Are there any good digraph width measures?[†]

Robert Ganian[‡] Petr Hliněný[‡] Joachim Kneis[§] Daniel Meister[§]
Jan Obdržálek[‡] Peter Rossmanith[§]
Somnath Sikdar[§]

[‡]Faculty of Informatics, Masaryk University, Brno, Czech Republic.
{ganian,hlineny,obdrzalek}@fi.muni.cz
[§]Theoretical Computer Science, RWTH Aachen University, Germany.
{kneis,meister,rossmani,sikdar}@cs.rwth-aachen.de

Abstract

Many width measures for directed graphs have been proposed in the last few years in pursuit of generalizing (the notion of) treewidth to directed graphs. However, none of these measures possesses, at the same time, the major properties of treewidth, namely,

1. being *algorithmically useful*, that is, admitting polynomial-time algorithms for a large class of problems on digraphs of bounded width (e.g. the problems definable in MSO_1);
2. having nice *structural properties* such as being (at least nearly) monotone under taking subdigraphs and some form of arc contractions (property closely related to characterizability by particular cops-and-robber games).

We investigate the question whether the search for directed treewidth counterparts has been unsuccessful by accident, or whether it has been doomed to fail from the beginning. Our main result states that any *reasonable* width measure for directed graphs which satisfies the two properties above must necessarily be similar to treewidth of the underlying undirected graph.

1 Introduction

An intensely investigated field in algorithmic graph theory is the design of graph *width parameters* that satisfy two seemingly contradictory requirements: (1) a large class of problems must be efficiently solvable on the graphs of bounded width and (2) graphs of bounded width should have a nice, reasonably rich and natural structure.

For undirected graphs, research into width parameters has been extremely successful with several algorithmically useful measures being proposed over the years, chief among them being

[†]This work has been supported by the Deutsche Forschungsgemeinschaft (project RO 927/9) and the Czech Science Foundation (project P201/09/J021 till 2010). Currently, the Brno team is supported by the research centre Institute for Theoretical Computer Science (CE-ITI), project P202/12/G061.

treewidth [35], branchwidth [36] and clique-width [9] (see also [5, 21]). Many problems that are hard on general graphs turned out to be tractable on graphs of bounded treewidth. These results were combined and generalized into Courcelle's celebrated theorem which states that a (very) large class of problems, i.e. the class of all MSO_2-definable problems, is tractable on graphs of bounded treewidth [7]. Treewidth and branchwidth are closely related and their very nice structural properties are well known. Furthermore, the tractability of all MSO_2 problems closely characterizes the class of graphs of bounded treewidth [28].

However, for *directed graphs* no single known *width measure* is as successful as treewidth is for undirected graphs. We feel the reason for this is that all the currently known digraph width measures fail on one of the aforementioned conditions (1) and (2).

During the last decade, many digraph width measures inspired by treewidth were introduced, the prominent ones being directed treewidth [24], DAG-width [4, 32], and Kelly-width [23]. These width measures proved useful for some problems. For instance, one can obtain polynomial-time (XP to be more precise, see Section 2) algorithms for Hamiltonian Path on digraphs of bounded directed treewidth [24] and for Parity Games on digraphs of bounded DAG-width [4] and Kelly-width [23]. But there is the negative side, too. Hamiltonian Path, for instance, probably cannot be solved [29] on digraphs of directed treewidth, DAG-width, or Kelly-width at most k in time $\mathcal{O}(f(k) \cdot n^c)$, where c is a constant independent of k. Note that Hamiltonian Path *can* be solved in such a running time for undirected graphs of treewidth at most k [7].

Moreover, for the former ones and even new [15] measures DAG-depth and Kenny-width[1] which are much more restrictive than DAG-width; (1) problems such as Directed Dominating Set, Directed Cut, Oriented Chromatic Number 4, Max / Min Leaf Outbranching, or k-Path remain NP-complete on digraphs of constant width [15]. In contrast, clique-width and the recent digraph measure bi-rank-width [25] look more promising on the algorithmic side: A Courcelle-like [8] MSO_1 theorems exist for digraphs of bounded directed clique-width and bi-rank-width, and many other interesting problems can be solved in polynomial time (XP) on digraphs of bounded directed clique-width and bi-rank-width [19, 26, 18]. Yet, (2) the latter measures are not monotone even under taking subdigraphs. For a recent exhaustive survey on complexity results for DAG-width, Kelly-width, bi-rank-width, and other digraph measures, see [16].

In this paper, we show (under well-established complexity assumptions) that any *reasonable* digraph width measure that is algorithmically useful (*powerful*; Definition 3.1) and is closed under a notion of *directed topological minors* (Definition 4.4) must upper-bound the treewidth of the underlying undirected graph. In what follows, we explain and formalize this statement. We start with the notion of algorithmic powerfulness and note what it is that makes treewidth such a successful measure. Courcelle's theorem [7] states that all MSO_2-expressible problems are linear-time decidable on graphs of bounded treewidth. To us it seems that an algorithmically useful width measure must admit algorithms with running time $\mathcal{O}(n^{f(k)})$, at least for a smaller class of all MSO_1-expressible problems on n-vertex digraphs of width at most k, where f is some computable function (that is, XP running time).

Algorithmically powerful digraph width measures do indeed exist. Candidates include the number of vertices of the input graph and the treewidth of the underlying undirected graph.

[1]Kenny-width of [15] is a different measure than Kelly-width of [23].

However, the former is not interesting at all, and in the latter case one can apply the rich theory of undirected graphs of bounded treewidth but would not get anything substantially new for digraphs. As such, we are interested in digraph width measures that are substantially different (*incomparable*, Definition 3.3) from undirected treewidth.

We also take a look at the aforementioned structural side. To motivate our discussion of directed topological minors in Section 4 and their role in our main results, we recall [39] that treewidth has a nice alternative cops-and-robber game characterization. In fact, several digraph width measures such as DAG-width [4, 32], Kelly-width [23], and DAG-depth [15] admit some variants of this useful game-theoretic characterization. While there is no formal definition of a cops-and-robber game-based width measure, all versions of the cops-and-robber game that have been considered so far share a basic property in that shrinking induced paths does not generally help the robber. What we actually argue is that a directed width measure that is "cops-and-robber game based" should be nearly monotone under taking directed topological minors (Theorem 4.7). On the other hand, we repeat that there exist algorithmically useful measures more general than undirected treewidth—those are directed clique-width [9] and bi-rank-width [25]—which are not monotone even under taking subdigraphs.

To be more specific about the words "nearly monotone", we say that a digraph width measure δ is *closed under taking* directed topological minors if there is an absolute constant c such that, for each digraph D, the δ-width of any directed topological minor of D is at most $\delta(D) + c$ (to account for possible sporadic cases). In summary, our main conceptual contribution which is further explained in Section 3 and proved in Section 6 (Theorems 6.6 and 6.7), reads:

> A digraph width measure that admits XP-time algorithms for all MSO_1-problems wrt. the width as a parameter and is closed under taking directed topological minors, cannot be substantially different from ordinary undirected treewidth.

This in turn implies that in a search for an algorithmically useful digraph width measure, one has to resign on demanding also nice structural properties, e.g. those related to cops-and-robber game characterizations.

The paper is organized in six parts, starting with some standard definitions in Section 2. Then in Section 3, we formally establish and discuss the properties an algorithmically useful digraph width measure should have. In Section 4, we introduce the notion of a directed topological minor, and discuss its properties, game-theoretic characterizations, and consider separate complexity issues. In particular, we show that it is hard to decide for a fixed small digraph whether it is a directed topological minor of a given digraph. In Section 5, we continue with the technical prerequisites of our main result. There we show that the structure of hard MSO_1-definable graph problems is just as rich for subdivisions of planar graphs with degrees 1 or 3 as for general graphs. In Section 6, we prove our main results which have already been outlined above. We also provide examples showing that the prerequisites of our results cannot be substantially weakened. We end with some concluding remarks in Section 7.

2 Definitions and notation

The graphs (both undirected and directed) that we consider in this paper are *simple*, i.e. they do not contain loops and parallel edges. Given a graph G, we let $V(G)$ denote its vertex set and $E(G)$ denote its edge set, if G is undirected. We usually denote a directed graph (*digraph* shortly) by D and its arc set by $A(D)$. Given a directed graph D, the *underlying undirected graph* $U(D)$ of D is an undirected graph on the vertex set $V(D)$ and $\{u,v\}$ is an edge of $U(D)$ if and only if either $(u,v) \in A(D)$ or $(v,u) \in A(D)$. A digraph D is an *orientation* of an undirected graph G if $U(D) = G$. If H is a sub(di)graph of G, then we denote this fact by writing $H \subseteq G$.

For a vertex pair u,v of a digraph D, a sequence $P = (u = x_0, \dots, x_r = v)$ is called *directed* (u,v)*-path* of length $r > 0$ in D if the vertices $x_0, \dots, x_r$ are pairwise distinct and $(x_i, x_{i+1}) \in A(G)$ for every $0 \le i < r$. A *directed cycle* is defined analogously with the modification that $x_0 = x_r$. A digraph D is *acyclic* (a *DAG*) if D contains no directed cycle.

Following Downey–Fellows [11], a parameterized problem $\mathcal{Q}$ is defined as a subset of $\Sigma \times \mathbb{N}_0$, where Σ is a finite alphabet and $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. The following is folklore:

Definition 2.1 (parameterized tractability, [11]) *A parameterized problem $\mathcal{Q}$ is* fixed-parameter tractable *if there is an algorithm that given $\langle x,k\rangle \in \Sigma \times \mathbb{N}_0$ decides whether $\langle x,k\rangle$ is a yes-instance of $\mathcal{Q}$ in time $f(k) \cdot p(|x|)$, where f is some computable function of the* parameter k *alone, p is a polynomial and $|x|$ is the size measure of the input. The class of fixed-parameter tractable problems is denoted by* FPT.

Furthermore, let XP *denote the class of parameterized problems $\mathcal{Q}$ that admit an algorithm running in time $\mathcal{O}(|x|^{f(k)})$ for some computable function f, i.e. with polynomial run-time for every fixed value of the parameter k. We refer to such an algorithm for $\mathcal{Q}$ as to an* XP-time *algorithm with respect to k.*

Some of the most successful kinds of parameterized problems are those of a form $\langle x,k\rangle$ where an integer k is simply the value of a suitable structural width parameter of the input x. Such as, the problem input is $\langle G,k\rangle$ where G is a graph and the parameter k is its treewidth. For the sake of completeness, we review next the definitions of the aforementioned two basic (and well-known) structural width parameters—of treewidth and clique-width.

Definition 2.2 (treewidth, [35]) *A* tree-decomposition *of a graph G is a pair (T,β), where T is a tree and $\beta : V(T) \to 2^{V(G)}$ is a mapping (of the "bags") that satisfies the following;*

- *for each edge $e = uv \in E(G)$, there is $x \in V(T)$ such that $\{u,v\} \subseteq \beta(x)$,*
- *if $x \in V(T)$, and if $y,z \in V(T)$ are two nodes in distinct components of $T - x$, then $\beta(y) \cap \beta(z) \subseteq \beta(x)$ ("interpolation"),*
- $\bigcup_{x \in V(T)} \beta(x) = V(G)$.

The width of (T,β) is the maximal value of $|\beta(x)| - 1$ over all $x \in V(T)$. The smallest width over all tree-decompositions of the graph G is the treewidth $\mathbf{tw}(G)$ *of G.*

Definition 2.3 (clique-width, [9]) *Let k be a positive integer. A pair (G,γ) is a k-*labelled graph *if G is a simple graph and $\gamma : V(G) \to \{1,2,\dots,k\}$ is a mapping. A k-*expression *is a well formed expression t built using the four operators defined below. Let $1 \le i,j \le k$. Then*

- $[i]$ *is a nullary operator which represents a graph with a single vertex labelled* i*,*
- $\eta_{i,j}$*, for* $i \neq j$*, is a unary operator which adds edges between all pairs of vertices where one is labelled* i *and the other is labelled* j*,*
- $\rho_{i \to j}$ *is a unary operator which changes the labels of all vertices labelled* i *to* j*, and*
- $\oplus$ *is a binary operator which represents disjoint union of two* k*-labelled graphs.*

Each k*-expression* t *naturally generates a* k*-labelled simple graph* $G = G[t]$*. The smallest* k *such that there exists a* k*-expression generating* G *is the* clique-width *of* G*.*

Monadic second-order (MSO) logic is a language particularly suited for describing problems on "tree-like structured" graphs. For instance, the celebrated result of Courcelle [7], and of Arnborg, Lagergren, and Seese [2], states that all MSO_2-definable graph problems have linear-time FPT algorithms when parameterized by the treewidth. The expressive power of MSO_2 is very strong as it includes many natural graph problems. We are primarily interested in another logical dialect commonly abbreviated as MSO_1, whose expressive power is noticeably weaker (in MSO_1 one cannot speak about sets of edges) than that of MSO_2. The weaker expressive power is an advantage in this paper since we are going to use it to prove negative results.

Definition 2.4 (MSO_1 logic) *The language of* MSO_1 *contains the logical expressions that are built from the following elements:*

- *variables for elements (vertices) and their sets, and the predicate* $x \in X$*,*
- *the predicate* $adj(u, v)$ *with* u *and* v *vertex variables,*
- *equality for variables, the connectives* $\wedge, \vee, \neg, \to$*, and the quantifiers* $\forall, \exists$ *.*

Example 2.5 *For undirected graphs, the property of being 3-colorable can be expressed by the* MSO_1 *formula*

$$\exists V_1, V_2, V_3 \big[\forall v \, (v \in V_1 \vee v \in V_2 \vee v \in V_3) \, \wedge \bigwedge\nolimits_{i=1,2,3} \; \forall v, w \, (v \notin V_i \vee w \notin V_i \vee \neg \, adj(v, w)) \, \big] .$$

A decision graph property $\mathcal{P}$ is MSO_1-*definable* if there exists an MSO_1 formula ϕ such that $\mathcal{P}$ holds for any graph G if and only if $G \models \phi$, i.e., ϕ is true on the model G. The language of MSO_1 can also be used for describing *digraph properties* by substituting the predicate $arc(u, v)$ for $adj(u, v)$. Courcelle, Makowski and Rotics have shown that all MSO_1-definable graph problems have FPT algorithms when parameterized by the clique-width of the input graph [8]. Since clique-width and rank-width are equivalent measures in the sense that a graph has bounded clique-width if and only if it has bounded rank-width [33], the latter result also holds when the parameter is the rank-width of the input graph.

3 Desirable digraph width measures

A *digraph width measure* is a function δ that assigns each digraph a non-negative integer. To stay reasonable, we expect that infinitely many non-isomorphic digraphs are of bounded width. Following the informal discussion in Section 1, we consider what properties a useful width measure is expected to have. Importantly, one must be able to solve a rich class of problems on digraphs of bounded width. But what does "rich" mean?

If we consider existing algorithmic metatheorems for undirected graphs, then we can see a *good balance* between the class of problems considered (say, MSO_2- or MSO_1-definable) and the existence of positive algorithmic results (linear-time decidability on graphs of bounded treewidth or clique-width, respectively). To achieve such a balance for digraphs, it seems to us that the class of MSO_1-definable problems presents the most appropriate conservative choice. For if we consider any logical language $\mathcal{L}$ over digraphs that is powerful enough to deal with sets of singletons (i.e. of monadic second order) and that can identify the adjacent pairs of vertices of the digraph, then $\mathcal{L}$ can naturally interpret the MSO_1 logic of the underlying graph (notice that this applies already to undirected MSO_1). That is why the following specification appears to be the most natural common conservative denominator in our context:

Definition 3.1 (algorithmic powerfulness) *A digraph width measure δ is* powerful *if, for every MSO_1-definable undirected property $\mathcal{P}$ (Definition 2.4), there exists an* XP*-time algorithm deciding $\mathcal{P}$ on all digraphs D (more formally, on $U(D)$) with respect to the parameter $\delta(D)$.*

The traditional measures treewidth, branchwidth, clique-width, and the more recent rank-width, are all powerful for undirected graphs [7, 8]. For directed graphs, unfortunately, exactly the opposite holds. All the width measures suggested in recent years as possible extensions of treewidth—including directed treewidth [24], D-width [38], DAG-width [32, 4], and Kelly-width [23]—are not powerful. More generally, it holds:

Proposition 3.2 *Assume a digraph width measure δ achieving only bounded values on the class of all acyclic digraphs (DAGs). If* P $\neq$ NP*, then δ is not powerful.*

Proof. Let $\mathcal{P}$ be any NP-complete MSO_1-definable property of undirected graphs, say, 3-colourability. We construct the digraph property $\mathcal{P}'$ by replacing every occurrence of the predicate $adj(x,y)$ in $\mathcal{P}$ by $\big(arc(x,y) \vee arc(y,x)\big)$. Clearly, an undirected graph has the property $\mathcal{P}$ if and only if any digraph D that is an orientation of G has the property $\mathcal{P}'$. If δ was powerful, then by Definition 3.1, the property $\mathcal{P}'$ would be decidable on all DAGs in polynomial time. Hence for any input graph G, we could decide whether $G \models \mathcal{P}$ in polynomial time by first constructing an acyclic orientation D of G, and then deciding whether $D \models \mathcal{P}'$ (in polynomial time). This would imply that P = NP. ■

Another important property one would like a digraph width measure to possess is that it should not be comparable to the treewidth of the underlying undirected graph. To be more precise, we do not want the width measure δ to upper-bound the treewidth of the underlying undirected graph. This makes sense because any measure δ which makes the undirected treewidth bounded is automatically powerful, but such δ would not help solve any more inputs than what we already can with traditional undirected measures.

Definition 3.3 *A digraph width measure δ is called* treewidth-bounding *if there exists a computable function b such that for every digraph D with $\delta(D) \leq k$, we have $\mathbf{tw}(U(D)) \leq b(k)$.*

Disparity between measures. In the rest of the section, we consider disparity between the treewidth-like digraph measures – DAG-width [32, 4] and Kelly-width [23], and the algorithmically successful ones – directed clique-width [9] and bi-rank-width [25]. None of these measures are treewidth-bounding. Since the definitions of DAG-width and Kelly-width are quite involved, we skip them here and refer to [32, 4, 23]. Instead, we note that both DAG- and Kelly-width share some common properties important for us:

- Acyclic digraphs (DAGs) have DAG-width 0 and Kelly-width 1 (cf. Proposition 3.2).
- If we replace each edge of a graph of treewidth k by a pair of opposite arcs, then the resulting digraph has DAG-width k and Kelly-width $k + 1$.
- Each of these measures can be exactly characterized by some natural directed counterpart of the cops-and-robber game for undirected treewidth [39] (cf. Proposition 4.2).

As for another type of digraph measures, aforementioned clique-width was originally defined for undirected graphs [9], but Definition 2.3 readily extends to digraphs; simply replace the operator $\eta_{i,j}$ by the operator $\alpha_{i,j}$ creating directed edges (arcs) from each vertex with label i to each vertex with label j. Bi-rank-width was introduced by Kanté [25] and is related to directed clique-width in the sense that one is bounded on a digraph class if and only if the other is. We refer to [25, 19] for its definition and properties.

Theorem 3.4 (via Courcelle, Makowsky, and Rotics [8]) *Directed clique-width, and hence bi-rank-width, are powerful digraph measures. Actually, the problem of deciding any fixed directed MSO_1-definable property of digraphs belongs to the class FPT wrt. their directed clique-width or bi-rank-width.*

For a better understanding of the situation, we stress one important but elusive fact: Bounding the *undirected* clique-width or rank-width of the underlying undirected graph does not generally help solve directed graph problems. Precisely, in analogy with Proposition 3.2 and its proof, we can claim:

Proposition 3.5 *If* $\mathrm{P} \neq \mathrm{NP}$*, then there exist MSO_1-definable digraph properties that have* ***no*** XP*-time algorithms with respect to undirected clique-width or rank-width.*

Proof. We argue that there exist directed problems $\mathcal{P}$ which are NP-complete even on tournaments (orientations of complete graphs), i.e. on digraphs of undirected clique-width 2 and rank-width 1. If such $\mathcal{P}$ had an XP-time algorithm, then this would immediately imply $\mathrm{P} = \mathrm{NP}$. As a (random) example of the problem $\mathcal{P}$ we refer [6]; the problem whether a tournament has a partition into two acyclic subtournaments which clearly has a directed MSO_1 definition. ■

Proposition 3.5 is in a *sharp contrast* to the situation with treewidth where bounding the treewidth of the underlying undirected graph allows all the algorithmic machinery to work also on digraphs. A brief informal explanation of this antagonism lies in the facts that a “bag” in a tree decomposition has bounded size and so there is only a bounded number of possible orientations of the edges in it, while a single $\eta_{i,j}$ (edge-addition) operation in a clique-width expression creates a bipartite clique of an arbitrary size which admits an unbounded number of possible orientations. As of now, there is no known non-trivial relationship between undirected measures and their directed generalizations.

A resolution? From Propositions 3.2 and 3.4, it seems that directed clique-width and bi-rank-width are better possible candidates for a good digraphs width measure. Unfortunately, clique-width and bi-rank-width do not possess the nice structural properties common to the various treewidth-like measures, such as being subgraph- or contraction-monotone. This is due to symmetric orientations of complete graphs all having clique-width two while their subdigraphs include all digraphs, even those with arbitrarily high clique-width. This seems to be a drawback and a possible reason why clique-width- and rank-width-like measures are not so widely accepted.

The natural question now is: can we take *the better of each of the two worlds*? In our search for an answer, we do not study specific digraph width measures but focus on desirable properties of all possible width measures in general. The core contribution of this paper, summarized in Theorems 6.6 and 6.7, then answers this question negatively.

- One *cannot* have a digraph width measure that is powerful (Definition 3.1), not treewidth-bounding (Definition 3.3) and closed under taking directed topological minors (Definition 4.4) at the same time.
- Consequently, "algorithmically useful" digraph width measures cannot possess "nice" structural properties at the same time.

This strong and conceptually new result holds modulo well-established computational complexity assumptions, namely, that NP $\neq$ P and NP $\not\subseteq$ P/poly. Further discussion of the necessary technical details of this result can be all found in Section 6.

4 Games and directed topological minors

The third requirement (after being powerful and not treewidth-bounding) we impose above on a desired "good" digraph width measure is for it to possess some *nice structural properties* similar to those we often see in undirected graph measures.

Having an alternative characterization in terms of a cops-and-robber game is a trait common to many widely known measures and proved to be very useful, often significantly simplifying some of the proofs. Most of the directed graph measures (in fact all of the major ones), which were proposed as a directed counterpart to treewidth, have a characterization in terms of some variant of this cops and robber game, and indeed they were explicitly designed to have this characterization.

The treewidth game. The original cops-and-robber game was introduced in [39] as an alternative characterization of treewidth. The robber stands on a vertex of the graph, and can at any time run at great speed to any other vertex along a path of the graph. He is not permitted to run through a cop, however. There are k cops, each of whom at any time either stands on a vertex or is in a helicopter. The goal of the player controlling the cops is to land a cop via a helicopter onto a vertex currently occupied by the robber, and the robber's objective is to elude capture. (The point of the helicopter is that cops are not constrained to move along the edges of the graph.) The robber can see the helicopter landing and may run to a new vertex before it actually lands.

More formally, the game is played on a graph G by two players: the cop player, and the robber player. We denote (X, r) the position of the game where cops are placed at $X \subseteq V(G)$ and the robber is in $r \in V(G)$. The game is played according to the following rules: At the beginning the robber player chooses a vertex $r_0 \in V(G)$, giving us an initial game position $(\emptyset, r_0)$. Given a position (X, r), the cop player chooses a set $X' \subseteq [V]^{\leq k}$, and then the robber player chooses a vertex $r' \in V(G) \setminus X'$ such that both r and r' lie in the same connected component of the graph $G \setminus (X \cap X')$, giving us the next position (X', r'). A play is a maximal sequence γ of positions formed from an initial game position according to the rule above.

The play is winning for the cop player if it is finite—i.e., for the final position (X, r) of the play γ, it is true that there is $X' \in [V]^{\leq k}$ such that no vertex of the graph $V(G) \setminus X'$ is in the same connected component of the graph $G \setminus \{X \cap X'\}$ as r (this immediately implies $r \in X$). On the other hand the robber player wins if the play is infinite. The *strategy* for the cop player is a function $\sigma : (X, r; \gamma) \mapsto X'$, which, for a position (X, r) of the game, gives us the next position X' of the cops. Notice that in general a strategy σ depends also on the *history* γ which is a subsequence of positions from the initial one to (X, r) in the current play. Without loss of generality we can restrict ourselves to strategies where $|X \triangle X'| = 1$, i.e. we either place or remove exactly one cop. The following theorem (often called the *treewidth duality theorem*) relates cops-and-robber games and treewidth:

Theorem 4.1 ([39]) *Graph G has treewidth k iff the minimum number of cops required to win the cops-and-robber game is $k + 1$.*

For directed graphs the game is naturally modified by additionally requiring that the robber follows the direction of arcs: for any two consecutive robber positions r and r', as above, there is a *directed path* from r to r' in $G \setminus \{X \cap X'\}$. Several variants of the directed cops-and-robber game were proposed for different digraph width measures. The most important variants restrict the movement of cops and/or the robber. The cops may be required to play *monotonely*, meaning that they can never revisit a vertex they have already left. The robber can be *invisible* to the cops, requiring the cops to search the whole graph. Finally the robber can be *lazy*, i.e. he can move only when a cop is just about to land on the vertex currently occupied by the robber (as opposite to the *eager* robber in the original game). Formally, a game is;

- *monotone* if, whenever positions $(X_i, r_i), \ldots (X_j, r_j), \ldots (X_k, r_k)$ occur in this order (not necessarily consecutive) in a play, then $X_j \supseteq X_i \cap X_k$;
- *robber-invisible* if the cop strategy, for any two valid play histories γ and γ' which differ (from each other) only in robber positions, fulfils $\sigma(X, r; \gamma) = \sigma(X, r'; \gamma')$.
- *robber-lazy* if a game position (X, r) is always succeeded by (X', r) (for some X') unless $r \in X'$.

Proposition 4.2 ([32, 4, 23, 3]) *The most common digraph width measures are characterized by the following cops-and-robber game variants:*

measure	*cops*	*robber*
DAG-width	*monotone*	*visible*
Kelly-width	*monotone*	*invisible,lazy*
directed path-width	*monotone*	*invisible*

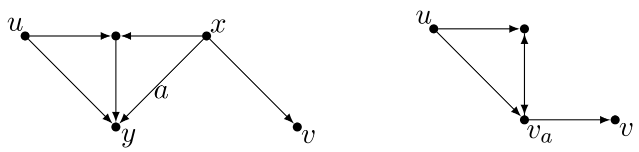

Figure 1: Arc contraction: digraphs D (left) and D/a.

Introducing directed minors. In the realm of undirected graphs, characterizability by a cops and robber game is closely related to monotonicity under taking minors. Recall that a graph H is a *minor* of a graph G if it can be obtained by a sequence of applications of three operations: vertex deletion, edge deletion and edge contraction. (See e.g. [10].) A measure is *monotone* under taking minors if the measure of a minor is never larger than the measure of the graph itself. Note that treewidth is monotone under taking minors. The relationship between cops-and-robber games and taking a minor of a graph should now be obvious: Taking a subgraph can never improve robber's chances of evading k cops and, since the robber can move infinitely fast and cops use helicopters, neither can edge contraction.

It is therefore only natural to expect that a "good" digraph measure, characterizable by a directed cops-and-robber game, should also be (at least nearly) monotone under some notion of a directed minor. However, there is currently no widely agreed definition of a directed minor in general. One published (but perhaps too restrictive on subdivisions as we will see later) notion is the *butterfly minor* [24]. Similarly, a traditional and more relaxed notion of *digraph immersion* is still too restrictive on subdivisions for our purpose. We now refer to Proposition 6.9 for more details.

To deal with directed minors, we first need a formal notion of an *arc contraction* for digraphs:

Definition 4.3 *Let D be a digraph and $a = (x, y) \in A(D)$ be an arc. The digraph obtained by* contracting arc *a, denoted by D/a, is the digraph on the vertex $(V(D) \setminus \{x, y\}) \cup \{v_a\}$ where v_a is a new vertex, and the arc set A' such that $(u, v) \in A'$ iff one of the following holds*

$$\begin{aligned} &(u, v) \in A(D \setminus \{x, y\}), \text{ or} \\ &v = v_a, \text{ and } (u, x) \in A \text{ or } (u, y) \in A, \text{ or} \\ &u = v_a, \text{ and } (x, v) \in A \text{ or } (y, v) \in A. \end{aligned}$$

See Fig. 1 for an example of a contraction. Note that contraction always produces simple digraphs (that is, no arcs of the form (x, x)). The result of a contraction does not depend on the orientation of the contracted arc, and we treat contraction of a pair of opposite arcs (x, y) and (y, x) as a contraction of a single bidirectional arc. For simplicity, we also sometimes identify the newly created vertex v_a with one of the former vertices x, y (depending on the context).

An important decision point when defining a minor is; which arcs do we *allow to contract*? In the case of undirected graph minors, any edge can be contracted. However, the situation is not so obvious in the case of digraphs. Look again at Figure 1. If we contract the arc a, we actually introduce a directed path from u to v which was not present before, whereas in undirected graphs no new (undirected) path is ever created by the edge contraction. On the other hand, simply never introducing a new directed path (that is the butterfly minor of [24]) is

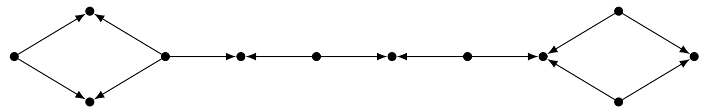

Figure 2: Any arc contraction in this digraph introduces a new directed path.

not a good strategy either – since one can easily construct, see Figure 2, digraphs in which no arc can be contracted without introducing a new directed path. Yet such digraphs can be "very simple" with respect to usual cops-and-robber games, and arc contractions do not significantly help the robber in the depicted situation of Figure 2.

In order to deal with the mentioned issue of contractibility of arcs, and to remain as general as possible at the same time, we actually consider a specific kind of minors, i.e. *directed topological minors* where we allow only those arc contractions that do not introduce any new directed path between vertices of degree at least three (cf. Figure 2 again). This is again inspired by a related notion in the realm of undirected graphs. A *topological minor* for undirected graphs is defined similarly to the usual minor, but the edge contraction can only be applied to an edge $e = \{u, v\}$ such that u or v has exactly two neighbours. In other words, graph H is a topological minor of G iff a subdivision of H is isomorphic to a subgraph of G.

Note that topological minors for digraphs have been defined before. In his thesis Hunter [22] considers a definition of directed topological minor where an arc is contractible iff at least one end vertex has both out- and in-degree one. This approach coincides with the notion of directed subdivisions used, e.g., by Mader [30]. However, this definition deeply suffers from the same problems (i.e. no arc is contractible in Fig. 2) as the butterfly minors. We will argue that our definition is more relevant in the context of digraph widths and their game characterizations, and justify our choice by the subsequent claims (see Lemma 4.6 and Theorem 4.7).

For simplicity, we sometimes write $u \to_D v$ for $(u, v) \in A(D)$, and $u \to_D^+ v$ for the transitive closure of $\to_D$ (a directed path from u to v). The index D may be omitted if it is clear from the context. We also write $u \to_D^* v$ if either $u \to_D^+ v$ or $u = v$.

Definition 4.4 *Let $V_3(D) \subseteq V(D)$ denote the subset of vertices having at least three neighbours in D. An arc $a = (u, v) \in A(D)$ is* 2-contractible *in a digraph D if*

- *u or v has exactly two neighbours (particularly, $\{u, v\} \not\subseteq V_3(D)$), and*
- *$(v, u) \in A(D)$ or there is no pair of vertices $x, y \in V_3(D)$, possibly $x = y$, such that $x \to_{(D\setminus a)}^* v$ and $u \to_{(D\setminus a)}^* y$.*

A digraph H is a directed topological minor *of D if there exists a sequence of digraphs $D_0, \ldots, D_r$ such that $D_0 \subseteq D$ and $D_r \cong H$, and for all $0 \le i \le r - 1$, one can obtain D_{i+1} from D_i by contracting a 2-contractible arc.*

The essential idea in the above definition is that an arc is 2-contractible if its contraction does not result in the creation of a new directed path between vertices of degree at least 3. Robustness of the definition is justified as follows.

Proposition 4.5 *Given a digraph D, let D' be obtained from D by a sequence of vertex deletions, arc deletions and contractions of 2-contractible arcs (in any order). Then D' is a directed topological minor of D.*

Proof. Let ϱ be the sequence of vertex deletions, arc deletions and contractions of 2-contractible arcs used to obtain D'. We process all three types of operation in a fixed order. We start by vertex deletions. Let W be the smallest subset of $V(D)$ created as follows: for each deletion of a vertex v in ϱ such that $v \in V(D)$ the set W contains v. Otherwise v was created by a contraction and W includes all vertices of $V(D)$ which were contracted to form v'. We now put $D_1 = D \setminus W$. Next we perform the arc deletions in ϱ (note that some of the arcs may have already been removed by the previous step). Let us call the resulting graph D_2. Finally we contract all remaining 2-contractible arcs which were contracted in ϱ, in the same order, and call the resulting graph D_3.

Obviously D_2 is a subgraph of D, and D_3 was obtained from D_2 by a series of contractions of 2-contractible arcs. We claim that D_3 is isomorphic to D'. It is easy to check that all vertices and arcs deleted in ϱ are not present in D_3, and no extra vertices and arcs were removed. The rest follows from the fact that by Definition 4.4 a 2-contractible arc cannot become non-2-contractible by vertex and arc deletions. ■

As mentioned above, the driving force in our definition of directed topological minors is to always allow contraction of "long subdivisions", as e.g. in Figure 2. We shall make this formally precise now. Let D be a digraph and $P = (x_0, \dots, x_k)$ a sequence of vertices of D. Then P is a *2-path* (of length k) in D if P is a path in the underlying graph $U(D)$ and all its internal vertices x_i for $0 < i < k$ have exactly two neighbours (degree 2) in $U(D)$. Obviously not every 2-path is a directed path. The following lemma explains the close relationship between 2-paths and directed topological minors.

Lemma 4.6 *Let D be a digraph and $S = (x_0, \dots, x_k)$ a 2-path of length $k > 2$ in D. Then there exists a sequence of 2-contractions of arcs of S in D turning S into a 2-path of length two (or even of length one if S is a directed path).*

Proof. If S is a directed path in D, then any arc of S is 2-contractible. Otherwise S is not a directed path in D and therefore there exists an index $0 < i < k$, such that $(x_i, x_{i-1}), (x_i, x_{i+1}) \in A(D)$ or $(x_{i-1}, x_i), (x_{i+1}, x_i) \in A(D)$. Then all other arcs of S are 2-contractible, as their contraction cannot produce a new directed path between any two vertices in D with more than two neighbours. ■

Monotonicity of the measures. The natural question is whether the known digraph measures are monotone under taking directed topological minors. This is not exactly true for certain sporadic examples as shown, e.g., in Figure 3, but we prove that the measures (i.e., those mentioned by Proposition 4.2) cannot grow by more than two after taking any such minor:

Theorem 4.7 *Let D be a digraph such that in the DAG-width (Kelly-width, directed path-width) game, $k \geq 1$ cops are enough to catch the robber. Let H be a directed topological minor of D. Then at most $k + 2$ cops are needed to catch the robber on H in that same game.*

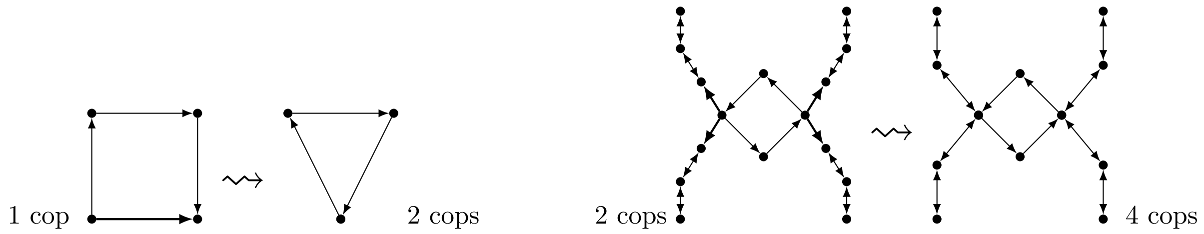

Figure 3: Two examples when the required number of cops grows by 1 and by 2, respectively, after taking a directed topological minor. The right-hand side applies only to the path-width game, and is a worst-case example according to Theorem 4.7.

We actually believe that a bound of $k+1$ cops is enough in Theorem 4.7 for the DAG-width and Kelly-width games (and perhaps even a bound of k cops if $k \geq 3$). The precise proof is, however, surprisingly nontrivial and so we stick here with a much simpler argument for the bound $k+2$, which simultaneously covers all the three games.

To state our claims, we need some more notation. Throughout the proof, for a game position (X, v) we denote by $Reach^D(v, X) = \{x \in V(D) \setminus X \mid v \to^*_{D \setminus X} x\}$ the set of vertices reachable by the robber placed in v with cops positioned in X. The vertices not in $Reach^D(v, X)$ are shortly called *clear*. Let γ denote a play $\big((X_i, r_i) : i = 0, 1, 2, \dots\big)$ of our respective game, which must be monotone by Proposition 4.2 (hence once cleared vertex will stay clear forever in a cop-win play). For $v \in V(D)$, we define $index_\gamma(v)$ as the least index k such that $v \in X_k$ in the play γ. By $Margin^D_\gamma(v)$ we denote the set of those vertices $x \in V(D) \setminus V_3(D)$ for which v is the last (wrt. the play γ) non-clear vertex of $V_3(D)$ such that x is reachable from v along a 2-path (including $x = v$). Formally, for $v \in V_3(D)$, and setting $k = index_\gamma(v)$ and $Reach^D_2(v) = Reach^D\big(v, V_3(D) \setminus \{v\}\big)$, we have got $Margin^D_\gamma(v) = \big\{x \in Reach^D_2(v) \mid \forall w \in V_3(D) : x \in Reach^D_2(w) \to w \notin Reach^D(r_k, X_k)\big\}$. As a special case, $Margin^D(\emptyset)$ denotes the vertices of D not reachable from any vertex of $V_3(D)$.

Note that $Margin^D(\emptyset)$ contains only vertices of degree ≤ 2 in $U(D)$. $Margin^D_\gamma(v)$ consists of the vertices of some 2-paths starting at v (up to the point where orientation of arcs admits reachability from v). For such a 2-path $P = (x_0 = v, x_1, \dots, x_k)$, we say that the play γ clears P in a *leapfrogging manner from* v if the following hold for some index ℓ of γ: It is $index_\gamma(v) < \ell$ and $v, x_1 \in X_\ell$. Then, for $i = 1, \dots, k-1$, it is $X_{\ell+2i-1} = X_{\ell+2i-2} \cup \{x_{i+1}\}$ and $X_{\ell+2i} = X_{\ell+2i-1} \setminus \{x_i\}$, unless possibly $x_{i+1} \in X_{\ell+2i-2}$ in which case the former one step is skipped. Finally, $X_{\ell+2k-1} = X_{\ell+2k-2} \setminus \{x_k\}$.

Now, a monotone play $\gamma = \big((X_i, r_i) : i = 0, 1, 2, \dots\big)$ is called *2-special* if (1) whole $Margin^D(\emptyset)$ is cleared before any other vertex of D; and (2) for every $v \in V_3(D)$, each of the 2-paths forming $Margin^D_\gamma(v)$ is cleared by γ in a leapfrogging manner from v right after a cop lands on v (precisely, before a cop subsequently lands on another vertex of $V_3(D)$). The order of clearing between these 2-paths is irrelevant. A monotone cop strategy is *2-special* if every play according to this strategy is 2-special.

Lemma 4.8 *Let D be a digraph such that in the DAG-width (Kelly-width, directed path-width) game, $k \geq 1$ cops are enough to catch the robber. Then $k+2$ cops have (in the respective game)* *a 2-special monotone winning strategy against the robber in D.*

Proof. We do the same proof for all the three games, and hence we neither look at robber's position, nor assume him to be lazy. We also do not explicitly invoke play history in subsequent modifications. The proof closely follows the definition of 2-special. Let c_1, c_2 be the two additional cops. Namely, for a monotone winning cop strategy σ (not using c_1, c_2), we first add to the beginning of σ an initial episode clearing whole $Margin^D(\emptyset)$ with c_1, c_2 and one of the original cops. This is always possible (even with an invisible robber) since the subgraph of $U(D)$ induced by $Margin^D(\emptyset)$ consists of paths and cycles, and it is not reachable from the rest of D. Then we continue with original σ, while for each $v \in V_3(D)$ we iteratively add "leapfrog episodes" (in the aforementioned way) which clear whole $Margin^D_\gamma(v)$ using c_1, c_2 in a leapfrogging manner from v, right after an original cop lands on v (irrespectively of the robber moves).

Note that after each added episode, the cops c_1, c_2 are again free. Subsequently to an added episode as above, we also skip from σ all (potential) future cop moves to the vertices of $Margin^D(\emptyset)$ or $Margin^D_\gamma(v) \setminus \{v\}$, respectively. Consequently, this makes our new strategy monotone if σ is such.

So, the new strategy σ' is 2-special monotone by its definition, and it remains to prove that each play by σ' is cop-win. The whole game would clearly be cop-win (but non-monotone) if we only added episodes. Then, firstly, the initial episode clears $Margin^D(\emptyset)$, and the robber can never return there by its definition which justifies skipping of subsequent cop moves to $Margin^D(\emptyset)$ in σ'. Secondly, consider a play γ by σ' and a vertex $x \in Margin^D_\gamma(v) \setminus \{v\}$ where $v \in V_3(D)$. Then x is cleared during the leapfrog episode at v, and x cannot be reached from $V_3(D)$ otherwise then from v or possibly from another vertex cleared in γ before v, by the definition of $Margin^D$. Hence it follows from assumed monotonicity of σ that x will not be reachable again by the robber in γ. Hence the robber is finally caught in γ. ■

Lemma 4.9 *Let D be a digraph such that in the DAG-width (Kelly-width, directed path-width) game, $k \geq 3$ cops have a 2-special monotone winning strategy against the robber in D. If $D' = D/(s,t)$ is obtained by contracting a 2-contractible arc (s,t) in D, then k cops have a 2-special monotone winning strategy against the robber in D', too.*

Proof. We refer to the proof of Lemma 4.8. If $s, t \in Margin^D(\emptyset)$, then the claim is clear since 3 cops are always enough for the initial episode of a 2-special strategy. If $s, t \in Margin^D_\gamma(v)$ for some $v \in V_3(D)$ (and possibly $s = v$ or $t = v$), then the case is similarly easy for each such play γ; the leapfrog episode from v simply gets shorter. Now assume $s \in Margin^D(\emptyset)$ while $t \in Margin^D_\gamma(v)$; in D', the set $Margin^{D'}_\gamma(v)$ increases by the vertices of $Margin^D(\emptyset)$ reachable from s, but that again is no problem for the leapfrog episode from v.

It remains to consider the interesting last case of $s \in Margin^D_\gamma(v)$ and $t \in Margin^D_\gamma(w)$ for some $v \neq w \in V_3(D)$ in a play γ. Then $(t,s) \notin E(D)$. Since $t \in Reach^D_2(v)$, by definition, w is cleared in γ later than v. Let q be the vertex of D' resulting from contraction of (s,t). By Definition 4.4, $q \not\to^*_{(D' \setminus w)} v$. Hence, for $S = Reach^{D'}_2(q) \setminus Margin^D_\gamma(w)$, the new leapfrog episode from v in D' can safely skip the vertices of S; it is now $Margin^{D'}_\gamma(v) = Margin^D_\gamma(v) \setminus S$ and $Margin^{D'}_\gamma(w) = Margin^D_\gamma(w) \cup S$. This already gives the desired 2-special monotone winning strategy in D'. ■

Proof. (Theorem 4.7) If k cops can catch the robber in a digraph D using strategy σ, they can do so in any subdigraph $H \subseteq D$ using the strategy $\sigma|H$, i.e. the strategy σ restricted to H. We therefore focus on the case of arc contractions (cf. Proposition 4.5). We first build a 2-special monotone winning strategy for $k + 2 \geq 3$ cops as in Lemma 4.8. Then we iteratively apply Lemma 4.9 onto this case, and so finish the proof. ∎

For the last claim, recall that "closed" means there is a constant c such that, for each digraph D, the δ-width of any directed topological minor of D is at most $\delta(D) + c$.

Corollary 4.10 *DAG-width, Kelly-width and directed path-width are closed under taking directed topological minors.*

The résumé of this section is a strong suggestion that the property of being closed under taking directed topological minors is indeed a *natural requirement* for any "cops-and-robber based" digraph width measure. See Section 6 for the consequences.

Complexity of testing directed topological minors. In the rest of this section, we conclude our discussion of directed minors with considering the complexity of deciding whether a given digraph is a directed topological minor of another digraph. We show that this problem is hard by giving a reduction from the 2-Linkage problem, which is the following problem. Let D be a digraph and let s_1, s_2, t_1, t_2 be pairwise different vertices of D. A *2-linkage* for $\{(s_1, t_1), (s_2, t_2)\}$ is a pair (P_1, P_2) of vertex-disjoint directed paths where P_i is a (s_i, t_i)-path in D for $i \in \{1, 2\}$.

Proposition 4.11 ([13]) *The* 2-Linkage *problem, given a digraph D and $\{(s_1, t_1), (s_2, t_2)\}$ where s_1, s_2, t_1, t_2 are pairwise different vertices of D, to decide whether D has a 2-linkage for $\{(s_1, t_1), (s_2, t_2)\}$, is* NP*-complete.*

Theorem 4.12 *There exists a digraph H such that the problem, given a digraph D, to decide whether H is directed topological minor of D, is* NP*-complete.*

Proof. Since the problem is clearly solvable in non-deterministic polynomial time, it remains to show that the problem is NP-hard. We reduce from the 2-Linkage problem. The proof goes in two steps. We first show that the 2-Linkage problem remains hard on digraphs where every vertex has at most three neighbours.

Let $D, \{(s_1, t_1), (s_2, t_2)\}$ be an instance of the 2-Linkage problem. Let D' be a digraph such that $V(D') = V(D) \cup \{s'_1, s'_2, t'_1, t'_2\}$, where s'_1, s'_2, t'_1, t'_2 are new vertices, and $A(D') = A(D) \cup \{(s'_1, s_1), (s'_2, s_2), (t_1, t'_1), (t_2, t'_2)\}$. Obviously D has a 2-linkage for $\{(s_1, t_1), (s_2, t_2)\}$ if and only if D' has a 2-linkage for $\{(s'_1, t'_1), (s'_2, t'_2)\}$. Next, we modify large-degree vertices. We obtain digraph D'' from D' by iteratively executing, for every vertex x with $\mathrm{d}_{D'}(x) \geq 4$, the following sequence of operations: delete x, introduce $\mathrm{d}_{D'}(x)$ new vertices $x_1, \ldots, x_{\mathrm{d}_{D'}(x)}$, add the arcs (x_i, x_{i+1}) for $1 \leq i < \mathrm{d}_{D'}(x)$, assign the in-neighbours of x as in-neighbours to the vertices $x_1, \ldots, x_j$ where $j = |N^{\mathrm{in}}_{D'}(x)|$, and assign the out-neighbours of x as out-neighbours to the remaining vertices. Observe that x is replaced by a digraph with vertices of degree at most 3, and reachability is preserved. In particular, it holds that D' has a 2-linkage for $\{(s'_1, t'_1), (s'_2, t'_2)\}$

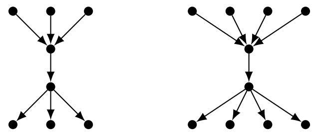

Figure 4: Digraph H from the proof of Theorem 4.12.

if and only if D'' has a 2-linkage for $\{(s_1', t_1'), (s_2', t_2')\}$. Since the construction of D'' requires only polynomial time, this shows that 2-Linkage is NP-complete on digraphs of maximum degree at most 3 due to Proposition 4.11.

For the second step of the proof, we continue with D'' and $\{(s_1', t_1'), (s_2', t_2')\}$. We introduce three new vertices and make them in-neighbours of s_1'. Similarly, we make three new vertices out-neighbours of t_1', and four new vertices become in-neighbours of s_2' and another four new vertices become out-neighbours of t_2'. Let D^* be the resulting digraph. It holds that s_1', s_2', t_1', t_2' are the only vertices of D^* of degree more than 3. We want to show that D^* has a 2-linkage for $\{(s_1', t_1'), (s_2', t_2')\}$ if and only if the digraph H depicted in Figure 4 is a directed topological minor of D^*. For the first implication, let D^* have a 2-linkage for $\{(s_1', t_1'), (s_2', t_2')\}$. Thus, there are a directed (s_1', t_1')-path P_1 and a directed (s_2', t_2')-path P_2 in D^* that are vertex-disjoint. Hence, D^* has a subgraph F that contains the vertices of P_1 and P_2 and the fourteen new vertices for D^*, that are connected only to s_1', s_2', t_1', t_2'. By Lemma 4.6 we can contract both P_1 and P_2 to a single arc each. Therefore H is a directed topological minor of D^*. For the converse, let H be a directed topological minor of D^*. Due to the definition, there is a subgraph H' of D^* such that H is isomorphic to a digraph that is obtained from H' by only contracting 2-contractible arcs. Since all vertices of H have degree different from 2, all arcs of H are either arcs of H' or result of a contraction. In particular, the two vertices of degree 4 of H correspond to s_1' and t_1' of D^*, and the two vertices of degree 5 of H correspond to s_2' and t_2' of D^*. That is because these four vertices are the only vertices of D^* of degree larger than 3. By definition of 2-contractible arc, a path between two vertices of a degree greater than two can be contracted to an arc only if it directed path between these two vertices. Hence, since H is obtained from H' by only contracting contractible arcs, there are a directed (s_1', t_1')-path P_1 and a directed (s_2', t_2')-path P_2 in H' that are vertex-disjoint. Since H' is a subgraph of D^*, P_1 and P_2 are directed paths in D^*, and therefore, D^* has a 2-linkage for $\{(s_1', t_1'), (s_2', t_2')\}$. This completes the proof of the theorem. ■

The complexity result of Theorem 4.12 shows that it is already difficult for relatively simple digraphs to decide whether they are directed topological minor of some given digraph. I.e. the "directed topological minor" decision problem is not fixed-parameter tractable with the number of vertices of the minor as a parameter.

A natural question is to ask how the "directed topological minor" problem behaves on restricted input digraphs. It has been shown that the generalization of the 2-Linkage problem to arbitrary numbers of given pairs, that is called the Linkage problem (given a digraph D and pairs of pairwise different vertices, decide whether the pairs can be joined by vertex-disjoint directed paths) is NP-complete on acyclic digraphs [40]. It is not difficult to see that the proof of Theorem 4.12 can be extended to prove the next result. In particular, if the input digraph D

is acyclic, all construction steps yield again acyclic digraphs.

Theorem 4.13 *The problem, given two acyclic digraphs D and H, to decide whether H is directed topological minor of D, is* NP*-complete.*

5 Hard MSO_1 problems for $\{1,3\}$-regular planar graphs

In the proof of the core result of the paper, Theorem 6.6, we make use of the fact that there exist MSO_1-definable problems that are NP-hard even on a very restricted graph class, the class of $\{1,3\}$-regular planar graphs. We call an undirected graph $\{1,3\}$*-regular planar* if it is planar and all its vertices have degree either one or three.

Many problems that are hard on general graphs admit efficient algorithms on planar graphs and particularly on $\{1,3\}$-regular planar graphs; some variants of the Graph Colouring problem being good examples. Recall that for $k \geq 1$, a graph G is k*-colorable* if each vertex of G can be assigned one of k colours such that adjacent vertices receive different colours. It is well-known that every planar graph is 4-colorable and that it is NP-hard to decide whether a planar graph is 3-colorable [20], while Brooks' Theorem says that all $\{1,3\}$-regular graphs are 3-colorable with a single exception of K_4.

On the other hand, some other traditional hard problems such as Maximum Independent Set remain NP-hard on $\{1,3\}$-regular planar graphs. However, the decision version of this problem is not MSO_1-definable; it belongs to a wider class of problems called EMSO_1—see [8].

To give an example of a natural graph problem that is hard on $\{1,3\}$-regular planar graphs, we observe the following: Deciding whether the square of a $\{1,3\}$-regular planar graph is 4-colorable is NP-hard [12]. Squares of planar graphs may not be planar, but the 4-colorability problem of the square of a graph admits a reformulation as a colouring problem for the input graph itself: given a graph G, decide whether G admits a 4-colouring such that the neighbours of each vertex receive pairwise different colours. We call this NP-hard problem Fair 4-Colouring of G, and it is easy to formulate it in the MSO_1 language (cf. Definition 2.4).

Moreover, a similar topic has already been addressed by Makowsky and Mariño [31]; showing that $3SAT$ reduces to a certain MSO_1-definable property on planar graphs of degree ≤ 3. Their reduction, however, uses also vertices of degree 2 while we need to avoid them here.

In this section, we thus show a stronger result covering a wide range of problems. We prove that every MSO_1 property φ efficiently translates into another MSO_1 property ψ such that a graph H satisfies φ if and only if a certain $\{1,3\}$-regular planar graph G satisfies ψ. The graph H can be an arbitrary graph, and the graph G is efficiently computable from H. The constructed property ψ is, importantly, invariant under subdivisions of G. Consequently, every MSO_1-definable problem that is hard on general graphs translates into a hard problem for $\{1,3\}$-regular planar graphs and their arbitrary subdivisions. The formal statement is given by Theorem 5.1 which we prove in this section.

Theorem 5.1 *For every pair (H,φ) where H is an undirected graph and φ is an MSO_1-definable property, there exists a pair (G,ψ) that satisfies the following conditions:*

a) G is a $\{1,3\}$-regular planar graph and ψ is an MSO_1-definable property,

$$\begin{array}{ccc} \phi \in \text{MSO over } \mathscr{K} & \xrightarrow{\;I\;} & \phi^I \in \text{MSO over } \mathscr{L} \\ H \in \mathscr{K} & & G \in \mathscr{L} \\[1em] G^I \cong H & \xleftarrow{\;I\;} & G \\ (\text{s.t. } G^I \models \phi) & & (\text{s.t. } G \models \phi^I) \end{array}$$

Figure 5: The concept of an interpretation I of $\mathrm{Th}_{\mathrm{MSO}}(\mathscr{K})$ into $\mathrm{Th}_{\mathrm{MSO}}(\mathscr{L})$.

b) $H \models \varphi$ if and only if $G \models \psi$,
c) for every subdivision G_1 of G: $G_1 \models \psi$ if and only if $G \models \psi$,
d) and ψ depends only on φ and $|\psi| = \mathcal{O}(|\varphi|)$.

Furthermore, the pair (G, ψ) can be computed from (H, φ) by a polynomial-time algorithm.

Notice, particularly, the important technical condition c) of the theorem, which states that the artificially constructed property ψ is invariant under any subdivisions of the graph G.

Our main tool for the proof of this theorem is the classical interpretability of logic theories [34]. To describe its simplified setting, assume that two classes of *relational structures* $\mathscr{K}$ and $\mathscr{L}$ are given. The basic idea of an *interpretation* I of the theory $\mathrm{Th}_{\mathrm{MSO}}(\mathscr{K})$ into $\mathrm{Th}_{\mathrm{MSO}}(\mathscr{L})$ is to transform MSO formulas ϕ over $\mathscr{K}$ into MSO formulas ϕ^I over $\mathscr{L}$ in such a way that "truth is preserved":

- First one chooses a formula $\alpha(x)$ intended to define in each structure $G \in \mathscr{L}$ a set of individuals (new domain) $G[\alpha] := \{a : a \in dom(G) \text{ and } G \models \alpha(a)\}$, where $dom(G)$ denotes the set of individuals (domain) of G.

- Then one chooses for each s-ary relational symbol R from $\mathscr{K}$ a formula $\beta^R(x_1, \ldots, x_s)$, with the intended meaning to define a corresponding relation $G[\beta^R] := \{(a_1, \ldots, a_s) : a_1, \ldots, a_s \in dom(G) \text{ and } G \models \beta^R(a_1, \ldots, a_s)\}$. With the help of these formulas one can define for each structure $G \in \mathscr{L}$ the relational structure $G^I := \big(G[\alpha], G[\beta^R], \ldots\big)$ intended to correspond with structures in $\mathscr{K}$.

- Finally, there is a natural way to translate each formula ϕ (over $\mathscr{K}$) into a formula ϕ^I (over $\mathscr{L}$), by induction on the structure of formulas. The atomic formulas are substituted by corresponding chosen formulas (such as β^R) with the corresponding variables. Then one proceeds via induction simply as follows:

$$\begin{array}{rclcrcl} (\neg\phi)^I & \mapsto & \neg(\phi^I) & , & (\phi_1 \wedge \phi_2)^I & \mapsto & (\phi_1)^I \wedge (\phi_2)^I, \\ \big(\exists x\, \phi(x)\big)^I & \mapsto & \exists y \,\big(\alpha(y) \wedge \phi^I(y)\big) & , & \big(\exists X\, \phi(X)\big)^I & \mapsto & \exists Y\, \phi^I(Y). \end{array}$$

The concept is briefly illustrated in Figure 5.

Definition 5.2 *Let $\mathscr{K}$ and $\mathscr{L}$ be classes of relational structures. Theory* $\mathrm{Th}_{\mathrm{MSO}}(\mathscr{K})$ *is* interpretable *in theory* $\mathrm{Th}_{\mathrm{MSO}}(\mathscr{L})$ *if there exists an interpretation I as above, such that the following two conditions are satisfied:*

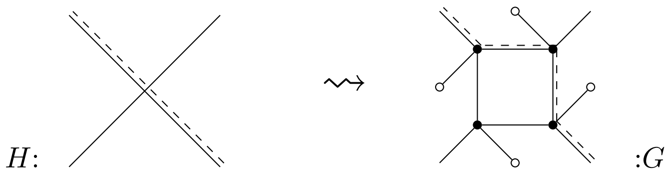

Figure 6: The *crossing-gadget* modelling, in planar G, an edge crossing of $G^I \cong H$.

i) For every structure $H \in \mathscr{K}$*, there is* $G \in \mathscr{L}$ *such that* $G^I \cong H$*, and*
ii) for every $G \in \mathscr{L}$*, the structure* G^I *is isomorphic to some structure of* $\mathscr{K}$*.*

Furthermore, $\mathrm{Th}_{\mathrm{MSO}}(\mathscr{K})$ *is* efficiently interpretable *in* $\mathrm{Th}_{\mathrm{MSO}}(\mathscr{L})$ *if the translation of each* ϕ *into* ϕ^I *is computable in polynomial time, and also the structure* $G \in \mathscr{L}$ *such that* $G^I \cong H$ *can be computed from any* $H \in \mathscr{K}$ *in polynomial time.*

Since interpretability is clearly a transitive concept, we prove Theorem 5.1 in the following sequence of three relatively easy claims.

Lemma 5.3 *The* MSO_1 *theory of all simple (undirected) graphs is efficiently interpretable in the* MSO_1 *theory of simple planar graphs.*

We believe that Lemma 5.3 is a known statement, but since we have not found an explicit reference, we present an illustrating proof for it. (Alternatively, one can adjust the reduction of [31] to achieve a similar goal.)

Proof. For start, we define the formula $deg_1(x) \equiv \forall y,z\,[(adj(x,y) \wedge adj(x,z)) \rightarrow y = z] \wedge \exists y\; adj(x,y)$ expressing that x is of degree 1 (in the model G). Furthermore, the formula

$$con(u,v,X) \equiv \forall Y \subseteq X\; \exists y,z\, \big[(y \in Y \vee y = u) \wedge (z \notin Y \vee z = v) \wedge adj(y,z)\big] \tag{1}$$

expresses the property that the subgraph induced by $X \cup \{u,v\}$ connects u to v, and

$$mcon(u,v,X) \equiv\; con(u,v,X) \wedge \forall Y(Y \subsetneq X \rightarrow \neg\, con(u,v,Y)) \tag{2}$$

says that X is a minimal connection (a path) between u, v.

Let $\mathscr{G}$ be the class of all simple graphs, and $\mathscr{P}$ the class of planar simple graphs. Following Definition 5.2, the formula α_1 defining the domain (vertex set in this case) of a graph $G^I \cong H \in \mathscr{G}$ inside $G \in \mathscr{P}$ is given as

$$\alpha_1(v) \;\equiv\; \neg\, deg_1(v) \wedge \forall x(adj(x,v) \rightarrow \neg\, deg_1(x))\,. \tag{3}$$

The underlying idea is that we would like to use some vertices of a planar graph $G \in \mathscr{P}$ to model edge "crossings" of $H \cong G^I$, see in Figure 6. Hence we "mark" each of such supplementary vertices with a new neighbour of degree 1.

It remains to interpret the adjacency relation $\beta_1^{adj}(u,v)$ for G^I. Notice that the crossing gadget in Figure 6 uses a unique "marked" 4-cycle to model each crossing, and we can identify all such 4-cycles in G with a formula $crgadg(C) \equiv \sigma \wedge (|C| = 4) \wedge \forall x \in C\, \neg\alpha_1(x)$ where σ

routinely describes the possible edge sets of a 4-cycle on a given set of vertices C. The shortcut $|\,.\,| = 4$ has an obvious implementation in MSO. Then we use

$$\beta_1^{adj}(u,v) \equiv \exists X \big[\, \forall x \in X\ (\neg\alpha_1(x)) \wedge mcon(u,v,X) \wedge \wedge \forall C \big((crgadg(C) \wedge X \cap C \neq \emptyset) \to |X \cap C| = 3 \big) \big] . \tag{4}$$

The meaning of $\beta_1^{adj}(u,v)$ is that there exists a path P between u,v using only (besides u,v) marked internal vertices X, and such that P intersects every crossing-gadget 4-cycle in exactly three vertices which ensure that P is not "making a turn" at a crossing.

The last step in the proof is to verify the two conditions of Definition 5.2. While the second condition is trivially true since β_1^{adj} is a symmetric binary relation, the first one requires an efficient algorithm constructing, for each $H \in \mathscr{K}$, a graph $G_H \in \mathscr{L}$ such that $G_H^I \cong H$. This is done as follows:

i. A "nice" drawing of H in the plane is found (and fixed) such that no two edges cross more than once, no three edges cross in one point, and no edge passes through another vertex.
ii. For every degree-1 vertex $w \in V(H)$, the unique edge $\{x,w\} \in E(H)$ is replaced with a path of length 3 on $\{x,w,w_1,w_2\}$ where w_1, w_2 are new vertices. (This is needed since degree-1 vertices have special meaning in the interpretation.)
iii. Finally, every edge crossing is naturally replaced with a copy of the gadget from Figure 6. The resulting planar graph is named G_H.

It is routine to verify that $G_H^I \cong H$. ∎

Lemma 5.4 *The MSO_1 theory of all simple undirected graphs is efficiently interpretable in the MSO_1 theory of simple $\{1,3\}$-regular graphs. Moreover, this interpretation preserves planarity.*

Proof. Let $\mathscr{G}$ be the class of all simple graphs, and $\mathscr{R}$ denote the class of all simple $\{1,3\}$-regular graphs. We start the proof by showing a construction, for $H \in \mathscr{G}$, of the graph $G = G_H \in \mathscr{R}$ such that $G^I \cong H$ in the intended interpretation I.

For each vertex $v \in V(H)$ of degree d, we create a new vertex r_v adjacent to two (three if $d = 0$) other new vertices of degree 1. If $d > 0$, then we also create a new bicolored (black–white) cycle C_v of length $2d + 2$ such that each its black vertex is adjacent to a new vertex of degree 1, one of the white vertices is adjacent to r_v, and the d edges formerly incident with v in H are now one-to-one attached to the remaining d white vertices of C_v. See Figure 7. The resulting graph is our G_H. Notice that if H is planar, then the cyclic order of edges incident with each C_v can be preserved, and so G_H will also be planar.

The domain of H can be identified within G_H with the formula

$$\alpha_2(v) \equiv \exists x,y \left[\, x \neq y \wedge adj(v,x) \wedge adj(v,y) \wedge \forall z \big((adj(z,x) \vee adj(z,y)) \to z = v \big) \right] \tag{5}$$

meaning simply that v has (at least) two neighbours of degree 1.

Before interpreting the adjacency relation of H in G_H, we have to identify the cycles C_v from our construction. Notice that these are the only induced cycles of G_H with the property that every second of their consecutive vertices has a neighbour of degree 1 (for instance, each

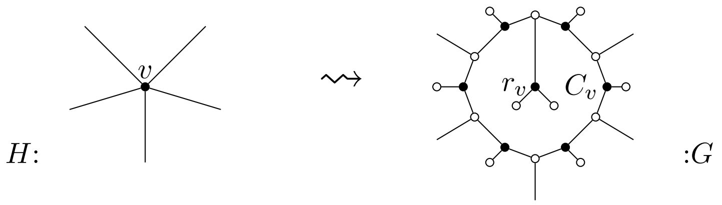

Figure 7: The *vertex-gadget* of Lemma 5.4, replacing vertices of H with $\{1,3\}$-regular parts in G.

edge of G_H coming from an edge of H has both ends with all neighbours of degree 3). In this sense we write

$$\varrho(U) \equiv \mathit{cycle}(U) \wedge \forall x, y \in U \left[\mathit{adj}(x,y) \rightarrow \bigvee_{z=x,y} \exists w \left(\mathit{adj}(z,w) \wedge \forall t \left(\mathit{adj}(t,w) \rightarrow t = z\right)\right) \right] \tag{6}$$

where $\mathit{cycle}(U)$ is an MSO$_1$ predicate saying that U induces a cycle in the graph, and finally,

$$\beta_2^{adj}(u,w) \equiv \exists U, W\, \exists u_1, u_2 \in U, w_1, w_2 \in W \; \varrho(U) \wedge \varrho(W) \wedge \mathit{adj}(u_1,w_1) \wedge \mathit{adj}(u_2,u) \wedge \mathit{adj}(w_2,w)\,. \tag{7}$$

We have finished description of the intended interpretation I, and it is now straightforward that $G_H^I \cong H$ for every $H \in \mathscr{G}$, cf. Definition 5.2. The proof is complete. ■

Lemma 5.5 *For every MSO$_1$ formula φ there exists an MSO$_1$ formula φ_1 such that the following holds: For every $\{1,3\}$-regular graph G and every subdivision G_1 of G, it is $G_1 \models \varphi_1$ if and only if $G \models \varphi$. Furthermore, $|\varphi_1| = \mathcal{O}(|\varphi|)$ and φ_1 is computable from φ in polynomial time.*

Proof. An alternative view of the situation is that we are interpreting the MSO$_1$ theory of $\{1,3\}$-regular graphs in the class of all their subdivisions. We construct φ_1 as the formula φ^I in such an interpretation I: Since G is $\{1,3\}$-regular, we simply identify the domain of G (its vertex set) inside G_1 with $\alpha_3(v) \equiv \neg\, \mathit{deg}_2(v)$ where $\mathit{deg}_2(v)$ routinely expresses that v is of degree two in G_1.

We moreover recall (1) the MSO$_1$ formula $\mathit{con}(u,v,X)$ meaning that u is connected to v via the vertices of X (in G_1). The adjacency relation of G is then replaced with $\beta_3^{adj}(u,v) \equiv \exists X \big(\mathit{con}(u,v,X) \wedge \forall y \in X\ \mathit{deg}_2(y)\big)$. Clearly, $G_1 \models \beta_3^{adj}(u,v)$ if and only if u and v are connected with a path in G_1 created by subdividing an edge $\{u,v\}$ of G. The rest follows trivially. ■

Proof of Theorem 5.1. We apply the chain of interpretations I_1, I_2, I_3 from Lemmas 5.3, 5.4, and 5.5 in this order to the formula φ, and obtain the resulting formula $\psi \equiv ((\varphi^{I_1})^{I_2})^{I_3}$. Following the constructive proofs of the aforementioned lemmas, we construct a graph G such that $H \cong (G^{I_2})^{I_1}$. Notice that the interpretations are applied in reverse order (cf. Figure 5). Then, part b) of the statement of Theorem 5.1 follows from Definition 5.2, and part c) follows from Lemma 5.5. Part d) is true since the interpretations I_1, I_2, I_3 are all efficient, i.e., are computable in polynomial time, and each of the translated formulas grows linearly in size. ■

6 Nonexistence of good digraph width measures

In this core section we finally prove some nearly optimal negative answers to the questions raised in the Introduction and at the end of Section 3. To recapitulate, we have asked whether it is possible to define a digraph width measure that is closed under some reasonable notion of a directed minor (e.g., Definition 4.4) and that is still powerful (Definition 3.1) analogously to ordinary treewidth. We also recall the property of being treewidth-bounding (which we want to avoid) from Definition 3.3.

Recall that a digraph width measure δ is closed under taking directed topological minors if there is an absolute constant c such that, for each digraph D, the δ-width of any directed topological minor of D is at most $\delta(D)+c$. The major existing ones are such by Corollary 4.10. We moreover give the following relaxed definition to make our negative results slightly stronger:

Definition 6.1 *A digraph width measure δ is* weakly closed under taking *directed topological minors if there exists a computable function w such that, for each digraph D, the δ-width of any directed topological minor of D is at most $w\big(\delta(D)\big)$.*

Although the aforementioned three requirements are enough to state the core nonexistence result in latter Theorem 6.7, we mention one more technical property that a reasonable digraph width measure should possess—we do not want to allow the measure to "keep computationally excessive" information in the orientation of edges. This requirement is important in intermediate Theorem 6.6. Formally:

Definition 6.2 *A digraph width measure δ is* efficiently orientable *if there exist a computable function h, and a polynomial-time computable function $r : \mathscr{G} \to \mathscr{D}$ (from the class of all graphs to that of digraphs), such that for every undirected graph $G \in \mathscr{G}$, we have $U(r(G)) = G$ and*

$$\delta(r(G)) \leq h(\,\min\{\delta(D) : D \text{ a digraph s.t. } U(D) = G\})\,.$$

To explain, a digraph width measure δ is efficiently orientable if, given any undirected graph G, one can orient its edges in time polynomial in $|G|$ to obtain a digraph with near-optimal δ-width. Quite many width measures possess this natural property, as the following proposition shows.

Proposition 6.3 *DAG-width, Kelly-width, digraph clique-width, and bi-rank-width are all efficiently orientable.*

Proof. As noted in Section 3, DAG-width and Kelly-width attain their globally minimum values on DAGs. On the other hand, clique-width and bi-rank-width attain an optimal value on symmetric orientations of graphs (replacing each edge by a pair of opposite arcs). Therefore given an undirected graph, one can easily orient its edges to obtain a digraph whose width is optimal wrt. to each of these width measures. ∎

Furthermore, to better illustrate the practical meaning of the words "to keep computationally excessive information in the orientation of edges", we point out in advance the counteractive example of Theorem 6.10: It is possible to encode the 3-colorability of a graph in a low-width orientation of its edges, and in this way to "cheat" when claiming fixed parameter tractability of the 3-Colouring problem on general graphs wrt. such an artificial parameter.

Our coming proofs also rely on some ingredients from the Graph Minors:

Theorem 6.4 ([37]) *If H is a planar undirected graph, then there exists a number n_H such that for every G of treewidth at least n_H, H is a minor of G.*

Proposition 6.5 (folklore) *If H is a minor of G and the maximum degree of H is three, then H is a topological minor of G.*

With all the ingredients at hand, we now state and prove our first main result.

Theorem 6.6 *Let δ be a digraph width measure with the following properties*

a) δ is not treewidth-bounding;

b) δ is weakly closed under taking directed topological minors;

c) δ is efficiently orientable.

Then δ is not powerful *unless* P = NP.

Proof. We assume that δ is powerful and show that for every MSO_1-definable property φ of undirected graphs there exists a polynomial-time algorithm that decides, given as input an undirected graph H, whether $H \models \varphi$. Since, by Example 2.5, there are MSO_1 properties φ such that deciding whether $H \models \varphi$ is NP-hard, this would imply that P = NP.

Given an MSO_1-formula φ and an undirected graph H, we construct a $\{1,3\}$-regular planar graph G and an MSO_1-formula ψ as in Theorem 5.1. Let G_1 be the 1-subdivision of G (i.e. replacing every edge of G with a path of length two). We claim that, under the assumptions (a) and (b), there exists an orientation D of G_1 such that $\delta(D) \leq k$, for some constant k dependent only on δ.

We postpone the proof of this claim, and show its implications first. Since δ is efficiently orientable, by Definition 6.2 and the existence of such D, we can efficiently construct another orientation D_1 of G_1 such that $\delta(D_1) \leq h(k)$, for some computable function h. Note that since k is a constant, the width of D_1 is at most a constant. Let ψ_1 be the (directed) MSO_1 formula obtained from ψ by replacing $adj(u,v)$ with $(arc(u,v) \vee arc(v,u))$. Then, by Theorem 5.1, $H \models \varphi$ iff $D_1 \models \psi_1$, and hence we have a polynomial reduction of the problems $H \models \varphi$ onto $D_1 \models \psi_1$. Since δ is assumed to be powerful, the latter problem can be solved by an XP algorithm wrt. the constant parameter $h(k)$, that is, in polynomial time.

We now return to our claim. Since δ is not treewidth-bounding, there is $k' \geq 0$ such that the class $\mathcal{U} = \{U(D) \colon \delta(D) \leq k'\}$ has unbounded treewidth. By Theorem 6.4, there exists D_0 such that $\delta(D_0) \leq k'$ and $U(D_0)$ contains a G_1-minor. Since the maximum degree of G_1 is three, by Proposition 6.5, G_1 is a topological minor of $U(D_0)$ and hence some subdivision G_2 of G_1 is a subgraph of $U(D_0)$. Therefore there exists D_2, a subdigraph of D_0, with $U(D_2) = G_2$. Finally, by Lemma 4.6 one can contract 2-paths in D_2, if necessary, to obtain a digraph D_3 with $U(D_3) = G_1$. Clearly D_3 is a directed topological minor of D_0 and since δ is weakly closed under taking directed topological minors, we have $\delta(D_3) \leq k = h'(k')$ for some computable function h' dependent on δ. This completes the proof of our claim and the theorem. ∎

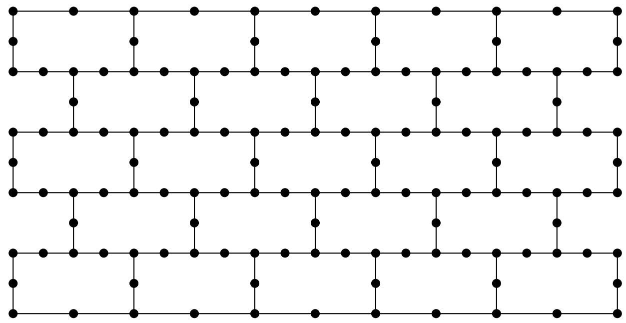

Figure 8: Graph H_6, the 1-subdivision of a 6×6 cubic grid.

Secondly, we further strengthen Theorem 6.6 by removing even the assumption (c) of efficient orientability as follows.

Theorem 6.7 *Let δ be a digraph width measure with the following properties*

a) δ is not treewidth-bounding;

b) δ is weakly closed under taking directed topological minors.

Then δ is not powerful *unless* NP $\subseteq$ P/poly.

A small price we have to pay for the stronger formulation in Theorem 6.7 is the need for a stronger complexity assumption, namely that NP $\not\subseteq$ P/poly instead of NP $\neq$ P. Recall that P/poly denotes the polynomial-time complexity class with a *polynomially-bounded advice function*, i.e. the class of languages that have polynomial-size circuits. By the Karp-Lipton theorem [27], NP $\subseteq$ P/poly would imply that the polynomial hierarchy collapsed to the level Σ_2^P (which is not considered likely).

Proof. (Theorem 6.7) We start the proof similarly to that of Theorem 6.6—assume that δ is powerful and show that then for every MSO_1-definable property ψ of undirected graphs there exists a polynomial-time algorithm that, given as input an undirected graph G and a polynomial-size advice depending only on $|V(G)|$, decides whether $G \models \psi$. By Theorem 5.1, we may assume G to be $\{1,3\}$-regular planar and ψ invariant under subdivisions of G.

To informally see how one can trade efficient orientability of the width measure δ for a polynomially-bounded advice function A depending only on the input size, it is helpful to deconstruct the proof of Theorem 6.6: The crux of the former proof is in showing that, given a $\{1,3\}$-regular planar graph G, there exists an orientation D of the 1-subdivision G_1 of G with small δ-width. While the former existential proof of D subsequently needed efficient orientability of δ to actually construct some (possibly different) orientation D_1 with small δ-width, we now bypass that point with a constructive argument for D. Precisely, we ask the advice A to give us a δ-optimal orientation of a suitably-sized grid (a "generic picture") in which we then find our particular D as a topological minor.

Firstly, we recall the notion of a cubic grid (also known as a wall or hexagonal grid [10]), and denote by H_r the 1-subdivision of the $r \times r$ cubic grid (see Figure 8). Note that H_r is always planar. The fact (a) that δ is not treewidth-bounding together with Theorem 6.4 and Proposition 6.5 imply there exists a constant $k' \geq 0$ such that members of the class $\mathcal{U} = \{U(D): \delta(D) \leq k'\}$ contain H_r as topological minors for arbitrarily large values of r. So, for any r some orientation D'_r of a suitable subdivision of H_r is a subdigraph of some $D_r^+ \in \mathcal{U}$.

Secondly, though this D'_r may be very large, we can again use Lemma 4.6 about contracting 2-paths to argue that there exists a directed topological minor D_r of D'_r such that $U(D_r) = H_r$. Using (b), we get that $\delta(D_r) \leq k = w(k')$ is bounded irrespectively of r, and so is $\delta(D^o) \leq k$ for any directed topological minor D^o of D_r using Proposition 4.5.

And, thirdly, it is known that any cubic planar graph G on n vertices is a topological minor of a $cn \times cn$ cubic grid where c is a universal constant [1], and that one can embed G in such a grid in time $\mathcal{O}(n \log n)$ [14]. This clearly extends to $\{1, 3\}$-regular planar graphs G. We can now tie up the loose threads as follows.

Let G be a given $\{1, 3\}$-regular planar graph on n vertices. We ask our polynomially-bounded advice function A to give us a digraph $F = A(n)$ such that $U(F) = H_r$ where $r = cn$, and F is a directed topological minor of some member of the aforementioned $\mathcal{U}$. Then we construct an embedding of G into H_r, and consequently obtain a subdigraph D_1 of D such that $U(D_1)$ is a subdivision of G. Hence $\delta(D_1) \leq k$. Let, moreover, ψ_1 be the (directed) MSO_1 formula obtained from ψ by replacing $adj(u, v)$ with $(arc(u, v) \vee arc(v, u))$. Since our ψ is invariant under subdivisions of $\{1, 3\}$-regular graphs, $G \models \psi$ is equivalent to $D_1 \models \psi_1$.

So far, the construction of D_1 and ψ_1 clearly falls into the complexity class P/poly. If δ was powerful, then we could decide $D_1 \models \psi_1$ in XP time wrt. constant width $\delta(D_1) \leq k$, i.e. in P time. Hence we would get a P/poly time algorithm for deciding $G \models \psi$ (equivalent to $D_1 \models \psi_1$). To finish the proof it remains to note that (Theorem 5.1) there exists abundance of formulas ψ for which the decision problem $G \models \psi$ is NP-hard. ■

Remark 6.8 *Note that the proof of Theorem 6.7 also shows another possible "intermediate" reformulation of Theorem 6.6; one in which the condition (c) of efficient orientability is required to hold only for subdivisions of cubic grids (instead of for all graphs).*

Necessity of the assumptions. Discussing once again the core outcome of this section—that a powerful digraph width measure essentially "cannot be stronger" than ordinary undirected treewidth, unless NP $\not\subseteq$ P/poly—brings us to a natural question of whether the technical assumptions of the result are necessary, or, put differently, whether our result can be strengthened by weakening (some of) its essential assumptions. We address this question in the remainder of this section.

First of all, the assumption (a) of δ not being treewidth-bounding is obviously unavoidable in order to deal with the richer universe of digraphs (cf. Section 3).

Next, the assumption (b) of δ being closed under taking directed topological minors has been intensively discussed in Section 4, namely in connection with the undirected case and cops-and-robber game characterizations. Here we moreover argue that both Theorems 6.6 and 6.7 are the strongest possible in that one cannot relax the condition (b) to "closed under subdigraphs" (or even under well-studied butterfly minors or digraph immersions).

For that purpose we construct an example of an artificial width measure which attains its "power" by subdividing every edge with a tower-exponential number of new vertices (and thus giving sufficient time margin for lengthy computation—an idea occurring, slightly differently in the undirected setting, also in [31]). Since this is certainly not a desirable behaviour of a width measure, it provides additional strong justification for condition (b) of both Theorems 6.6 and 6.7.

Proposition 6.9 *There exists a* ***powerful*** *digraph width measure* δ *such that*

a) δ *is not treewidth-bounding;*

b) δ *is monotone under taking subdigraphs;*

c) δ *is efficiently orientable.*

Moreover, the same remains true if we replace (b) with

b') δ *is monotone under taking subdigraphs of* D *and such contractions of arcs* $a \in A(D)$ *that create no new directed paths in* D/a *(i.e., under* butterfly minors*), or with*

b") δ *is monotone under taking digraph immersions of* D*. A digraph* H *is an* immersion *of* D *if* $V(H)$ *is mapped injectively into* $V(D)$ *such that the edges of* H *are mapped into edge-disjoint directed paths of* D *connecting the respective vertex images.*

Proof. We are going to apply the following modified version of Courcelle's Theorem [7]: *There exists a computable function* g *such that for all digraphs* D *and* MSO_2 *definable digraph properties* φ*, one can decide whether* $D \models \varphi$ *in time* $\mathcal{O}\big(g(|\varphi| + |V_3(D)|) \cdot |V(D)|\big)$*.* Recall that $V_3(D) \subseteq V(D)$ denotes the subset of those vertices having at least three neighbours in D. The original version of Courcelle's theorem states something stronger—it uses the quantifier depth (rank) of φ instead of $|\varphi|$ and the treewidth of $U(D)$ instead of $|V_3(D)|$—but for undirected graphs. The extension to digraph MSO_2 is straightforward (and not strictly needed here).

We give an explicit definition of δ. For an undirected graph G, we denote by $\mathrm{dist}_G(u, v)$ the length of a shortest path between vertices u and v in G; if there is no such path in G then $\mathrm{dist}_G(u, v) = \infty$. Let g be the function as in Courcelle's theorem as stated previously. Without loss of generality, we can assume that g is non-decreasing. For a digraph D, we define

$$\delta(D) = \begin{cases} 1, & \text{if } \mathrm{dist}_{U(D)}(u, v) \geq g(2 \cdot |V_3(D)|) \text{ for all pairs } u \neq v \in V_3(D); \\ |V(D)|, & \text{otherwise}. \end{cases} \tag{8}$$

We first show that δ fulfils the claimed properties. First, notice that δ does not depend on the orientation of edges; that is, $U(D_1) = U(D_2)$ readily implies $\delta(D_1) = \delta(D_2)$. Hence, (c) δ is efficiently orientable in linear time. Second, if we take any undirected graph G (of arbitrarily large treewidth) and subdivide every edge of G with $g(2 \cdot |V(G)|)$ vertices, then $\delta(D) = 1$ holds for every orientation D of $G = U(D)$. Therefore, (a) δ cannot be treewidth-bounding.

Third, concerning (b), let D be a digraph and let F be a subdigraph of D. We have to show that $\delta(D) \geq \delta(F)$. This is clearly true if $\delta(D) = |V(D)|$, and so assume $\delta(D) = 1$. Take any vertex pair $u, v \in V_3(F) \subseteq V_3(D)$. Then by our assumption, $\mathrm{dist}_{U(D)}(u, v) \geq g(2 \cdot |V_3(D)|)$, and $g(2 \cdot |V_3(D)|) \geq g(2 \cdot |V_3(F)|)$ by assumed monotonicity of g. Hence $\mathrm{dist}_{U(F)}(u, v) \geq \mathrm{dist}_{U(D)}(u, v) \geq g(2 \cdot |V_3(F)|)$, and consequently $\delta(F) = 1 \leq \delta(D)$.

It remains to show that δ is powerful. Let φ be an MSO_1-definable (undirected) property, and let D be an input digraph. We simply apply Courcelle's theorem to decide $U(D) \models \varphi$, and prove that this is an XP (even FPT) algorithm wrt. the parameter $\delta(D)$. If $\delta(D) = |V(D)|$, then indeed $\mathcal{O}\big(g(|\varphi| + |V_3(D)|) \cdot |V(D)|\big) = \mathcal{O}\big(g(\delta(D)) \cdot |V(D)|\big)$ for every fixed φ. So assume $\delta(D) = 1$. If every component of $U(D)$ contains at most one cycle, then the treewidth of D is at most two and the case follows trivially. Otherwise, some two vertices of $V_3(D)$ are connected by a path and so $|V(D)| > g(2 \cdot |V_3(D)|)$. Then the run-time bound of Courcelle's theorem gives $\mathcal{O}\big(g(|\varphi| + |V_3(D)|) \cdot |V(D)|\big) = \mathcal{O}\big(\max\{g(2|\varphi|), |V(D)|\} \cdot |V(D)|\big) = \mathcal{O}(|V(D)|^2)$ for fixed φ.

As for the condition (b'), we simply extend the previous arguments. Note that if a 2-path P between $u, v \in V_3(D)$ has "alternately oriented" arcs (as in Figure 2), then no arc of P is contractible without introducing new directed paths. If, informally, this happened for all possible 2-paths in D, then taking butterfly minors in D would be almost equivalent to taking subdigraphs of D.

Formally, let $S(D) \subseteq V(D)$ denote the subset of those vertices s in D that either s has no in-neighbours (i.e., s is a *source*) or s has no out-neighbours (i.e., s is a *sink*). Instead of (8) we give the following modified definition:

$$\delta'(D) = \begin{cases} 1, & \text{if for all pairs } u \neq v \in V_3(D) \text{, every 2-path in } D \text{ between } u, v \\ & \text{has at least } g(2 \cdot |V_3(D)|) \text{ internal vertices in S(D);} \\ |V(D)|, & \text{otherwise.} \end{cases} \tag{9}$$

This measure clearly again satisfies (a) and (c), and Courcelle's algorithm on D is in FPT wrt. the width $\delta'(D)$, too.

To prove (b') for δ', we proceed as follows. Let D be a digraph and let F be a subdigraph of D. If $\delta(D) = |V(D)|$, then clearly $\delta(F) \leq \delta(D)$, and so assume $\delta(D) = 1$. Take any 2-path P between a pair $u, v \in V_3(F) \subseteq V_3(D)$. Then P is obtained by contracting some arcs in a path (undirected) $Q \subseteq D$. Since Q is a union of (at least one) 2-paths in D, by the assumption $\delta(D) = 1$ we have $|V(Q) \cap S(D)| \geq g(2 \cdot |V_3(D)|) \geq g(2 \cdot |V_3(F)|)$. Notice that if a is an arc in D with both ends in $S(D)$, then its contraction necessarily creates a new path in D/a and so such a contraction of a is not allowed in (b'). Therefore, $|V(P) \cap S(F)| \geq |V(Q) \cap S(D)| \geq g(2 \cdot |V_3(F)|)$, and consequently $\delta(F) = 1 \leq \delta(D)$.

Finally, (b") has the same proof as (b') since a vertex which is a source or a sink cannot be immersed in a digraph by definition. ■

Lastly, we take a closer look at the property of δ being efficiently orientable. Though this condition (c) of Theorem 6.6 is, after all, completely avoided in stronger Theorem 6.7, it nevertheless deserves further discussion in our opinion: It is not unreasonable to assume a digraph width measure to be efficiently orientable since most known digraph measure are, e.g. Proposition 6.3. Furthermore, efficient orientability prevents digraph measures from "keeping excessive information" in the orientation of arcs, such as in the following example:

Proposition 6.10 *There exists a digraph width measure δ such that*

a) δ is not treewidth-bounding;

b) δ is monotone under taking directed topological minors;

c) for every 3-colorable graph G there exists an orientation D, $U(D) = G$, such that $\delta(D) = 1$;

d) and for every $k \geq 1$, on any digraph D with $\delta(D) \leq k$, one can decide in (FPT) time $\mathcal{O}(3^k \cdot n^2)$ whether $U(D)$ is 3-colorable, and find a 3-colouring if it exists.

To briefly comment on this result, we emphasize the following. On the one hand, there is nothing specially interesting in solving the 3-colorability problem on digraphs—this is taken just as an example of an NP-complete problem, and it is no surprise that solutions to NP-complete problems can be *somehow* encoded in the orientation of arcs.

On the other hand, it is in our opinion really unexpected that one can naturally encode an "excessive information" (an NP-completeness oracle, to be precise) in the orientation of arcs such that this encoding is invariant (b) under taking directed topological minors.

Proof. (Theorem 6.10) We start by defining our digraph width measure δ. For a digraph D,

$$\delta(D) = \begin{cases} 1, & \text{if the arcs of } D \text{ encode a 3-colouring of } U(D); \\ |V(D)|, & \text{otherwise.} \end{cases} \tag{10}$$

We say that the arcs of a digraph D *encode a 3-colouring* if, for every directed (s,t)-path in D with $s,t \in V_3(D)$, we have that either s has no in-neighbours (i.e. s is a *source*) or t has no out-neighbours (i.e. t is a *sink*).

The crucial property of this definition (10) is that if $\delta(D) = 1$, then $U(D)$ is 3-colorable. To see this, assume that the arcs of D encode a 3-colouring, and let S_1, S_2, S_3 be a tripartition of $V_3(D)$ such that S_1 is the set of all sources, S_3 is the set of all sinks, and S_2 the remaining vertices. The sets S_1 and S_3 are obviously independent in $U(D)$ as they contain only source/sink vertices. If S_2 was not independent, then there would be an arc $a = (u,v)$ forming a directed path of length 1 between $u,v \in S_2 \subseteq V_3(D)$, and so one of u,v would actually belong to $S_1 \cup S_3$, a contradiction. As the remaining vertices of $V(D) \setminus V_3(D)$ have at most two neighbours each, we can use a straightforward greedy algorithm to extend the partition (S_1, S_2, S_3) of $V_3(D)$ into a partition of $V(D)$ into three independent sets. Therefore $U(D)$ is 3-colorable.

We now have to show that δ fulfils all the four conditions of Theorem 6.10. We start with (c); assume a graph G with a proper 3-colouring $c : V(G) \to \{1,2,3\}$. We construct an acyclic orientation D, $U(D) = G$, by directing every edge $\{u,v\} \in E(G)$ as $(u,v) \in A(D)$ such that $c(u) < c(v)$. Then $\delta(D) = 1$ since D actually has no directed path of length three. Consequently, for the complete bipartite graph $K_{n,n}$ there is an orientation D' such that $\delta(D') = 1$ and $U(D') = K_{n,n}$, but the treewidth of $K_{n,n}$ is n. Hence also (a); δ is not treewidth-bounding.

To prove (d); let D be any digraph. If $\delta(D) = k \geq 2$, then $n = |V(D)| = k$. By trying all possible 3-colourings we can solve the task in time $\mathcal{O}(3^k \cdot n^2)$. On the other hand if $\delta(D) = 1$, then (10) D encodes a 3-colouring, and we can compute a valid 3-colouring of $U(D)$ in time $\mathcal{O}(n^2)$, as outlined above.

The last step is to show (b); that δ is monotone under taking directed topological minors for every $k \geq 1$. Let D be a digraph. If $\delta(D) \geq 2$, then $\delta(D) = |V(D)|$ and thus $\delta(F) \leq \delta(D)$ for every directed topological minor F of D. Therefore let $\delta(D) = 1$, i.e. D encodes a 3-colouring. For a directed topological minor F of D, assume any directed (s,t)-path in F with $s,t \in V_3(F) \subseteq V_3(D)$. By Definition 4.4, no (s,t)-path can be created by contractions of 2-contractible arcs in D, and so a directed (s,t)-path exists in D. Then, up to symmetry, s is a source in D since $\delta(D) = 1$. Again by the definition of 2-contractible arcs, s must be a source in F, too. Consequently $\delta(F) = 1 \leq \delta(D)$, and the condition is proved. ■

7 Conclusions

Looking at the main result of this paper (Theorem 6.7) from a yet slightly different perspective than in Section 3, one can conclude the following:

> Any algorithmically useful digraph width measure δ that is substantially different from undirected treewidth must possess the following property: There exist digraphs of low δ-width such that δ grows very high on their directed topological minors.

Since "standard" cops-and-robber games remain closed on directed topological minors, we conclude that a digraph width measure that allows efficient decisions of MSO_1-definable digraph properties on classes of bounded width cannot be defined or (even asymptotically) characterized using such games. All this gives more weight to the argument [15] that bi-rank-width [25] and its generalizations are the best (though not optimal) currently known candidates for a *good* digraph width measure from the algorithmic perspective.

Nevertheless, the area of treewidth-like digraph width measures (and of the related cops-and-robber games) remains a very interesting topic in pure combinatorics, with solid structural foundations and abundance of fundamental open questions. We believe that the results and suggestions contained in our paper will also lead to new ideas and research directions in this combinatorial area of digraph width measures—an area that seems to be stuck at this moment.

Finally, our main result also leaves room for other ways of overcoming the problems with the currently existing digraph width measures. We have asked for width measures that are powerful, i.e., all MSO_1-definable digraph properties are decidable in polynomial time on digraphs of bounded width. What happens if we relax this requirement? We can ask for more time, like subexponential running time, or we can ask for restricted classes of MSO_1-definable digraph properties. Currently, we are not aware of any noticeable progress in this direction.

On the other hand, could it still be possible to relax a bit the requirement of monotonicity under taking directed topological minors? Theorem 6.9 suggests a strict negative answer, but one can also take an approach analogous to Kreutzer–Tazari's [28] for undirected graphs; say, a directed measure δ is *nearly treewidth-bounding* if $\mathbf{tw}(U(D)) \leq b(\delta(D)) \cdot \log^c |V(D)|$ for some constant c and computable function b. Then, with a stronger condition (a') that δ is not nearly treewidth-bounding and weaker (b') δ is monotone under taking subdigraphs, a claim analogous to Theorem 6.7 could still be true. Indeed, the latter has been proved true just very recently in [17] with a suitable stronger complexity assumption.

References

[1] A. Ageev. Complexity of the network median problem on planar grids. *Siberian Adv. Math.*, 5:1–9, 1995.

[2] S. Arnborg, J. Lagergren, and D. Seese. Easy problems for tree-decomposable graphs. *J. Algorithms*, 12(2):308–340, 1991.

[3] J. Barát. Directed path-width and monotonicity in digraph searching. *Graphs and Combin.*, 22(2):161–172, 2006.

[4] D. Berwanger, A. Dawar, P. Hunter, and S. Kreutzer. DAG-width and parity games. In *STACS'06*, volume 3884 of *LNCS*, pages 524–536. Springer, 2006.

[5] H. Bodlaender. Treewidth: Characterizations, Applications, and Computations. In *WG'06*, volume 4271 of *LNCS*, pages 1–14. Springer, 2006.

[6] X. Chen, X. Hu, and W. Zang. A min-max theorem on tournaments. *SIAM J. Comput.*, 37:923–937, 2007.

[7] B. Courcelle. The monadic second order logic of graphs I: Recognizable sets of finite graphs. *Inform. and Comput.*, 85:12–75, 1990.

[8] B. Courcelle, J. A. Makowsky, and U. Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000.

[9] B. Courcelle and S. Olariu. Upper bounds to the clique width of graphs. *Discrete Appl. Math.*, 101(1-3):77–114, 2000.

[10] R. Diestel. *Graph Theory*, volume 173 of *Graduate texts in mathematics*. Springer, New York, 2005.

[11] R. Downey and M. Fellows. *Parameterized complexity.* Monographs in Computer Science. Springer, 1999.

[12] Z. Dvořák, R. Škrekovski, and M. Tancer. List coloring squares of sparse subcubic graphs. *SIAM J. Discrete Math.*, 22:139–159, 2008.

[13] S. Fortune, J. E. Hopcroft, and J. Wyllie. The directed subgraph homeomorphism problem. *Theoret. Comput. Sci.*, 10:111–121, 1980.

[14] H. Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10:41–51, 1990.

[15] R. Ganian, P. Hliněný, J. Kneis, A. Langer, J. Obdržálek, and P. Rossmanith. On digraph width measures in parametrized algorithmics. In *IWPEC'09*, volume 5917 of *LNCS*, pages 161–172. Springer, 2009.

[16] R. Ganian, P. Hliněný, J. Kneis, A. Langer, J. Obdržálek, and P. Rossmanith. Digraph width measures in parametrized algorithmics. Submitted. Available from: `http://www.fi.muni.cz/~hlineny/Research/papers/kenny-full-v2.pdf`, 2012.

[17] R. Ganian, P. Hliněný, A. Langer, J. Obdržálek, P. Rossmanith, and S. Sikdar. Lower bounds on the complexity of MSO1 model-checking. In *STACS'12*, volume 14 of *LIPIcs*, pages 326–337. Dagstuhl Publishing, 2012.

[18] R. Ganian, P. Hliněný, and J. Obdržálek. Clique-width: When hard does not mean impossible. In *STACS'11*, volume 9 of *LIPIcs*, pages 404–415. Dagstuhl Publishing, 2011.

[19] R. Ganian, P. Hliněný, and J. Obdržálek. Unified approach to polynomial algorithms on graphs of bounded (bi-)bank-width. *European J. Combin.*, 2012. To appear. Available from `http://dx.doi.org/10.1016/j.ejc.2012.07.024`.

[20] M. Garey, D. Johnson, and L. Stockmeyer. Some simplified NP-complete graph problems. *Theor. Comput. Sci.*, 1:237–267, 1976.

[21] P. Hliněný, S. Oum, D. Seese, and G. Gottlob. Width parameters beyond tree-width and their applications. *The Computer Journal*, 51(3):326–362, 2008.

[22] P. Hunter. *Complexity and Infinite Games on Finite Graphs*. PhD thesis, University of Cambridge, 2007.

[23] P. Hunter and S. Kreutzer. Digraph measures: Kelly decompositions, games, and orderings. *Theoret. Comput. Sci.*, 399(3):206–219, 2008.

[24] T. Johnson, N. Robertson, P. D. Seymour, and R. Thomas. Directed tree-width. *J. Combin. Theory Ser. B*, 82(1):138–154, 2001.

[25] M. Kanté. The rank-width of directed graphs. arXiv:0709.1433v3, March 2008.

[26] M. Kanté and M. Rao. F-rank-width of (edge-colored) graphs. In *CAI'11*, volume 6742 of *LNCS*, pages 158–173. Springer, 2011.

[27] R. M. Karp and R. J. Lipton. Some connections between nonuniform and uniform complexity classes. In *STOC'80*, pages 302–309. ACM, 1980.

[28] S. Kreutzer and S. Tazari. Lower bounds for the complexity of monadic second-order logic. In *LICS'10*, pages 189–198. IEEE, 2010.

[29] M. Lampis, G. Kaouri, and V. Mitsou. On the algorithmic effectiveness of digraph decompositions and complexity measures. In *ISAAC*, volume 5369 of *LNCS*, pages 220–231. Springer, 2008.

[30] W. Mader. On topological tournaments of order 4 in digraphs of outdegree 3. *J. Graph Theory*, 21(4):371–376, 1996.

[31] J. A. Makowsky and J. Mariño. Tree-width and the monadic quantifier hierarchy. *Theoret. Comput. Sci.*, 303(1):157–170, 2003.

[32] J. Obdržálek. DAG-width – connectivity measure for directed graphs. In *SODA'06*, pages 814–821. ACM-SIAM, 2006.

[33] S. Oum and P. D. Seymour. Approximating clique-width and branch-width. *J. Combin. Theory Ser. B*, 96(4):514–528, 2006.

[34] M. O. Rabin. A simple method for undecidability proofs and some applications. In Y. Bar-Hillel, editor, *Logic, Methodology and Philosophy of Sciences*, volume 1, pages 58–68. North-Holland, Amsterdam, 1964.

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

[36] N. Robertson and P. D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory Ser. B*, 52(2):153–190, 1991.

[37] N. Robertson, P. D. Seymour, and R. Thomas. Quickly excluding a planar graph. *J. Combin. Theory Ser. B*, 62:323–348, 1994.

[38] M. Safari. D-width: A more natural measure for directed tree-width. In *MFCS'05*, volume 3618 of *LNCS*, pages 745–756. Springer, 2005.

[39] P. D. Seymour and R. Thomas. Graph searching and a min-max theorem for tree-width. *J. Combin. Theory Ser. B*, 58(1):22–33, 1993.

[40] A. Slivkins. Parameterized tractability of edge-disjoint paths on directed acyclic graphs. In *ESA'03*, volume 2832 of *LNCS*, pages 482–493. Springer, 2003.