - 1]$, there are at most $n^{\mathcal{O}(k)}$ distinct tuples $\mathfrak{u}_i \in \boldsymbol{B}_G(\pi, k)$. Indeed, for each $i \in [n - 1]$, the cut $S_G(\pi, i)$ has at most $2i(n - i) \leq n^2/2$ directed

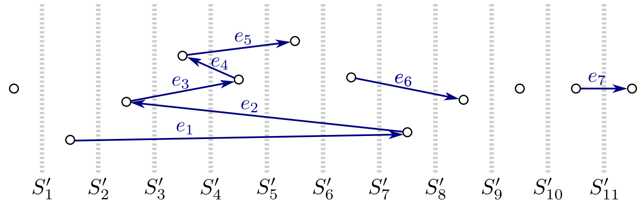

Fig. 3. Example of a compatible subcut sequence: $\gamma' = \langle S'_1, \ldots, S'_{11}\rangle$, where $S'_1 = \emptyset$, $S'_2 = \{e_1\}$, $S'_3 = \{e_1, e_2, e_3\}$, $S'_4 = \{e_1, e_2, e_3, e_4, e_5\}$, $S'_5 = \{e_1, e_2, e_5\}$, $S'_6 = \{e_1, e_2\}$, $S'_7 = \{e_1, e_2, e_6\}$, $S'_8 = \{e_6\}$, $S'_9 = \emptyset$, $S'_{10} = \emptyset$, and $S'_{11} = \{e_7\}$.

edges, which is the maximum possible number of directed edges between the vertices belonging to $\{u_1, \ldots, u_i\}$ and the vertices belonging to $\{u_{i+1}, \ldots, u_n\}$. Thus, $S_G(\pi, i)$ has at most $\binom{n^2/2}{k+1} = \mathcal{O}(n^{2k+2})$ distinct subcuts S'_i of cardinality at most $k+1$, and each such a subcut S'_i admits at most $(k+1)^{\mathcal{O}(k)}$ distinct partitions.

Let $i \in [n-1]$ and $\mathfrak{p}_i = \left(1, S'_1, \phi_1, \varphi_1, \mathcal{S}_1, \tau_1, \psi_1\right), \ldots, \left(i, S'_i, \phi_i, \varphi_i, \mathcal{S}_i, \tau_i, \psi_i\right)$ be a sequence of tuples, such that S'_j is compatible with S'_{j+1} for each $j \in [i-1]$. Then, let H_i be the subgraph of G with vertex set $V(H_i) = \{u_1, \ldots, u_i\} \cup X_i$ and edge set $E(H_i) = S'_1 \cup \cdots \cup S'_i$, where X_i denotes the set of endpoints of the edges in $S'_1 \cup \cdots \cup S'_i$. Note that H_i may contain isolated vertices.

Intuitively, the ternary flag ϕ_i (the ternary flag φ_i, resp.) informs whether there exist zero, one, or more than one vertices from $\{u_1, \ldots, u_i\}$ that are source vertices (target vertices, resp.) of H_i.

The partition $\mathcal{S}_i$ represents the set of all non-trivial weakly connected components of H_i, *restricted to the subcut* S'_i, that are defined by only taking into account the vertices from $\{u_1, \ldots, u_i\}$. In other words, two edges $e, e' \in S'_i$ belong to a same part of $\mathcal{S}_i$ if and only if there exists an undirected path of H_i between an endpoint of e and an endpoint of e' that only uses vertices from $\{u_1, \ldots, u_i\}$. For instance, consider the compatible subcut sequence γ' illustrated in Fig. 3. In this example, $\mathcal{S}_1 = \emptyset$, $\mathcal{S}_2 = \{\{e_1\}\}$, $\mathcal{S}_3 = \{\{e_1\}, \{e_2, e_3\}\}$, $\mathcal{S}_4 = \{\{e_1\}, \{e_2, e_3\}, \{e_4, e_5\}\}$, $\mathcal{S}_5 = \{\{e_1\}, \{e_2, e_5\}\}$, $\mathcal{S}_6 = \{\{e_1\}, \{e_2\}\}$, $\mathcal{S}_7 = \{\{e_1\}, \{e_2\}, \{e_6\}\}$, $\mathcal{S}_8 = \{\{e_6\}\}$, $\mathcal{S}_9 = \emptyset$, $\mathcal{S}_{10} = \emptyset$, and $\mathcal{S}_{11} = \{\{e_7\}\}$.

The ternary flag τ_i informs whether there exist zero, one, or more than one non-trivial weakly connected components of H_i that do not contain any of the vertices from $\{u_{i+1}, \ldots, u_n\}$. For instance, consider again the compatible subcut sequence γ' illustrated in Fig. 3. In this example, $\tau_i = 0$ for each $i \in \{1, \ldots, 7\}$, $\tau_8 = 1$, and $\tau_i = 2$ for each $i \in \{9, 10, 11\}$.

Finally, the boolean flag ψ_i informs whether or not there exists a subcut of width $k+1$ among the subcuts $S'_1, \ldots, S'_i$.

3.2.1. Initial, final and intermediary tuples

Now, we present the formal definitions of the notions of *initial*, *final* and *intermediary* tuples, which precisely comprise the vertex set of $D_G(\pi, k)$. More specifically, the initial tuples correspond to the initial vertices of $D_G(\pi, k)$, the final tuples correspond to the final vertices of $D_G(\pi, k)$, and the intermediary tuples correspond to the intermediary vertices of $D_G(\pi, k)$.

Let $\mathfrak{u} = \left(i, S'_i, \phi_i, \varphi_i, \mathcal{S}_i, \tau_i, \psi_i\right)$ be a tuple in $\boldsymbol{B}_G(\pi, k)$.

We say that $\mathfrak{u}$ is *initial* if $i = 1$ and the following conditions are satisfied:

1. The vertex u_1 has at most one in-edge and at most one out-edge in S'_1;
2. If u_1 has no in-edge and one out-edge in S'_1, then $\phi_1 = 1$ and $\varphi_1 = 0$;
3. If u_1 has one in-edge and no out-edge in S'_1, then $\phi_1 = 0$ and $\varphi_1 = 1$;
4. If u_1 has one in-edge and one out-edge in S'_1 or does not have any incident edge in S'_1, then $\phi_1 = 0$ and $\varphi_1 = 0$;
5. If $S'_1 = \emptyset$, then $\mathcal{S}_1 = \emptyset$; otherwise, $\mathcal{S}_1 = \left\{S'_1\right\}$;
6. $\tau_1 = 0$;
7. If $|S'_1| = k+1$, then $\psi_1 = 1$; otherwise, $\psi_1 = 0$.

On the other hand, we say that $\mathfrak{u}$ is *final* if $i = n-1$ and the following conditions are satisfied:

1. The vertex u_n has at most one in-edge and at most one out-edge in S'_{n-1};
2. If u_n has no in-edge and one out-edge in S'_{n-1}, then $\phi_{n-1} = 0$ and $\varphi_{n-1} = 1$;
3. If u_n has one in-edge and no out-edge in S'_{n-1}, then $\phi_{n-1} = 1$ and $\varphi_{n-1} = 0$;
4. If u_n has one in-edge and one out-edge in S'_{n-1} or does not have any incident edge in S'_{n-1}, then $\phi_{n-1} = 1$ and $\varphi_{n-1} = 1$;
5. $|\mathcal{S}_{n-1}| \leq 2$, and if $|\mathcal{S}_{n-1}| = 1$, then $|S'_{n-1}| = 1$;
6. $\tau_{n-1} \leq 1$, and if $\tau_{n-1} = 1$, then $S'_{n-1} = \emptyset$;
7. $\psi_{n-1} = 1$.

Intuitively, the tuple $\mathfrak{u}$ is called initial (final, resp.) if $i = 1$ ($i = n-1$, resp.) and the values assigned to the parameters $S'_i, \phi_i, \varphi_i, \mathcal{S}_i, \tau_i$ and ψ_i establish a valid configuration with respect to the semantic of each parameter itself and with respect to Conditions (DP1)–(DP4).

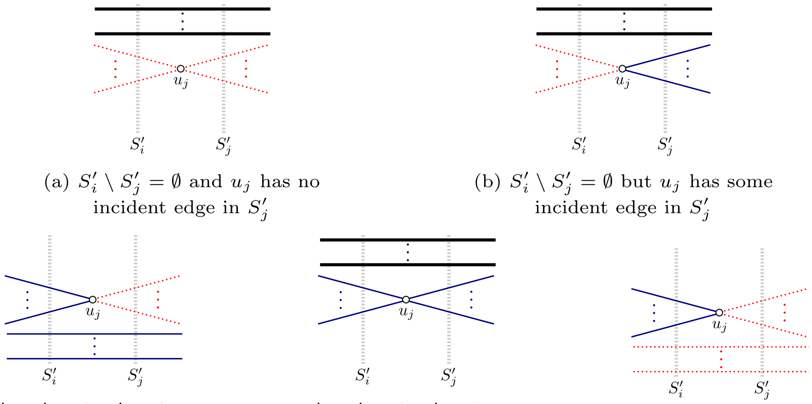

(a) $S'_i \setminus S'_j = \emptyset$ and u_j has no incident edge in S'_j

(b) $S'_i \setminus S'_j = \emptyset$ but u_j has some incident edge in S'_j

(c) $S'_i \setminus S'_j \neq \emptyset$, $S'_j \neq \emptyset$ and u_j has no incident edge in S'_j

(d) $S'_i \setminus S'_j \neq \emptyset$, $S'_j \neq \emptyset$ but u_j has some incident edge in S'_j

(e) $S'_i \setminus S'_j \neq \emptyset$ but $S'_j = \emptyset$

Fig. 4. Connectedness rules. (Red) dotted lines represent non-edges, (black) thicker lines represent non-mandatory edges, and (blue) normal style lines represent mandatory edges.

Finally, we say $\mathfrak{u}$ is *intermediary* if $i \in [n-1] \setminus \{1, n-1\}$.

3.2.2. Compatibility relation

Let $\mathfrak{u} = \left(i, S'_i, \phi_i, \varphi_i, \mathcal{S}_i, \tau_i, \psi_i\right)$ and $\mathfrak{v} = \left(j, S'_j, \phi_j, \varphi_j, \mathcal{S}_j, \tau_j, \psi_j\right)$ be a pair of tuples from $\boldsymbol{B}_G(\pi, k)$. We say that $\mathfrak{u}$ is *compatible* with $\mathfrak{v}$, and we denote such a fact by $\mathfrak{u} \rightsquigarrow \mathfrak{v}$, if $j = i + 1$, $S'_i \prec_{G,\pi} S'_j$, and $\mathfrak{u}$ and $\mathfrak{v}$ satisfy the Vertex degree, Connectedness and Minimum subcut width rules, which are presented below.

Vertex degree rules.

1. The vertex u_j has at most one in-edge and at most one out-edge in $S'_i \cup S'_j$.
2. If u_j has no in-edge and one out-edge in $S'_i \cup S'_j$, then $\phi_j = \min\{2, \phi_i + 1\}$ and $\varphi_j = \varphi_i$.
3. If u_j has one in-edge and no out-edge in $S'_i \cup S'_j$, then $\phi_j = \phi_i$ and $\varphi_j = \min\{2, \varphi_i + 1\}$.
4. If u_j has one in-edge and one out-edge in $S'_i \cup S'_j$ or does not have any incident edge in $S'_i \cup S'_j$, then $\phi_j = \phi_i$ and $\varphi_j = \varphi_i$.

Connectedness rules.

1. If $S'_i \setminus S'_j = \emptyset$ and u_j has no incident edge in S'_j (see Fig. 4(a)), then $\tau_j = \tau_i$ and $\mathcal{S}_j = \mathcal{S}_i$.
2. If $S'_i \setminus S'_j = \emptyset$ but u_j has some incident edge in S'_j (see Fig. 4(b)), then $\tau_j = \tau_i$ and $\mathcal{S}_j = \mathcal{S}_i \cup \left\{S'_j \setminus S'_i\right\}$.
3. If $S'_i \setminus S'_j \neq \emptyset$ and $S'_j \neq \emptyset$ (see Figs. 4(c) and 4(d)), then $\tau_j = \tau_i$ and $\mathcal{S}_j = (\mathcal{S}_i \setminus \mathcal{Q}'_j) \cup \mathcal{Q}_j$, where $\mathcal{Q}'_j$ denotes the collection of all sets in $\mathcal{S}_i$ that have at least one edge in S'_i with u_j as an endpoint, *i.e.*
$$\mathcal{Q}'_j = \left\{Q \in \mathcal{S}_i : Q \cap (S'_i \setminus S'_j) \neq \emptyset\right\},$$
and $\mathcal{Q}_j$ denotes the singleton collection whose set comprises all edges in S'_j with u_j as an endpoint, along with all edges in S'_j that belong to a set of $\mathcal{Q}'_j$, *i.e.*
$$\mathcal{Q}_j = \left\{(S'_j \setminus S'_i) \cup \left(\bigcup_{Q \in \mathcal{Q}'_j} Q \cap S'_j\right)\right\}.$$
In this case, we further require $\mathcal{Q}_j \neq \emptyset$.
Informally, $\mathcal{Q}'_j$ represents the set of non-trivial weakly connected components *restricted to* S'_i that have at least one edge in S'_i with u_j as an endpoint. Since, when considering the subcut S'_j, all such components contain u_j as a common vertex, they actually form a single non-trivial weakly connected component *restricted to* S'_j. This single component is represented by $\mathcal{Q}_j$, which, besides the edges that are already present in S'_i, contains all the edges in S'_j with u_j as an endpoint.
4. If $S'_i \setminus S'_j \neq \emptyset$ but $S'_j = \emptyset$ (see Fig. 4(e)), then $\tau_j = \min\{2, \tau_i + 1\}$ and $\mathcal{S}_j = \emptyset$.

Minimum subcut width rule.

(1) If $\psi_i = 1$ or $|S'_j| = k + 1$, then $\psi_j = 1$.

We notice that, for any sequence $\langle \mathfrak{u}_1, \ldots, \mathfrak{u}_{n-1} \rangle$ of tuples from $\boldsymbol{B}_G(\pi, k)$, such that $\mathfrak{u}_i$ is compatible with $\mathfrak{u}_{i+1}$ for each $i \in [n-2]$, there exists a unique associated compatible subcut sequence $\gamma' \in \boldsymbol{\Gamma}(\gamma_{G,\pi})$. Thus, the intuition behind the Vertex degree and Connectedness rules is ensuring that, if γ' is the subcut sequence associated with a directed path $\langle \mathfrak{u}_1, \ldots, \mathfrak{u}_{n-1} \rangle$ in $D_G(\pi, k)$, then $G[\gamma']$ satisfies Conditions (DP1)–(DP4). And, the intuition behind the Minimum subcut width rule is ensuring that the width of any such compatible subcut sequences γ' is at least $k+1$.

Now, we are finally able to formally define the notion of *compatibility graph* and then prove Theorem 1.

For each directed graph G, each bijection $\pi : V(G) \to [|G|]$ and each non-negative integer k, we define the *compatibility graph* of the triple (G, π, k) as the directed acyclic graph $D_G(\pi, k)$ with vertex set

$$V = \{\mathfrak{u} \in \boldsymbol{B}_G(\pi, k) : \mathfrak{u} \text{ is initial, intermediary or final}\}$$

and edge set

$$E = \{(\mathfrak{u}, \mathfrak{v}) \in V \times V : \mathfrak{u} \rightsquigarrow \mathfrak{v}\}.$$

Lemma 2 states that deciding whether there exists a compatible subcut sequence of $\gamma_{G,\pi}$ of width at least $k+1$ whose associated directed graph is a directed path is equivalent to deciding whether there exists a directed path of $D_G(\pi, k)$ with $n-1$ vertices. Then, based on this characterization and on Lemma 1, we prove in Lemma 3 that deciding whether $\mathsf{zn}(G, \pi) \geq k+1$ is reducible to the Reachability problem in $D_G(\pi, k)$.

Section 3.3 is devoted to present the proof of Lemma 2.

Lemma 2. *Let G be a directed graph, $\pi : V(G) \to [|G|]$ be a bijection and k be a non-negative integer. There exists a compatible subcut sequence $\gamma' \in \boldsymbol{\Gamma}(\gamma_{G,\pi})$ such that $\omega(\gamma') \geq k+1$ and $G[\gamma']$ is a directed path if and only if there exists a directed path of $D_G(\pi, k)$ with $|G| - 1$ vertices.*

Lemma 3. *Given a directed graph G, a bijection $\pi : V(G) \to [|G|]$ and a non-negative integer k, one can deterministically decide in time $|G|^{\mathcal{O}(k)}$ whether $zn(G, \pi) \leq k$.*

Proof. First, we construct the directed graph $D_G(\pi, k)$. Note that, for each tuple

$$\mathfrak{u}_i = \left(i, S'_i, \phi_i, \varphi_i, \mathcal{S}_i, \tau_i, \psi_i\right) \in \boldsymbol{B}_G(\pi, k),$$

the subcut S'_i has at most $k+1$ distinct elements. As a result, one can easily check in time polynomial in k if $\mathfrak{u}_i$ is an *initial*, an *intermediary* or a *final* tuple. Moreover, since there are $|G|^{\mathcal{O}(k)}$ distinct tuples in $\boldsymbol{B}_G(\pi, k)$, the vertex set of $D_G(\pi, k)$ can be determined in time $|G|^{\mathcal{O}(k)} \cdot \mathsf{poly}(k) = |G|^{\mathcal{O}(k)}$. Regarding the edge set of $D_G(\pi, k)$, we have by definition that there exists a directed edge from a vertex $\mathfrak{u}_i = \left(i, S'_i, \phi_i, \varphi_i, \mathcal{S}_i, \tau_i, \psi_i\right)$ to a vertex $\mathfrak{u}_j = \left(j, S'_j, \phi_j, \varphi_j, \mathcal{S}_j, \tau_j, \psi_j\right)$ of $D_G(\pi, k)$ if and only if $\mathfrak{u}_i$ is *compatible with* $\mathfrak{u}_j$, *i.e.*, $j = i+1$, $S'_i \prec_{G,\pi} S'_j$, and $\mathfrak{u}$ and $\mathfrak{v}$ satisfy the Vertex degree, Connectedness and Minimum subcut width rules. Since $|S'_i| \leq k+1$ and $|S'_j| \leq k+1$, the satisfaction of the Vertex degree, Connectedness and Minimum subcut width rules by $\mathfrak{u}$ and $\mathfrak{v}$ can be clearly checked in time polynomial in k. In addition, one can verify whether $S'_i \prec_{G,\pi} S'_j$ in time polynomial in $|G|$. Thus, it can be checked in time $\mathsf{poly}(|G|, k)$ whether there should exist in $D_G(\pi, k)$ a directed edge from $\mathfrak{u}$ to $\mathfrak{v}$. This implies that the edge set of $D_G(\pi, k)$ can be determined in time $|G|^{\mathcal{O}(k)} \cdot |G|^{\mathcal{O}(k)} \cdot \mathsf{poly}(|G|, k) = |G|^{\mathcal{O}(k)}$. Therefore, $D_G(\pi, k)$ can be wholly constructed in time $|G|^{\mathcal{O}(k)}$.

Then, by using an algorithm for the Reachability problem, we decide in time linear in the number of vertices and edges of $D_G(\pi, k)$, *i.e.* in time $|G|^{\mathcal{O}(k)}$, whether there is a directed path of $D_G(\pi, k)$ with $|G| - 1$ vertices. By Lemmas 1 and 2, such a path exists if and only if $\mathsf{zn}(G, \pi) \geq k+1$. Therefore, we can decide in time $|G|^{\mathcal{O}(k)}$ whether $\mathsf{zn}(G, \pi) \leq k$. □

As a result, we obtain that deciding whether $\mathsf{zn}(G) \leq k$ is in NP for each fixed $k \geq 0$, concluding thereby the proof of Theorem 1.

3.3. Proof of Lemma 2

Assume that $V(G) = \{u_1, \ldots, u_n\}$ and, for each $u_i, u_j \in V(G)$, $i < j$ if and only if $u_i <_\pi u_j$. Consider the following auxiliary claim.

Claim 1. *Let $\mathfrak{p} = \langle \mathfrak{u}_1, \ldots, \mathfrak{u}_{n-1} \rangle$ be a sequence of tuples such that*

1. *for each $i \in [n-1]$, $\mathfrak{u}_i = \left(i, S'_i, \phi_i, \varphi_i, \mathcal{S}_i, \tau_i, \psi_i\right) \in \boldsymbol{B}_G(\pi, k)$;*
2. *$\mathfrak{u}_1$ is initial;*
3. *for each $i \in [n-2]$, $\mathfrak{u}_i \rightsquigarrow \mathfrak{u}_{i+1}$;*

and let $\gamma' = \langle S'_1, \ldots, S'_{n-1} \rangle$ be the compatible subcut sequence corresponding to $\mathfrak{p}$. Then, for each $\ell \in [n-2]$, we have that any two edges $e, e' \in S'_\ell$ belong to a same part of $\mathcal{S}_\ell$ if and only if there exists in $G[\gamma']$ an undirected path between x and y that only contains vertices from $\{u_1, \ldots, u_\ell\}$, where x and y denote the endpoints of e and e', respectively, that belong to $\{u_1, \ldots, u_\ell\}$.

Proof of claim. The proof is by induction on ℓ.

Base case. Suppose that $\ell = 1$. In this case, $x = y$. Then, trivially, there exists in $G[\gamma']$ an undirected path between x and y that only contains vertices from $\{u_1, \dots, u_\ell\}$. Moreover, it follows from the fact that $\mathfrak{u}_1$ is an initial tuple that $\mathcal{S}_1 = \left\{S'_\ell\right\}$. Thus, e and e' belong to a same part of $\mathcal{S}_1$.

Inductive hypothesis. Suppose that there exists $\iota \in [|G| - 1]$ such that the claim holds for each $\ell \in [\iota - 1]$.

Inductive step. Suppose that $\ell = \iota > 1$. First, consider $x = y$, and let $i = \pi(x) \leq \iota$. Similarly to the base case, there trivially exists an undirected path of $G[\gamma']$ between x and y that only contains vertices from $\{u_1, \dots, u_\iota\}$. Moreover, since $\mathfrak{u}_{i-1} \rightsquigarrow \mathfrak{u}_i$, it follows from the Connectedness rules that e and e' belong to a same part of $\mathcal{S}_i$. As a result, we obtain that e and e' also belong to a same part of $\mathcal{S}_\iota$. Thus, in what follows, consider $x \neq y$. Additionally, assume without loss of generality that $\pi(x) < \pi(y)$.

First, we prove that, if e and e' belong to a same part of $\mathcal{S}_\iota$, then there is in $G[\gamma']$ an undirected path between x and y that only contains vertices from $\{u_1, \dots, u_\iota\}$. Thus, suppose e and e' belong to a same part of $\mathcal{S}_\iota$. Note that, if e and e' belong to a same part of $\mathcal{S}_i$ for some $i < \iota$, then the result immediately follows from the inductive hypothesis. Thus, assume that ι is the least integer $j \in \{\pi(y), \dots, n-1\}$ such that e and e' belong to a same part of $\mathcal{S}_j$.

Consider $\iota = \pi(y)$. Since $\mathfrak{u}_{\iota-1} \rightsquigarrow \mathfrak{u}_\iota$, it follows from the Connectedness rules that there exists an edge $e'' \in S'_{\iota-1} \setminus S'_\iota$ such that e and e'' belong to a same part of $\mathcal{S}_{\iota-1}$, otherwise e and e' would belong to distinct parts of $\mathcal{S}_\iota$. Then, let z be the endpoint of e'' that belongs to $\{u_1, \dots, u_{\iota-1}\}$. By the inductive hypothesis, there exists in $G[\gamma']$ an undirected path between x and z that only contains vertices from $\{u_1, \dots, u_{\iota-1}\}$. Moreover, since $\iota = \pi(y)$ and $e'' \in S'_{\iota-1} \setminus S'_\iota$, we have that y is an endpoint of e''. Therefore, there exists in $G[\gamma']$ an undirected path between x and y that only contains vertices from $\{u_1, \dots, u_\iota\}$.

Now, consider $\iota > \pi(y)$. By the Connectedness rules, there exist two distinct edges $e''_x, e''_y \in S'_{\iota-1} \setminus S'_\iota$, such that e and e''_x belong to a same part of $\mathcal{S}_{\iota-1}$ and e' and e''_y belong to a same part of $\mathcal{S}_{\iota-1}$, otherwise e and e' would belong to distinct parts of $\mathcal{S}_\iota$. Then, let z_x and z_y be the endpoints of e''_x and e''_y, respectively, that belong to $\{u_1, \dots, u_{\iota-1}\}$. By the inductive hypothesis and by the minimality of ι, there exist in $G[\gamma']$ two vertex-disjoint undirected paths $P_x = \langle x, \dots, z_x\rangle$ and $P_y = \langle z_y, \dots, y\rangle$ that only contain vertices from $\{u_1, \dots, u_{\iota-1}\}$, where P_x is a path between x and z_x, and P_y is a path between z_y and y. Therefore, since z''_x and z''_y are both neighbors of u_ι, $P_x + \langle u_\iota\rangle + P_y = \langle x, \dots, z_x, u_\iota, z_y, \dots, y\rangle$ is an undirected path between x and y that only use vertices from $\{u_1, \dots, u_\iota\}$.

Now, we prove the converse part, *i.e.* if there is in $G[\gamma']$ an undirected path between x and y that only contains vertices from $\{u_1, \dots, u_\iota\}$, then e and e' belong to a same part of $\mathcal{S}_\iota$. Thus, suppose that there is in $G[\gamma']$ an undirected path P between x and y that only contains vertices from $\{u_1, \dots, u_\iota\}$. Also, assume that ι is the least integer in $\{\pi(y), \dots, n-1\}$ holding such a property.

Consider $\iota = \pi(y)$. In this case, one can verify that there exists exactly one edge e'' in the set $E(P) \cap S'_{\iota-1} \setminus S'_\iota$. Let z be the endpoint of e'' that belongs to $\{u_1, \dots, u_{\iota-1}\}$. Note that, $P - y = P - u_\iota$ is an undirected path between x and z that only contains vertices from $\{u_1, \dots, u_{\iota-1}\}$. By the inductive hypothesis, e and e'' belong to a same part of $\mathcal{S}_{\iota-1}$. Therefore, we obtain by Connectedness rule 3 that e and e' belong to a same part of $\mathcal{S}_\iota$.

Now, consider $\iota > \pi(y)$. In this case, there exist exactly two distinct edges e''_x, e''_y in the set $E(P) \cap S'_{\iota-1} \setminus S'_\iota$. Let z_x and z_y be the endpoints of e''_x and e''_y, respectively, that belong to $\{u_1, \dots, u_{\iota-1}\}$. Note that, $P - u_\iota$ consists of two undirected path P_x and P_y that only contain vertices from $\{u_1, \dots, u_{\iota-1}\}$, where P_x is an undirected path between x and z_x, and P_y is an undirected path between z_y and y. Thus, it follows from the inductive hypothesis that e and e''_x belong to a same part of $\mathcal{S}_{\iota-1}$, and that e' and e''_y belong to a same part of $\mathcal{S}_{\iota-1}$. Therefore, by Connectedness rule 3, e and e' belong to a same part of $\mathcal{S}_\iota$. ■

Now, we are finally able to properly prove Lemma 2.

First, suppose that there exists $\gamma' = \langle S'_1, \dots, S'_{n-1}\rangle \in \boldsymbol{\Gamma}(\gamma_{G,\pi})$ such that $P = G[\gamma']$ is a directed path and $\omega(\gamma') \geq k+1$. Let $\mathfrak{u}_1 = \left(1, S'_1, \phi_1, \varphi_1, \mathcal{S}_1, \tau_1, \psi_1\right)$ be the initial tuple in $\boldsymbol{B}_G(\pi, k)$ obtained from the subcut S'_1. We notice that, given the subcut S'_1, the parameters ϕ_1, φ_1, $\mathcal{S}_1$, τ_1 and ψ_1 are uniquely determined according to the definition of initial tuple. Thus, $\mathfrak{u}_1$ is well-defined. Additionally, note that, according to the Vertex degree, Connectedness and Minimum subcut width rules, for each $i \in \{2, \dots, n-1\}$ and each tuple

$$\mathfrak{u}_{i-1} = \left(i-1, S'_{i-1}, \phi_{i-1}, \varphi_{i-1}, \mathcal{S}_{i-1}, \tau_{i-1}, \psi_{i-1}\right) \in \boldsymbol{B}_G(\pi, k),$$

there exists exactly one tuple $\mathfrak{u}_i = \left(i, S'_i, \phi_i, \varphi_i, \mathcal{S}_i, \tau_i, \psi_i\right) \in \boldsymbol{B}_G(\pi, k)$ such that $\mathfrak{u}_{i-1}$ is compatible with (*i.e.* $\mathfrak{u}_{i-1} \rightsquigarrow \mathfrak{u}_i$).

Consequently, there exists a unique sequence $\langle \mathfrak{u}_1, \dots, \mathfrak{u}_{n-1}\rangle$ of tuples from $\boldsymbol{B}_G(\pi, k)$ that can be obtained from γ' and satisfies the conditions of $\mathfrak{u}_1$ being initial and of $\mathfrak{u}_i$ being compatible with $\mathfrak{u}_{i+1}$ for each $i \in [n-2]$.

We claim that such a sequence $\langle \mathfrak{u}_1, \dots, \mathfrak{u}_{n-1}\rangle$ corresponds to a directed path of $D_G(\pi, k)$ with $n-1$ vertices. To prove this, we just need to show that the tuple $\mathfrak{u}_{n-1} = \left(n-1, S'_{n-1}, \phi_{n-1}, \varphi_{n-1}, \mathcal{S}_{n-1}, \tau_{n-1}, \psi_{n-1}\right)$ is final.

Since by hypothesis P is a directed path, every vertex of P has in-degree at most one and out-degree at most one. Moreover, P contains exactly one source vertex $u_\iota \in V(P)$ for some $\iota \in [n]$. Thus, one can verify that: if $\iota < n$, then $\phi_i = 0$ for each $i \in [\iota - 1]$ and $\phi_j = 1$ for each $j \in \{\iota, \dots, n-1\}$; and, if $\iota = n$, then $\phi_i = 0$ for each $i \in [n-1]$. Similarly, P contains exactly one target vertex $u_{\iota'} \in V(P)$ for some $\iota' \in [n]$. Thus, one can verify that: if $\iota' < n$, then $\varphi_i = 0$ for each $i \in [\iota' - 1]$ and $\varphi_j = 1$ for each $j \in \left\{\iota', \dots, n-1\right\}$; and, if $\iota' = n$, then $\varphi_i = 0$ for each $i \in [n-1]$. As a result, if $\iota \neq n$ and $\iota' \neq n$, then $\phi_{n-1} = 1$ and $\varphi_{n-1} = 1$. On the other hand, note that $\iota \neq \iota'$. Hence, if $\iota = n$ and $\iota' < n$, then $\phi_{n-1} = 0$ and $\varphi_{n-1} = 1$; and if $\iota < n$ and $\iota' = n$, then $\phi_{n-1} = 1$ and $\varphi_{n-1} = 0$.

Moreover, since $\omega(\gamma') \geq |S'_\iota| = k+1$ for some $\iota \in [n-1]$, one can easily verify that $\psi_i = 0$ for each $i \in [\iota - 1]$ and $\psi_j = 1$ for each $j \in \{\iota, \dots, n-1\}$.

Now, let a and b be the least and the greatest integers in $[n-1]$, respectively, such that $S'_a \neq \emptyset$ and $S'_b \neq \emptyset$. Since by hypothesis $\omega(\gamma') \geq k+1 \geq 1$, such integers a and b are well-defined. Thus, it follows from the fact that P is weakly connected that the following properties hold.

1. For each $i \in [a-2]$, $S'_i \setminus S'_{i+1} = \emptyset$ and u_{i+1} has no incident edge in S'_{i+1}, which implies $\tau_{i+1} = \tau_i = 0$.
2. $S'_{a-1} \setminus S'_a = \emptyset$ but u_a has some incident edge in S'_a, which implies $\tau_a = \tau_{a-1} = 0$.
3. For each $i \in \{a, \dots, b-1\}$, $S'_{i+1} \neq \emptyset$, which implies $\tau_{i+1} = \tau_i = 0$. In particular, if $b = n-1$, then $\tau_{n-1} = 0$. On the other hand, if $b < n-1$, then $S'_b \setminus S'_{b+1} \neq \emptyset$ and $S'_{b+1} = \emptyset$, which implies $\tau_{b+1} = \tau_b + 1 = 1$; moreover, $S'_{i-1} \setminus S'_i = \emptyset$ and $S'_i = \emptyset$ for each $i \in \{b+2, \dots, n-1\}$, which implies $\tau_i = \tau_{i-1} = 1$.

As a result, we obtain that $\tau_{n-1} \leq 1$, and that $\tau_{n-1} = 1$ implies $S'_{n-1} = \emptyset$.

Since u_n has at most one in-edge and at most one out-edge in S'_{n-1}, it is immediate that $|S'_{n-1}| \leq 2$ and $|\mathcal{S}_{n-1}| \leq 2$. Moreover, it follows from Claim 1 that $|\mathcal{S}_{n-1}| = 1$ implies $|S'_{n-1}| = 1$, otherwise P would not be a directed path.

Now, we prove the converse of Lemma 2. Suppose that there is in $D_G(\pi, k)$ a directed path $\mathfrak{p} = \langle \mathfrak{u}_1, \dots, \mathfrak{u}_{n-1} \rangle$ with $n-1$ vertices. One can verify that, for each $i \in [n-1]$, $\mathfrak{u}_i$ necessarily consists in a tuple in $\boldsymbol{B}_G(\pi, k)$ of the form

$$\mathfrak{u}_i = \left(i, S'_i, \phi_i, \varphi_i, \mathcal{S}_i, \tau_i, \psi_i\right).$$

Note that, $\gamma' = \langle S'_1, \dots, S'_{n-1} \rangle \in \boldsymbol{\Gamma}(\gamma_{G,\pi})$. Moreover, it follows from the definition of $D_G(\pi, k)$ that the tuple $\mathfrak{u}_{n-1}$ is final. As a result, $\psi_{n-1} = 1$, and thus $\omega(\gamma') \geq k+1$, since by the Minimum subcut width rule we have that $\psi_{n-1} = 1$ if and only if $|S'_i| = k+1$ for some $i \in [n-1]$. Thus, it just remains to prove that $G[\gamma']$ is a directed path. We prove in the following claims that $G[\gamma']$ satisfies each of Conditions (DP1)–(DP4), respectively.

Claim 2. *$G[\gamma']$ contains exactly one source vertex.*

Proof of claim. Since $\mathfrak{u}_{n-1}$ is final, either $\phi_{n-1} = 0$ or $\phi_{n-1} = 1$.

Suppose that $\phi_{n-1} = 1$. By the Vertex degree rules, $\phi_i \leq \phi_{i+1}$ for each $i \in [n-2]$. Hence, there exists $\iota \in [n-1]$ such that $\phi_i = 0$ for each $i \in [\iota - 1]$ and $\phi_j = 1$ for each $j \in \{\iota, \dots, n-1\}$. Since $\mathfrak{u}_1$ is initial, if $\iota = 1$, then u_ι has no in-edge and one out-edge in S'_1. On the other hand, if $\iota > 1$, then, by the Vertex degree rules, u_ι has no in-edge and one out-edge in $S'_{\iota-1} \cup S'_\iota$. Consequently, u_ι is a source vertex of $G[\gamma']$. Moreover, one can verify that, for any $i \in [n] \setminus \{\iota\}$, there is no vertex u_i that has in-degree 0 and out-degree 1 in $G[\gamma']$, otherwise $\phi_{n-1} = 2$. Therefore u_n is the only source vertex of $G[\gamma']$.

Now, suppose that $\phi_{n-1} = 0$. Then, u_n has no in-edge and one out-edge in S'_{n-1}, otherwise $\mathfrak{u}_{n-1}$ would not be final. Thus, u_n is a source vertex of $G[\gamma']$. In addition, note that $\phi_i = 0$ for each $i \in [n-1]$, otherwise $\phi_{n-1} \neq 0$. As a result, we obtain that, for any $i \in [n-1]$, there is no vertex u_i that has in-degree 0 and out-degree 1 in $G[\gamma']$. Therefore u_n is the only source vertex of $G[\gamma']$. ■

Claim 3. *$G[\gamma']$ contains exactly one target vertex.*

Proof of claim. The proof of this claim is analogous to the proof of Claim 2, following from the fact that $\varphi_i \leq \varphi_{i+1}$ for each $i \in [n-2]$, and from the fact that either $\varphi_{n-1} = 0$ or $\varphi_{n-1} = 1$, since $\mathfrak{u}_{n-1}$ is a final tuple. ■

Claim 4. *Let s and t be the source and target vertices of $G[\gamma']$, respectively, and let $u \in V(G[\gamma']) \setminus \{s, t\}$. Then, u has in-degree 1 and out-degree 1 in $G[\gamma']$.*

Proof of claim. By hypothesis $\mathfrak{u}_1$ is initial, $\mathfrak{u}_{n-1}$ is final and $\mathfrak{u}_i \rightsquigarrow \mathfrak{u}_{i+1}$ for each $i \in [n-2]$. This implies that every vertex of $G[\gamma']$ has in-degree at most 1 and out-degree at most 1. Moreover, it follows from the uniqueness of s and from the uniqueness of t that the vertex u is neither a source vertex nor a target vertex of $G[\gamma']$. Therefore, since $G[\gamma']$ does not contain isolated vertices, we obtain that u necessarily has in-degree 1 and out-degree 1 in $G[\gamma']$. ■

Claim 5. *$G[\gamma']$ is a weakly connected graph.*

Proof of claim. For the sake of contradiction, suppose that $G[\gamma']$ is not weakly connected. Then, there exist two distinct vertices $u_i, u_j \in V(G[\gamma'])$ such that there is no undirected path between them in $G[\gamma']$. Let $u_{i'}$ be a vertex of $G[\gamma']$ such that there exists in $G[\gamma']$ an undirected path between u_i and $u_{i'}$, for some $i' \in \{1, \dots, n\} \setminus \{i\}$. And, let $u_{j'}$ be a vertex of $G[\gamma']$ such that there exists in $G[\gamma']$ an undirected path between u_j and $u_{j'}$, for some $j' \in \{1, \dots, n\} \setminus \{j\}$. Note that, such vertices $u_{i'}$ and $u_{j'}$ necessarily exist, since by definition $G[\gamma']$ does not contain any isolated vertices. Assume without loss of generality that $i' > i$ and $j' > j$. Additionally, assume that i' is the greatest integer belonging to $\{i+1, \dots, n\}$ that holds the property of existing in $G[\gamma']$ an undirected path between u_i and $u_{i'}$. Analogously, assume that j' is the greatest

(a) Case 1 (b) Case 2

Fig. 5. Proof that $G[\gamma']$ is weakly connected. (Red) dotted lines represent non-edges, (black) thicker lines represent possibly existing edges, (blue) waved lines represent undirected paths, and (blue) normal style lines represent mandatory edges. In these illustrations, we assume that $\iota = i'$ (the case $\iota = j'$ is symmetric).

integer belonging to $\{j+1, \ldots, n\}$ that holds the property of existing in $G[\gamma']$ an undirected path between u_j and $u_{j'}$. Let $\iota = \min\left\{i', j'\right\}$. By the maximalities of i' and j', there is no edge in S'_ι that has u_ι as an endpoint, *i.e.*

$$\left\{(x, y) \in S'_\iota : x = u_\iota \text{ or } y = u_\iota\right\} = \emptyset. \tag{1}$$

Moreover, one can readily verify that $S'_{\iota-1} \setminus S'_\iota \neq \emptyset$. We split the remainder of this proof into two cases.

Case 1. Suppose that $\left\{i, \ldots, i'\right\} \cap \left\{j, \ldots, j'\right\} \neq \emptyset$ (see Fig. 5(a)). Note that, necessarily $S'_\iota \neq \emptyset$. Then, it follows from Eq. (1) and from Claim 1 that there is no part $Q \in \mathcal{S}_{\iota-1}$ such that $Q \cap (S'_{\iota-1} \setminus S'_\iota) \neq \emptyset$ and $Q \cap S'_\iota \neq \emptyset$, otherwise there would exist in $G[\gamma']$ an undirected path between u_i and u_j. Thus, $\mathcal{Q}_\iota = \emptyset$. Therefore, since $S'_{\iota-1} \setminus S'_\iota \neq \emptyset$ and $S'_\iota \neq \emptyset$, we obtain by Connectedness rule 3 that $\mathfrak{u}_{\iota-1}$ is not compatible with $\mathfrak{u}_\iota$.

Case 2. Suppose that $\left\{i, \ldots, i'\right\} \cap \left\{j, \ldots, j'\right\} = \emptyset$ (see Fig. 5(b)). It follows from Eq. (1) and from Claim 1 that there is no part $Q \in \mathcal{S}_{\iota-1}$ such that $Q \cap (S'_{\iota-1} \setminus S'_\iota) \neq \emptyset$ and $Q \cap S'_\iota \neq \emptyset$, otherwise i' or j' would not be maximum with respect to the aforementioned properties. Thus, $\mathcal{Q}_\iota = \emptyset$. However, possibly $S'_\iota = \emptyset$. First, suppose that $S'_\iota \neq \emptyset$. Then, as in the previous case, it follows from the fact that $S'_{\iota-1} \setminus S'_\iota \neq \emptyset$ and from the Connectedness rules that the tuple $\mathfrak{u}_{\iota-1}$ is not compatible with the tuple $\mathfrak{u}_\iota$. On the other hand, suppose that $S'_\iota = \emptyset$. Then, $\tau_\iota \geq 1$. As a result, $\tau_l = \tau_{l-1} \geq 1$ for each $l \in \left\{\iota+1, \ldots, \max\left\{i', j'\right\} - 1\right\}$, and $\tau_l = 2$ for each $l \in \left\{\max\left\{i', j'\right\}, \ldots, n-1\right\}$. In particular, we obtain that:

1. either $n > \max\left\{i', j'\right\}$, and then $\tau_{n-1} = 2$;
2. or $n = \max\left\{i', j'\right\}$, and then $S'_{n-1} \neq \emptyset$ and $\tau_{n-1} = 1$.

In either case, $\mathfrak{u}_{n-1}$ is not a final tuple. Therefore, $G[\gamma']$ is weakly connected. ■

By the previous claims, we obtain that $G[\gamma']$ is indeed a directed path, and thereby we conclude the proof of Lemma 2.

4. NP-hardness

In this section, we prove that 2-ZIG-ZAG NUMBER is an NP-hard problem. For that, we present a polynomial-time reduction from POSITIVE NOT ALL EQUAL 3SAT, which is a well-known NP-complete problem [17], defined next.

POSITIVE NOT ALL EQUAL 3SAT (PNAE 3SAT)

Input:	Set X of variables and a collection $\mathcal{C}$ of clauses over X such that each clause has no negative literal and exactly three positive literals.
Question:	Is there a truth assignment $\alpha : X \to \{0, 1\}$ such that each clause in $\mathcal{C}$ has at least one true literal and at least one false literal under α?

Construction 1. *Let $I = (X, \mathcal{C})$ be an instance of* PNAE 3SAT *with variable set X and clause set $\mathcal{C}$. We let G_I be the directed graph obtained from I as follows.*

- *For each variable $x_i \in X$, add the vertices u_i^1, u_i^2 and u_i^3, and add the edges (u_i^1, u_i^2), (u_i^2, u_i^3) and (u_i^3, u_i^1).*
- *For each clause $C_j \in \mathcal{C}$, add the vertices v_j^1, v_j^2 and v_j^3, and add the edges (v_j^1, v_j^2), (v_j^2, v_j^3) and (v_j^3, v_j^1). Moreover, assuming $C_j = \left\{x_{l_1}, x_{l_2}, x_{l_3}\right\}$ with $l_1 < l_2 < l_3$, add the edges $\left(u_{l_1}^1, v_j^1\right)$, $\left(u_{l_1}^3, v_j^1\right)$, $\left(u_{l_2}^1, v_j^2\right)$, $\left(u_{l_2}^3, v_j^2\right)$, $\left(u_{l_3}^1, v_j^3\right)$ and $\left(u_{l_3}^3, v_j^3\right)$.*

For each variable $x_i \in X$, we let H_i denote the subgraph of G_I induced by the vertices in $\left\{u_i^1, u_i^2, u_i^3\right\}$. And, for each clause $C_j \in \mathcal{C}$, we let $\widetilde{H}_j$ denote the subgraph of G_I induced by the vertices in $\left\{v_j^1, v_j^2, v_j^3\right\}$. We remark that H_i and $\widetilde{H}_j$ are directed cycles of length 3.

Fig. 6 exemplifies the directed graph G_I, described in Construction 1.

We establish in Lemmas 4 and 6 that there exists a satisfying truth assignment for an instance I of PNAE 3SAT if and only if there exists a linear order of zig-zag number at most 2 for the vertices of G_I. The central idea of our proof

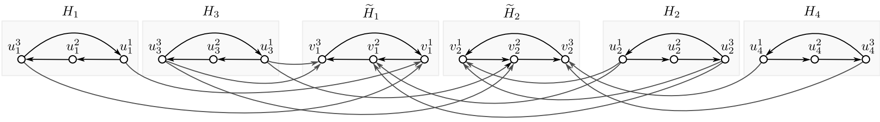

Fig. 6. Directed graph G_I obtained from the instance $I = (X, \mathcal{C})$ of PNAE 3SAT where $X = \{x_1, x_2, x_3, x_4\}$ and $\mathcal{C} = \{C_1 = \{x_1, x_2, x_3\}, C_2 = \{x_2, x_3, x_4\}\}$.

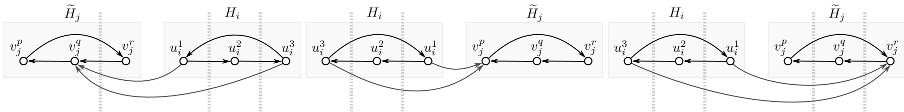

(a) $i = l_q$ (b) $i = l_p$ (c) $i = l_r$

Fig. 7. Case in which the clause $C_j = \{x_{l_1}, x_{l_2}, x_{l_3}\}$ has exactly one true literal under the truth assignment α, say x_{l_q} for some $q \in \{1, 2, 3\}$.

is to explore the possible internal relative orderings of the vertices of each directed cycle of G_I and, for each clause $C_j = \{x_{l_1}, x_{l_2}, x_{l_3}\} \in \mathcal{C}$, the possible ordered relative placements among the subgraphs H_{l_1}, H_{l_2}, H_{l_3}, and $\widetilde{H}_j$.

Lemma 4. *Let $I = (X, \mathcal{C})$ be an instance of* PNAE 3SAT*. If I is a yes instance of* PNAE 3SAT*, then $zn(G_I) \le 2$.*

Proof. Let $\alpha : X \to \{0, 1\}$ be a truth assignment such that each clause in $\mathcal{C}$ has at least one true literal and at least one false literal under α. In what follows, we define from α a linear order $<_\pi$ of the vertices of G_I such that $\mathrm{zn}(G_I, \pi) \le 2$.

Throughout this proof, consider $X = \{x_1, \dots, x_{|X|}\}$ and $\mathcal{C} = \{C_1, \dots, C_{|\mathcal{C}|}\}$.

For each variable $x_i \in X$, set

$$\begin{cases} u_i^1 <_\pi u_i^2 <_\pi u_i^3 & \text{if } \alpha(x_i) = 1 \\ u_i^1 >_\pi u_i^2 >_\pi u_i^3 & \text{otherwise.} \end{cases}$$

Let $V_0' \doteq \{u_i^1, u_i^2, u_i^3 : \alpha(x_i) = 0\}$ and $V_1' \doteq \{u_i^1, u_i^2, u_i^3 : \alpha(x_i) = 1\}$. Then, for each $y \in V_0'$ and each $z \in V(G_I) \setminus V_0'$, set $y <_\pi z$. And, for each $y \in V_1'$ and each $z \in V(G_I) \setminus V_1'$, set $y >_\pi z$.

Let C_j be a clause in $\mathcal{C}$. Assume that $C_j = \{x_{l_1}, x_{l_2}, x_{l_3}\}$ with $l_1 < l_2 < l_3$. There are two cases to be considered. First, suppose that C_j has exactly one true literal under α, say l_q for some $q \in \{1, 2, 3\}$. Then, set

$$v_j^p <_\pi v_j^q <_\pi v_j^r,$$

where $p = q \bmod 3 + 1$ and $r = (q + 1) \bmod 3 + 1$. Now, suppose that C_j has exactly two true literals under α. Thus, C_j has exactly one false literal under α, say l_q for some $q \in \{1, 2, 3\}$. Then, set

$$v_j^p >_\pi v_j^q >_\pi v_j^r,$$

where $p = q \bmod 3 + 1$ and $r = (q + 1) \bmod 3 + 1$.

Finally, for each pair of distinct variables $x_i, x_{i'} \in X$ with $i < i'$, such that $\alpha(x_i) = \alpha(x_{i'})$, set $u_i^p <_\pi u_{i'}^q$ for each $p, q \in \{1, 2, 3\}$. And, for each pair of distinct clauses $C_j, C_{j'} \in \mathcal{C}$ with $j < j'$, set $v_j^p <_\pi v_{j'}^q$ for each $p, q \in \{1, 2, 3\}$.

One can readily verify that $<_\pi$ is indeed a linear order of the vertices of G_I.

Now, we prove that $\mathrm{zn}(G_I, \pi) \le 2$. For the sake of contradiction, suppose that there exists a directed path P in G_I such that $\mathrm{zn}(G_I, \pi, P) \ge 3$. Assume without loss of generality that P is a minimal path with respect to the property that $\mathrm{zn}(G_I, \pi, P) \ge 3$. Recall that, for each variable $x_i \in X$, H_i is a directed cycle of length 3. Similarly, for each clause $C_j \in \mathcal{C}$, $\widetilde{H}_j$ is a directed cycle of length 3. Consequently, P is neither a subgraph of H_i nor a subgraph of $\widetilde{H}_j$, for any $x_i \in X$ and any $C_j \in \mathcal{C}$, otherwise $\mathrm{zn}(G_I, \pi, P) < 3$. Moreover, every edge of G_I is either an edge of one of these subgraphs H_i and $\widetilde{H}_j$ or is an edge from a vertex of H_i to a vertex of $\widetilde{H}_j$, for some $x_i \in X$ and some $C_j \in \mathcal{C}$. As a result, there exists precisely one variable $x_i \in X$ and there exists precisely one clause $C_j \in \mathcal{C}$ such that $V(P) \cap V(H_i) \ne \emptyset$ and $V(P) \cap V(\widetilde{H}_j) \ne \emptyset$. More specifically, P consists in a directed path on at most 4 vertices (by its minimality) from a vertex of H_i to a vertex of $\widetilde{H}_j$ that only contains vertices belonging to $V(H_i) \cup V(\widetilde{H}_j)$. Assume that $C_j = \{x_{l_1}, x_{l_2}, x_{l_3}\}$ with $l_1 < l_2 < l_3$.

First, consider the case in which C_j has exactly one true literal under α, and let x_{l_q} be such a literal for some $q \in \{1, 2, 3\}$. Thus, $v_j^p <_\pi v_j^q <_\pi v_j^r$, where $p = q \bmod 3 + 1$ and $r = (q+1) \bmod 3 + 1$. Consequently, if $i = l_q$, then $u_i^1 <_\pi u_i^2 <_\pi u_i^3$ and $V(H_i) >_\pi V(\widetilde{H}_j)$, which implies $\mathrm{zn}(G_I, \pi, P) < 3$ (see Fig. 7(a)). On the other hand, if $i = l_p$ or $i = l_r$, then $u_i^1 >_\pi u_i^2 >_\pi u_i^3$ and $V(H_i) <_\pi V(\widetilde{H}_j)$, which also implies $\mathrm{zn}(G_I, \pi, P) < 3$ (see Figs. 7(b) and 7(c)).

Now, consider the case in which C_j has exactly two true literals under α, and let l_q be the only false literal of C_j under α for some $q \in \{1, 2, 3\}$. Thus, $v_j^p >_\pi v_j^q >_\pi v_j^r$, where $p = q \bmod 3 + 1$ and $r = (q + 1) \bmod 3 + 1$. If $i = l_q$, then $u_i^1 >_\pi u_i^2 >_\pi u_i^3$ and $V(H_i) <_\pi V(\widetilde{H}_j)$, which implies $\mathrm{zn}(G_I, \pi, P) < 3$ (see Fig. 8(a)). On the other hand, if $i = l_p$ or $i = l_r$, then $u_i^1 <_\pi u_i^2 <_\pi u_i^3$ and $V(H_i) >_\pi V(\widetilde{H}_j)$, which also implies $\mathrm{zn}(G_I, \pi, P) < 3$ (see Figs. 8(b) and 8(c)).

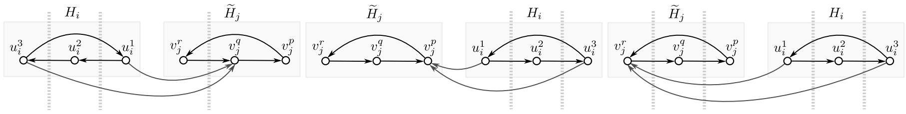

(a) $i = l_q$ (b) $i = l_p$ (c) $i = l_r$

Fig. 8. Case in which the clause $C_j = \{x_{l_1}, x_{l_2}, x_{l_3}\}$ has exactly one false literal under the truth assignment α, say x_{l_q} for some $q \in \{1, 2, 3\}$.

(a) (b)

Fig. 9. Case 1: $u_i^2 <_\pi u_i^1 <_\pi u_i^3$. (a) $v_j^q <_\pi u_i^1$. (b) $v_j^q >_\pi u_i^1$.

Therefore, such a path P does not exist in G_I, and consequently we obtain that $\mathrm{zn}(G_I) \leq \mathrm{zn}(G_I, \pi) \leq 2$. □

Lemma 5. *Let $I = (X, \mathcal{C})$ be an instance of* PNAE 3SAT, *$\pi : V(G_I) \to [|G_I|]$ be a bijection such that $zn(G_I, \pi) \leq 2$, and let $x_i \in X$. If $C_j \in \mathcal{C}$ is a clause containing x_i as a literal, then either $V(H_i) <_\pi V(\widetilde{H}_j)$ or $V(H_i) >_\pi V(\widetilde{H}_j)$. Furthermore, if there exists a clause $C_j \in \mathcal{C}$ containing x_i as a literal such that $V(H_i) <_\pi V(\widetilde{H}_j)$, then $V(H_i) <_\pi V(\widetilde{H}_{j'})$ for every other clause $C_{j'} \in \mathcal{C}$ containing x_i as a literal.*

Proof. Let $C_j \in \mathcal{C}$ be a clause containing x_i as a literal. For the sake of contradiction, suppose that neither $V(H_i) <_\pi V(\widetilde{H}_j)$ nor $V(H_i) >_\pi V(\widetilde{H}_j)$. Then, either there exist two vertices $v_j^p, v_j^{p'} \in V(\widetilde{H}_j)$ such that

$$\left\{v_j^p\right\} <_\pi V(H_i) <_\pi \left\{v_j^{p'}\right\}, \tag{2}$$

or there exists a vertex $v_j^p \in V(\widetilde{H}_j)$ such that, for some pair $a, b \in \{1, 2, 3\}$ with $a \neq b$,

$$u_i^a <_\pi v_j^p <_\pi u_i^b. \tag{3}$$

Hence, based on inequalities (2) and (3), one can verify that in either case there exist two (not necessarily distinct) vertices $v_j^p, v_j^{p'} \in V(\widetilde{H}_j)$ such that

$$v_j^p <_\pi \max{}_\pi V(H_i) \text{ and } v_j^{p'} >_\pi \min{}_\pi V(H_i) \tag{4}$$

In particular, we note that, $v_j^p = v_j^{p'}$ if and only if the second case – the one described by inequality (3) – holds.

Additionally, since by hypothesis x_i is a literal of C_j, there exists a vertex $v_j^q \in V(\widetilde{H}_j)$ such that $\left(u_i^1, v_j^q\right), \left(u_i^3, v_j^q\right) \in E(G_I)$. It is worth mentioning that, possibly, $v_j^p \neq v_j^q$ and $v_j^{p'} \neq v_j^q$. If $v_j^p \neq v_j^q$, then there exists a directed path $P_1' = \langle v_j^q, \ldots, v_j^p\rangle$ in $\widetilde{H}_j$ from v_j^q to v_j^p that has at least one edge. Analogously, if $v_j^{p'} \neq v_j^q$, then there exists a directed path $P_2' = \langle v_j^q, \ldots, v_j^{p'}\rangle$ in $\widetilde{H}_j$ from v_j^q to $v_j^{p'}$ that has at least one edge.

We split the remainder of this proof into three main cases.

(Case 1). Suppose that $u_i^2 <_\pi u_i^1 <_\pi u_i^3$. If $v_j^q <_\pi u_i^1$, then $P = \langle u_i^1, u_i^2, u_i^3, v_j^q\rangle$ is a directed path of G_I such that $\mathrm{zn}(G_I, \pi, P) = 3$ (see Fig. 9(a)). Similarly, if $v_j^q >_\pi u_i^1$, then $P = \langle u_i^2, u_i^3, u_i^1, v_j^q\rangle$ is a directed path of G_I such that $\mathrm{zn}(G_I, \pi, P) = 3$ (see Fig. 9(b)).

For the remainder cases, let $v_j^r \in V(\widetilde{H}_j)$ such that $\left(v_j^q, v_j^r\right) \in E(G_I)$. Note that, by construction, such a vertex v_j^r exists and, besides that, is well-defined.

(Case 2). Suppose that $u_i^3 <_\pi u_i^2 <_\pi u_i^1$. If $v_j^q <_\pi u_i^2$, then $P = \langle u_i^2, u_i^3, u_i^1, v_j^q\rangle$ is a directed path of G_I such that $\mathrm{zn}(G_I, \pi, P) = 3$ (see Fig. 10(a)). Assume that $v_j^q >_\pi u_i^2$. If $v_j^q >_\pi u_i^1$, then we obtain from (4) that $v_j^p \neq v_j^q$. Consequently, the path $P = \langle u_i^1, u_i^2, u_i^3, v_j^q, \ldots, v_j^p\rangle$ — obtained by concatenating $P'' = \langle u_i^1, u_i^2, u_i^3, v_j^q\rangle$ with P_1' — is a directed path of G_I such that $\mathrm{zn}(G_I, \pi, P) \geq 3$ (see Fig. 10(b)). Assume that $v_j^q <_\pi u_i^1$. If $v_j^r >_\pi v_j^q$, then $P = \langle u_i^3, u_i^1, v_j^q, v_j^r\rangle$ is a directed path of G_I such that $\mathrm{zn}(G_I, \pi, P) = 3$ (see Fig. 10(c)). Otherwise, if $v_j^r <_\pi v_j^q$, then $P = \langle u_i^1, u_i^2, u_i^3, v_j^q, v_j^r\rangle$ is a directed path of G_I such that $\mathrm{zn}(G_I, \pi, P) = 3$ (see Fig. 10(d)).

(Case 3). Suppose that $u_i^1 <_\pi u_i^3 <_\pi u_i^2$. If $v_j^q >_\pi u_i^3$, then $P = \langle u_i^1, u_i^2, u_i^3, v_j^q\rangle$ is a directed path of G_I such that $\mathrm{zn}(G_I, \pi, P) = 3$ (see Fig. 11(a)). Assume that $v_j^q <_\pi u_i^3$. If $v_j^q <_\pi u_i^1$, then we obtain from (4) that $v_j^{p'} \neq v_j^q$. Consequently, the path $P = \langle u_i^1, u_i^2, u_i^3, v_j^q, \ldots, v_j^{p'}\rangle$ — obtained by concatenating $P'' = \langle u_i^1, u_i^2, u_i^3, v_j^q\rangle$ with P_2' — is a directed path of G_I such that $\mathrm{zn}(G_I, \pi, P) \geq 3$ (see Fig. 11(b)). Assume that $v_j^q >_\pi u_i^1$. If $v_j^r <_\pi v_j^q$, then $P = \langle u_i^3, u_i^1, v_j^q, v_j^r\rangle$ is a directed path

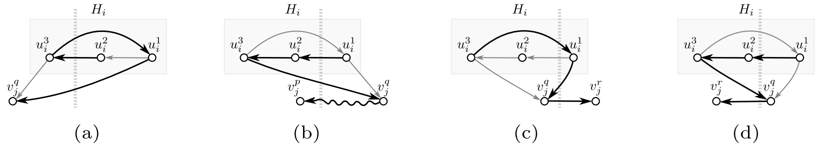

Fig. 10. Case 2: $u_i^3 <_\pi u_i^2 <_\pi u_i^1$. (a) $v_j^q <_\pi u_i^2$. (b) $v_j^q >_\pi u_i^2$ and $v_j^q >_\pi u_i^1$. (c) $u_i^2 <_\pi v_j^q <_\pi u_i^1$ and $v_j^r >_\pi v_j^q$. (d) $u_i^2 <_\pi v_j^q <_\pi u_i^1$ and $v_j^r <_\pi v_j^q$.

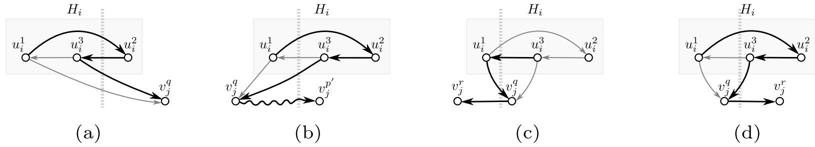

Fig. 11. Case 3: $u_i^1 <_\pi u_i^3 <_\pi u_i^2$. (a) $v_j^q >_\pi u_i^3$. (b) $v_j^q <_\pi u_i^3$ and $v_j^q <_\pi u_i^1$. (c) $u_i^1 <_\pi v_j^q <_\pi u_i^3$ and $v_j^r <_\pi v_j^q$. (d) $u_i^1 <_\pi v_j^q <_\pi u_i^3$ and $v_j^r >_\pi v_j^q$.

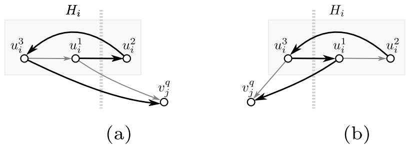

Fig. 12. Case in which $u_i^3 <_\pi u_i^1 <_\pi u_i^2$. (a) $v_j^q >_\pi u_i^1$. (b) $v_j^q <_\pi u_i^1$.

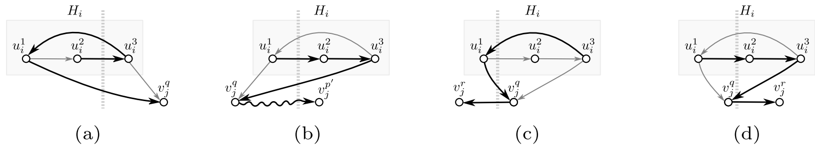

Fig. 13. Case in which $u_i^1 <_\pi u_i^2 <_\pi u_i^3$. (a) $v_j^q >_\pi u_i^2$. (b) $v_j^q <_\pi u_i^2$ and $v_j^q <_\pi u_i^1$. (c) $u_i^1 <_\pi v_j^q <_\pi u_i^2$ and $v_j^r <_\pi v_j^q$. (d) $u_i^1 <_\pi v_j^q <_\pi u_i^2$ and $v_j^r >_\pi v_j^q$.

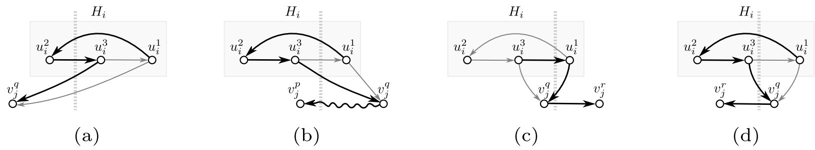

Fig. 14. Case in which $u_i^2 <_\pi u_i^3 <_\pi u_i^1$. (a) $v_j^q <_\pi u_i^3$. (b) $v_j^q >_\pi u_i^3$ and $v_j^q >_\pi u_i^1$. (c) $u_i^3 <_\pi v_j^q <_\pi u_i^1$ and $v_j^r >_\pi v_j^q$. (d) $u_i^3 <_\pi v_j^q <_\pi u_i^1$ and $v_j^r <_\pi v_j^q$.

of G_I such that $\mathrm{zn}(G_I, \pi, P) = 3$ (see Fig. 11(c)). Otherwise, if $v_j^r >_\pi v_j^q$, then $P = \langle u_i^1, u_i^2, u_i^3, v_j^q, v_j^r \rangle$ is a directed path of G_I such that $\mathrm{zn}(G_I, \pi, P) = 3$ (see Fig. 11(d)).

One can readily verify that the case in which $u_i^3 <_\pi u_i^1 <_\pi u_i^2$ is symmetric to Case 1 (see Fig. 12), the case in which $u_i^1 <_\pi u_i^2 <_\pi u_i^3$ is symmetric to Case 2 (see Fig. 13), and the case in which $u_i^2 <_\pi u_i^3 <_\pi u_i^1$ is symmetric to Case 3 (see Fig. 14). Additionally, note that, regardless of the existence of the vertex v_j^p, Case 1 and consequently the case in which $u_i^3 <_\pi u_i^1 <_\pi u_i^2$ do not consist in valid configurations, otherwise $\mathrm{zn}(G, \pi) \geq 3$ even if $V(H_i) <_\pi V(\widetilde{H}_j)$ or $V(H_i) >_\pi V(\widetilde{H}_j)$.

Thus, $V(H_i) <_\pi V(\widetilde{H}_j)$ or $V(H_i) >_\pi V(\widetilde{H}_j)$, otherwise $\mathrm{zn}(G, \pi) \geq 3$. Particularly, one can further verify that if $u_i^3 <_\pi u_i^2 <_\pi u_i^1$ or $u_i^2 <_\pi u_i^3 <_\pi u_i^1$, then necessarily $V(H_i) <_\pi V(\widetilde{H}_j)$. Analogously, we have that if $u_i^1 <_\pi u_i^3 <_\pi u_i^2$ or $u_i^1 <_\pi u_i^2 <_\pi u_i^3$, then $V(H_i) >_\pi V(\widetilde{H}_j)$. Therefore, if $V(H_i) <_\pi V(\widetilde{H}_j)$, then $V(H_i) <_\pi V(\widetilde{H}_{j'})$ for every other clause $C_{j'} \in \mathcal{C}$ containing x_i as a literal. $\square$

Lemma 6. *Let $I = (X, \mathcal{C})$ be an instance of* PNAE 3SAT. *If $zn(G_I) \leq 2$, then I is a* yes *instance of* PNAE 3SAT.

Proof. Let $\pi : V(G_I) \to [|G_I|]$ be a bijection such that $\mathrm{zn}(G_I, \pi) \leq 2$. It follows from Lemma 5 that, for each variable $x_i \in X$ and each clause $C_j \in \mathcal{C}$, if $V(H_i) >_\pi V(\widetilde{H}_j)$, then $V(H_i) >_\pi V(\widetilde{H}_{j'})$ for each clause $C_{j'} \in \mathcal{C}$ containing x_i as a literal.

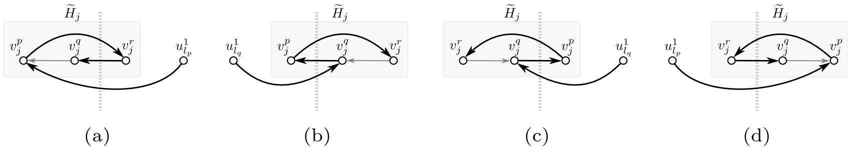

Fig. 15. (a) and (c) $\alpha(x_{l_1}) = \alpha(x_{l_2}) = \alpha(x_{l_3}) = 1$. (b) and (d) $\alpha(x_{l_1}) = \alpha(x_{l_2}) = \alpha(x_{l_3}) = 0$. (a) and (b) $v_j^p <_\pi v_j^q <_\pi v_j^r$. (c) and (d) $v_j^p >_\pi v_j^q >_\pi v_j^r$.

Thus, we let $\alpha : x \to \{0, 1\}$ be the truth assignment defined as follows: for each variable $x_i \in X$, $\alpha(x_i) = 1$ if and only if $V(H_i) >_\pi V(\widetilde{H}_j)$ for each clause $C_j \in \mathcal{C}$.

Now, we prove that each clause in $\mathcal{C}$ has at least one true literal and at least one false literal under α. For the sake of contradiction, suppose that there exists a clause $C_j = \left\{x_{l_1}, x_{l_2}, x_{l_3}\right\}$ in $\mathcal{C}$ such that $\alpha(x_{l_1}) = \alpha(x_{l_2}) = \alpha(x_{l_3})$. Let $q \in \{1, 2, 3\}$, $p = q \bmod 3 + 1$ and $r = (q + 1) \bmod 3 + 1$.

Suppose that $\alpha(x_{l_1}) = \alpha(x_{l_2}) = \alpha(x_{l_3}) = 1$. Thus, $\left\{u_{l_1}^1, u_{l_2}^1, u_{l_3}^1\right\} >_\pi V(\widetilde{H}_j)$. Consequently, if $v_j^p <_\pi v_j^q <_\pi v_j^r$, then $P = \langle u_{l_p}^1, v_j^p, v_j^r, v_j^q \rangle$ is a directed path of G_I such that $\mathsf{zn}(G_I, \pi, P) = 3$ (see Figs. 15(a)); on the other hand, if $v_j^p >_\pi v_j^q >_\pi v_j^r$, then $P = \langle u_{l_q}^1, v_j^q, v_j^p, v_j^r \rangle$ is a directed path of G_I such that $\mathsf{zn}(G_I, \pi, P) = 3$ (see Figs. 15(c)).

Suppose that $\alpha(x_{l_1}) = \alpha(x_{l_2}) = \alpha(x_{l_3}) = 0$. Thus, $\left\{u_{l_1}^1, u_{l_2}^1, u_{l_3}^1\right\} <_\pi V(\widetilde{H}_j)$. Consequently, if $v_j^p <_\pi v_j^q <_\pi v_j^r$, then $P = \langle u_{l_q}^1, v_j^q, v_j^p, v_j^r \rangle$ is a directed path of G_I such that $\mathsf{zn}(G_I, \pi, P) = 3$ (see Figs. 15(b)); on the other hand, if $v_j^p >_\pi v_j^q >_\pi v_j^r$, then $P = \langle u_{l_p}^1, v_j^p, v_j^r, v_j^q \rangle$ is a directed path of G_I such that $\mathsf{zn}(G_I, \pi, P) = 3$ (see Figs. 15(d)).

Therefore, each clause in $\mathcal{C}$ has at least one true literal and at least one false literal under α, and consequently I is a yes instance of PNAE 3SAT. □

Theorem 2. 2-ZIG-ZAG NUMBER *is NP-complete.*

Proof. By Theorem 1, 2-ZIG-ZAG NUMBER is in NP. It follows from Lemmas 4 and 6 that I is a yes instance of PNAE 3SAT if and only if $\mathsf{zn}(G_I) \leq 2$. Therefore, since G_I can be constructed in time polynomial in $|I|$, 2-ZIG-ZAG NUMBER is NP-complete. □

5. Zig-zag number and directed treewidth

It was proved in [13] that directed graphs of constant directed pathwidth have constant zig-zag number, and that there exist directed graphs of constant zig-zag number but unbounded directed pathwidth. Hence, the family of directed graphs of constant zig-zag number properly contains the family of directed graphs of constant directed pathwidth.

Nevertheless, it is unknown whether or not a similar result would hold with respect to zig-zag number and directed treewidth. In this section, we prove that there exist directed graphs of constant directed treewidth but unbounded zig-zag number. More specifically, we prove the following theorem.

Theorem 3. *There exist directed graphs on n vertices of constant directed treewidth but zig-zag number $\Omega(\log n)$.*

We remark that, although it was shown in [2] that there are directed graphs of constant directed treewidth but unbounded directed pathwidth, this result cannot be directly used to conclude the respective statement relating directed treewidth and zig-zag number.

Another interesting aspect of our result follows from the fact that directed graphs of constant directed treewidth have constant tree-zig-zag number [14]. Consequently, there are directed graphs of constant tree-zig-zag number but unbounded zig-zag number. Therefore, considering the fact that directed graphs of constant zig-zag number have constant tree-zig-zag number [14], we obtain that the family of directed graphs of constant tree-zig-zag number is strictly richer than the family of directed graphs of constant zig-zag number.

5.1. Basic definitions

A directed graph G is called *bidirected* if, for each two distinct vertices $u, v \in V(G)$, $(u, v) \in E(G)$ if and only if $(v, u) \in E(G)$. Note that, bidirected graphs may be regarded as undirected graphs. Based on that, we say that a pair of edges (u, v) and (u, v) of a bidirected graph G is a *bidirected edge* between u and v. A directed graph H is called an *undirected minor* of a bidirected graph G if H can be obtained from G by deleting vertices and edges, and by contracting bidirected edges.

An *undirected tree decomposition* of a directed graph G is a pair $\left(T, \{X_t\}_{t \in V(T)}\right)$ satisfying the following conditions:

1. T is a undirected tree;
2. $\bigcup_{t \in V(T)} X_t = V(G)$;

3. for each edge $(u, v) \in E(G)$, there exists a node $t \in V(T)$ such that $\{u, v\} \subseteq X_t$;
4. for each vertex $u \in V(G)$, the graph $T[\{t \in V(T) : u \in X_t\}]$ is connected.

In particular, $\left(T, \{X_t\}_{t\in V(T)}\right)$ is called an *undirected path decomposition* of G if T is an undirected path. The *width* of an undirected tree decomposition $\left(T, \{X_t\}_{t\in V(T)}\right)$ is defined as the integer $\max_{t\in V(T)}|X_t| - 1$. The *undirected treewidth* of a directed graph G is defined as the minimum width over all tree decompositions of G, and the *undirected pathwidth* of G is defined as the minimum width over all path decompositions of G.

Throughout this section we are mainly concerned with bidirected graphs. Thus, based on Lemmas 7 and 8, stated next, it suffices to define only the notions of undirected pathwidth and of undirected treewidth. For the definitions of directed pathwidth and directed treewidth, we refer to Refs. [12,15,18].

Lemma 7 ([1]). *If G is a bidirected graph, then the directed pathwidth of G is equal to its undirected pathwidth.*

Lemma 8 ([12]). *If G is a bidirected graph, then the directed treewidth of G is equal to its undirected treewidth.*

For the sake of simplicity, we also omit the formal definition of tree-zig-zag number, and we refer to Ref. [14]. Informally, the tree-zig-zag number of a directed graph G is defined similarly to the zig-zag number of G except for, instead of linear orders, considering *binary arboreal orders*. That is to say, partial orders $<_\pi \subseteq V(G) \times V(G)$ such that, for each vertex $v \in V(G)$, the following conditions hold: the set $\{u \in V(G) : u <_\pi v\}$ is linearly ordered by $<_\pi$ and there are at most two vertices $v' \in V(G)$ with $v <_\pi v'$ such that, for any $u \in V(G)$, $u <_\pi v'$ if and only if $u <_\pi v$.

5.2. Directed vertex separation number

Let G be a directed graph on n vertices and $\pi : V(G) \to [n]$ be a bijection. Assume that $V(G) = \{u_1, \ldots, u_n\}$ and, for each $u_i, u_j \in V(G)$, $i < j$ if and only if $u_i <_\pi u_j$. The *directed vertex separation number of G with respect to* π is defined as the maximum number of vertices in $\{u_{i+1}, \ldots, u_n\}$ that have some out-neighbor in $\{u_1, \ldots, u_i\}$, where the maximum is taken over all $i \in [n-1]$. More formally,

$$\mathsf{dvsn}(G, \pi) \doteq \max_{i\in[n-1]} \left|\left\{v \in \{u_{i+1}, \ldots, u_n\} : N^+_G(v) \cap \{u_1, \ldots, u_i\} \neq \emptyset\right\}\right|,$$

where $N^+_G(v)$ denotes the out-neighborhood of v in G. The *directed vertex separation number of G*, denoted by $\mathsf{dvsn}(G)$, is defined as the minimum $\mathsf{dvsn}(G, \pi)$ over all bijections $\pi : V(G) \to [n]$.

Lemma 9 ([19]). *Let G be a directed graph. The directed pathwidth of G is equal to the directed vertex separation number of G.*

Lemma 10. *If G_1 and G_2 are two directed graphs over a same vertex set X, and $\pi : X \to [|X|]$ is a bijection, then $\mathsf{dvsn}(G_1 \cup G_2, \pi) \le \mathsf{dvsn}(G_1, \pi) + \mathsf{dvsn}(G_2, \pi)$.*

Proof. Assume that $X = \{u_1, \ldots, u_n\}$ and, for each $u_i, u_j \in V(G)$, $i < j$ if and only if $u_i <_\pi u_j$. Let $i \in [n-1]$. Suppose that there exist ℓ_1 distinct vertices $v \in \{u_{i+1}, \ldots, u_n\}$ such that $N^+_{G_1}(v) \cap \{u_1, \ldots, u_i\} \neq \emptyset$. Analogously, suppose that there exist ℓ_2 distinct vertices $v \in \{u_{i+1}, \ldots, u_n\}$ such that $N^+_{G_2}(v) \cap \{u_1, \ldots, u_i\} \neq \emptyset$. As a result, there exist at most $\ell_1 + \ell_2$ distinct vertices $v \in \{u_{i+1}, \ldots, u_n\}$ such that $N^+_{G_1}(v) \cap \{u_1, \ldots, u_i\} \neq \emptyset$ or $N^+_{G_2}(v) \cap \{u_1, \ldots, u_i\} \neq \emptyset$. In other words, there exist at most $\ell_1 + \ell_2$ distinct vertices $v \in \{u_{i+1}, \ldots, u_n\}$ such that $N^+_{G_1\cup G_2}(v) \cap \{u_1, \ldots, u_i\} \neq \emptyset$. Therefore, $\mathsf{dvsn}(G_1 \cup G_2, \pi) \le \mathsf{dvsn}(G_1, \pi) + \mathsf{dvsn}(G_2, \pi)$. □

Lemma 11. *Let G be a directed graph, P be a directed path of G, and let H be the directed graph such that $V(H) = V(G)$ and $E(H) = E(P)$. Then, for each bijection $\pi : V(G) \to [|G|]$, $\mathsf{dvsn}(H, \pi) \le \mathsf{zn}(H, \pi) \le \mathsf{zn}(G, \pi)$.*

Proof. The second inequality follows from the fact that H is a subgraph of G. Now, we prove that the first inequality holds. Assume that $V(G) = \{u_1, \ldots, u_n\}$ and, for each $u_i, u_j \in V(G)$, $i < j$ if and only if $u_i <_\pi u_j$. Let $i \in [n-1]$. Since the set of edges of H induces a directed path, it follows from the definition of zig-zag number that there exist at most $\mathsf{zn}(H, \pi)$ edges in the cut $S_H(\pi, i)$. As a result, there exist at most $\mathsf{zn}(H, \pi)$ vertices $v \in \{u_{i+1}, \ldots, u_n\}$ such that $N^+_H(v) \cap \{u_1, \ldots, u_i\} \neq \emptyset$. Therefore, $\mathsf{dvsn}(H, \pi) \le \mathsf{zn}(H, \pi)$. □

Lemma 12. *Let p be a positive integer, G be a directed graph and $\pi : V(G) \to [|G|]$ be a bijection. If G can be described as the union of p directed paths, then $\mathsf{dvsn}(G, \pi) \le p \cdot \mathsf{zn}(G, \pi)$.*

Proof. Suppose that there are p directed paths $P_1, \ldots, P_p$ such that $G = \bigcup_{i\in[p]} P_i$. For each $i \in [p]$, let H_i be the directed graph with vertex set $V(H_i) = V(G)$ and edge set $E(H_i) = E(P_i)$. Note that, $G = \bigcup_{i\in[p]} H_i$. It follows from Lemma 11 that, for each $i \in [p]$, $\mathsf{dvsn}(H_i, \pi) \le \mathsf{zn}(H_i, \pi)$. Therefore, by Lemma 10, $\mathsf{dvsn}(G, \pi) \le p \cdot \mathsf{zn}(G, \pi)$. □

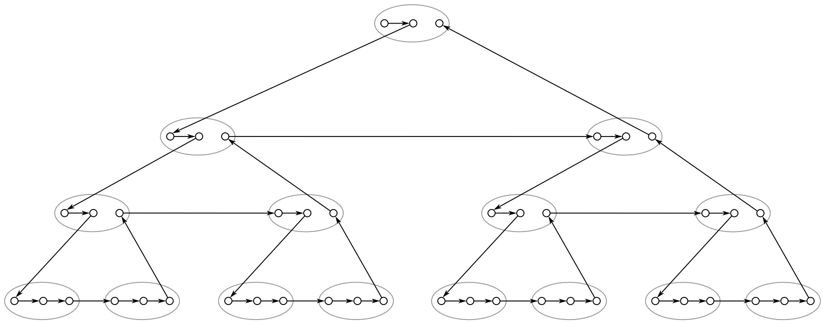

(a) Directed path P_1.

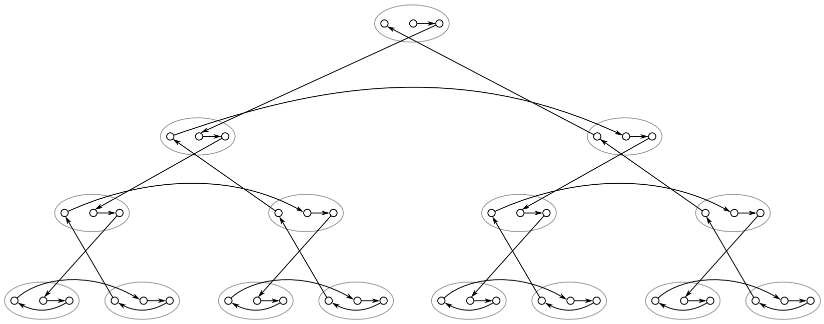

(b) Directed path P_2.

Fig. 16. Directed paths P_1 and P_2 obtained by Algorithm 1, respectively.

5.3. Proof of Theorem 3

Let B_n be a rooted oriented complete binary tree on n vertices. For each non-leaf vertex $u \in V(B_n)$, we write $\mathsf{left}(u)$ to denote the left child of u in B_n, and we write $\mathsf{right}(u)$ to denote the right child of u in B_n.

We let $\mathbf{B}_n$ be the bidirected graph with vertex set $V(\mathbf{B}_n) = V(B_n) \times \{0, 1, 2\}$ obtained by the union of two suitable directed paths P_1 and P_2, recursively defined in Algorithm 1, and their respective reverse directed paths P_1' and P_2'. More specifically, if r is the root of B_n, then P_1 and P_2 are defined as the directed paths returned by the function calls

- $\texttt{Construct-Path}(B_n, P = \langle\rangle, u = r, \texttt{idx} = 1)$ and
- $\texttt{Construct-Path}(B_n, P = \langle\rangle, u = r, \texttt{idx} = 2)$

of Algorithm 1, respectively, and P_1' is the reverse directed path of P_1 and P_2' is the reverse directed path of P_2. Therefore, $\mathbf{B}_n$ can be decomposed into four directed paths. Fig. 16 illustrates P_1 and P_2.

Algorithm 1: Construction of directed graph $\mathbf{B}_n$.

```
function Construct-Path(B_n, P, u, idx)
    a = (idx − 1) mod 3; b = idx mod 3; c = (idx + 1) mod 3
    P := P + ⟨(u, a) , (u, b)⟩            // concatenates P with the sequence ⟨(u, a) , (u, b)⟩
    if u is not a leaf of B_n then
        P := Construct-Path(B_n, P, left(u), idx)
        P := Construct-Path(B_n, P, right(u), idx)
    P := P + ⟨(u, c)⟩                      // concatenates P with the sequence ⟨(u, c)⟩
    return P
```

Lemma 13. *The complete binary tree B_n is an undirected minor of $\mathbf{B}_n$.*

Proof. First, we note that, for each $i \in \{1, 2\}$, if there exists a directed edge (u, v) in the directed path P_i, then there exists the directed edge (v, u) in the reverse directed path P_i' of P_i. This implies that, whenever (u, v) is a directed edge in P_i for some $i \in \{1, 2\}$, $\mathbf{B}_n$ contains a bidirected edge between the nodes u and v. As a result, in order to show the existence of a bidirected edge in $\mathbf{B}_n$ between nodes u and v, it is enough to show that either (u, v) or (v, u) is a directed edge in P_i for some $i \in \{1, 2\}$.

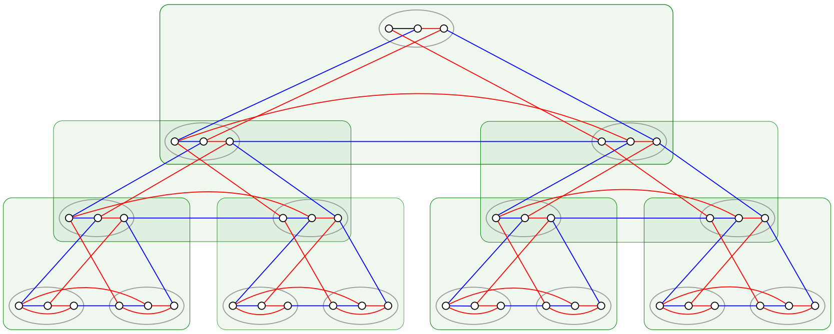

Fig. 17. Tree decomposition $\mathcal{T} = (T, \mathcal{X})$ of $\mathbf{B}_n$, where the rounded squares represent the bags of the nodes of T.

Based on that, we now remark that, for each node $t \in V(B_n)$, there exist in $\mathbf{B}_n$ a bidirected edge between the nodes $(t, 0)$ and $(t, 1)$, and a bidirected edge between the nodes $(t, 1)$ and $(t, 2)$. Indeed, this follows from the facts that $(t, 1)$ immediately succeeds $(t, 0)$ in the directed path P_1, and that $(t, 2)$ immediately succeeds $(t, 1)$ in the directed path P_2. In addition, it follows from construction of P_1 that, for each non-leaf node $t \in V(B_n)$, there exist in $\mathbf{B}_n$ an edge between the nodes $(t, 1)$ and $(\mathsf{left}(t), 0)$, and an edge between the nodes $(\mathsf{right}(t), 2)$ and $(t, 2)$. Similarly, it follows from construction of P_2 that, for each non-leaf node $t \in V(B_n)$, there exist in $\mathbf{B}_n$ an edge between the nodes $(t, 2)$ and $(\mathsf{left}(t), 1)$, and an edge between the nodes $(\mathsf{right}(t), 0)$ and $(t, 0)$.

Thus, let $\mathsf{contr}(\mathbf{B}_n)$ denote the bidirected graph obtained from $\mathbf{B}_n$ by contracting, for each node $t \in V(\mathbf{B}_n)$, the bidirected edge between $(t, 0)$ and $(t, 1)$, and the bidirected edge between $(t, 1)$ and $(t, 2)$. By the discussion above, $\mathsf{contr}(\mathbf{B}_n)$ is a subgraph of B_n. Therefore, B_n is an undirected minor of $\mathbf{B}_n$. □

Lemma 14. *$\mathsf{zn}(\mathbf{B}_n) = \Omega(\log n)$.*

Proof. By Lemma 13, B_n is an undirected minor of $\mathbf{B}_n$. As a result, we obtain that the undirected pathwidth of $\mathbf{B}_n$ is at least the undirected pathwidth of B_n *cf.* [4]. Moreover, it is well-known that complete binary trees on n vertices have undirected pathwidth $\Theta(\log n)$ *cf.* [2,14]. Thus, since $\mathbf{B}_n$ is a bidirected graph, it follows from Lemma 7 that the directed pathwidth of $\mathbf{B}_n$ is $\Omega(\log n)$. Consequently, by Lemma 9, $\mathsf{dvsn}(\mathbf{B}_n) = \Omega(\log n)$. Moreover, by construction, $\mathbf{B}_n$ can be described as the union of 4 directed paths. Then, it follows from Lemma 12 that $\mathsf{zn}(\mathbf{B}_n) \geq \frac{\mathsf{dvsn}(\mathbf{B}_n)}{4}$. Therefore, $\mathsf{zn}(\mathbf{B}_n) = \Omega(\log n)$. □

Lemma 15. *For each positive integer n, $\mathbf{B}_n$ has directed treewidth at most 8.*

Proof. Based on Lemma 8, it suffices to prove that $\mathbf{B}_n$ admits a undirected tree decomposition $\mathcal{T} = (T, \mathcal{X})$ of width 8, where $\mathcal{X} = (X_t)_{t \in V(T)}$.

We define T simply as the complete binary tree obtained from B_n by removing all its leaves. Then, for each node $t \in V(T)$, we define the bag of t as the set $X_t = \{t, \mathsf{left}(t), \mathsf{right}(t)\} \times \{0, 1, 2\}$. Fig. 17 illustrates the tree decomposition $\mathcal{T} = (T, \mathcal{X})$. One can verify that $\mathcal{T}$ is a tree decomposition of $\mathbf{B}_n$ of width 8.

Based on Lemmas 14 and 15, we conclude the proof of Theorem 3.

6. Concluding remarks

We have shed new light on the time complexity of computing the zig-zag number of a directed graph. Nonetheless, some questions still remain open.

More specifically, we have proved that one can non-deterministically decide whether a directed graph G admits zig-zag number at most k in time $|G|^{\mathcal{O}(k)}$, concluding that k-ZIG-ZAG NUMBER is in NP for each fixed $k \geq 0$. Nevertheless, it remains unknown whether k-ZIG-ZAG NUMBER admits a non-deterministic FPT-time algorithm. Another interesting question concerns to determine whether ZIG-ZAG NUMBER is also in NP for non-fixed k. It is worth mentioning that, to settle k-ZIG-ZAG NUMBER in NP, we have actually proved that, given a directed graph G and a bijection $\pi : V(G) \to [|G|]$, deciding whether $\mathsf{zn}(G, \pi) \leq k$ is polynomial-time solvable for fixed k. However, for non-fixed k, deciding whether $\mathsf{zn}(G, \pi) \leq k$ is coNP-complete. As a matter of fact, given a bipartite directed graph G with bipartition $V(G) = X \cup Y$, if $\pi : V(G) \to [|G|]$ is defined in such a way that $x <_\pi y$ for each $x \in X$ and each $y \in Y$, then deciding whether $\mathsf{zn}(G, \pi) \geq |G| - 1$ is equivalent to deciding whether G admits a Hamiltonian path, which is a well-known NP-complete problem [11].

Another intriguing question concerns to determine whether 1-ZIG-ZAG NUMBER is polynomial-time solvable. As already mentioned, every directed acyclic graph has zig-zag number at most 1, and every directed graph containing directed cycles of length at least 3 must have zig-zag number at least 2. However, there exist directed graphs that are not directed acyclic

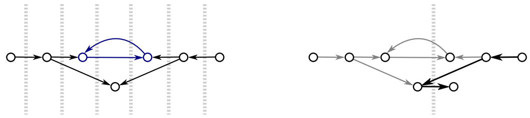

(a) $\mathsf{zn}(G) = 1$. (b) $\mathsf{zn}(G) \geq 2$.

Fig. 18. (a) Example of directed graph G that is not directed acyclic and has zig-zag number 1. (b) Example of directed graph G that does not contain directed cycles of length at least 3 and yet has zig-zag number 2.

but still have zig-zag number at most 1 (see Fig. 18(a)). Note that, such graphs can only contain directed cycles that are *digons*, *i.e.* directed cycles of length 2. Nevertheless, this is not a sufficient condition for a directed graph to have zig-zag number at most 1. In fact, there exist directed graphs that only contain directed cycles that are digons and yet have zig-zag number at least 2 (see Fig. 18(b)). A property that seems to be useful to resolve this problem is the fact that, for every directed graph G, $\mathsf{zn}(G) \leq 1$ if and only if there exists a bijection $\pi : V(G) \to [|G|]$ such that, for each three distinct vertices $a, b, c \in V(G)$, with $(a, b), (b, c) \in E(G)$, either $a <_\pi b <_\pi c$ or $c <_\pi b <_\pi a$.

Motivated by the NP-hardness of 2-ZIG-ZAG NUMBER, we additionally ask whether k-ZIG-ZAG NUMBER is NP-hard for $k \geq 3$. In particular, determining whether k-ZIG-ZAG NUMBER is polynomially reducible to $(k+1)$-ZIG-ZAG NUMBER is an elusive open problem. Generally, such a reduction must consist in constructing a directed graph H from a given directed graph G, such that $\mathsf{zn}(H) = \mathsf{zn}(G) + 1$. However, since for distinct bijections $\pi : V(G) \to [|G|]$ there might exist distinct directed paths P of G such that $\mathsf{zn}(G, \pi, P) = \mathsf{zn}(G, \pi)$, it is not obvious at all how G should be modified so as to produce a directed graph with zig-zag number exactly one unit greater than $\mathsf{zn}(G)$. Indeed, consider for instance the operation of adding a *universal vertex*, *i.e.* a vertex that is an out-neighbor and an in-neighbor of all the other vertices. There exist directed graphs G such that the addition of a universal vertex results in a directed graph with zig-zag number strictly greater than $\mathsf{zn}(G) + 1$; while there also exist directed graphs G such that the addition of a universal vertex results in a directed graph with zig-zag number equal to $\mathsf{zn}(G)$.

It is worth mentioning that, even if k-ZIG-ZAG NUMBER is proved to be NP-hard for every $k \geq 3$, zig-zag number is still a directed width measure of important theoretical and algorithmic interest. Indeed, besides the fact that zig-zag number is asymptotically upper bounded by directed pathwidth, there possibly exist efficient approximation algorithms with constant approximation factors for the k-ZIG-ZAG NUMBER problem, which remains an open question. Motivated by that, we ask for the existence of such approximation algorithms.

Finally, other important questions that are still open concern the establishment of relations between zig-zag number and distinct width measures. We have proved that there exist directed graphs of constant directed treewidth but unbounded zig-zag number. However, it is unknown whether the family of directed graphs of constant directed treewidth contains the family of directed graphs of constant zig-zag number. We remark that a counter-example for such containment would also imply the existence of directed graphs of constant tree-zig-zag number but unbounded directed treewidth, closing an open question from [14]. Related to this, we ask for the existence of a characterization of zig-zag number in terms of pursuit games.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by CAPES, Brazil (Finance Code: 001 and Grant Number : PDSE 88881.187636/2018-01), CNPq, Brazil (Grant Numbers: 303546/2016-6, 407635/2018-1, 140399/2017-8, and 303726/2017-2), FAPERJ, Brazil (Grant Numbers: E-26/202.793/2017 and E-26/203.272/2017), Bergen Research Foundation, and by the Research Council of Norway (Grant Number: 288761).

References

[1] J. Barát, Directed path-width and monotonicity in digraph searching, Graphs Comb. 22 (2) (2006) 161–172.
[2] D. Berwanger, A. Dawar, P. Hunter, S. Kreutzer, J. Obdrzálek, The DAG-width of directed graphs, J. Combin. Theory Ser. B 102 (4) (2012) 900–923.
[3] D. Berwanger, E. Grädel, L. Kaiser, R. Rabinovich, Entanglement and the complexity of directed graphs, Theoret. Comput. Sci. 463 (2012) 2–25.
[4] H.L. Bodlaender, A partial k-arboretum of graphs with bounded treewidth, Theoret. Comput. Sci. 209 (1–2) (1998) 1–45.
[5] B. Courcelle, The monadic second-order logic of graphs. I. Recognizable sets of finite graphs, Inform. and Comput. 85 (1) (1990) 12–75.
[6] B. Courcelle, J.A. Makowsky, U. Rotics, Linear time solvable optimization problems on graphs of bounded clique-width, Theory Comput. Syst. 33 (2) (2000) 125–150.
[7] R. Ganian, P. Hlinený, J. Kneis, A. Langer, J. Obdrzálek, P. Rossmanith, On digraph width measures in parameterized algorithmics, in: IWPEC, in: LNCS, vol. 5917, 2009, pp. 185–197.

[8] R. Ganian, P. Hliněný, J. Kneis, D. Meister, J. Obdržálek, P. Rossmanith, S. Sikdar, Are there any good digraph width measures? J. Combin. Theory Ser. B 116 (2016) 250–286.
[9] H. Gruber, M. Holzer, Finite automata, digraph connectivity, and regular expression size, in: ICALP, in: LNCS, vol. 5126, 2008, pp. 39–50.
[10] P. Hunter, S. Kreutzer, Digraph measures: Kelly decompositions, games, and orderings, Theoret. Comput. Sci. 399 (3) (2008) 206–219.
[11] A. Itai, C.H. Papadimitriou, J.L. Szwarcfiter, Hamilton paths in grid graphs, SIAM J. Comput. 11 (4) (1982) 676–686.
[12] T. Johnson, N. Robertson, P.D. Seymour, R. Thomas, Directed tree-width, J. Combin. Theory Ser. B 82 (1) (2001) 138–154.
[13] M. de Oliveira Oliveira, Subgraphs satisfying MSO properties on z-topologically orderable digraphs, in: IPEC, in: LNCS, vol. 8246, Springer, 2013, pp. 123–136.
[14] M. de Oliveira Oliveira, An algorithmic metatheorem for directed treewidth, Discrete Appl. Math. 204 (2016) 49–76.
[15] B.A. Reed, Introducing directed tree width, in: CTW, in: ENDM, vol. 3, 1999, pp. 222–229.
[16] M.A. Safari, D-width: A more natural measure for directed tree width, in: MFCS, in: LNCS, vol. 3618, 2005, pp. 745–756.
[17] T.J. Schaefer, The complexity of satisfiability problems, in: Proceedings of the Tenth Annual ACM Symposium on Theory of Computing, STOC'78, ACM, 1978, pp. 216–226.
[18] H. Tamaki, A polynomial time algorithm for bounded directed pathwidth, in: WG, in: LNCS, vol. 6986, 2011, pp. 331–342.
[19] B. Yang, Y. Cao, Digraph searching, directed vertex separation and directed pathwidth, Discrete Appl. Math. 156 (10) (2008) 1822–1837.