Twin-Width Is Linear in the Poset Width

Jakub Balabán
Masaryk University, Brno, Czech Republic

Petr Hliněný
Masaryk University, Brno, Czech Republic

Abstract

Twin-width is a new parameter informally measuring how diverse are the neighbourhoods of the graph vertices, and it extends also to other binary relational structures, e.g. to digraphs and posets. It was introduced just very recently, in 2020 by Bonnet, Kim, Thomassé and Watrigant. One of the core results of these authors is that FO model checking on graph classes of bounded twin-width is in FPT. With that result, they also claimed that posets of bounded width have bounded twin-width, thus capturing prior result on FO model checking of posets of bounded width in FPT. However, their translation from poset width to twin-width was indirect and giving only a very loose double-exponential bound. We prove that posets of width d have twin-width at most $8d$ with a direct and elementary argument, and show that this bound is tight up to a constant factor. Specifically, for posets of width 2 we prove that in the worst case their twin-width is also equal 2. These two theoretical results are complemented with straightforward algorithms to construct the respective contraction sequence for a given poset.

2012 ACM Subject Classification Mathematics of computing → Combinatorics; Mathematics of computing → Graph algorithms

Keywords and phrases twin-width, digraph, poset, FO model checking, contraction sequence

Digital Object Identifier 10.4230/LIPIcs.IPEC.2021.6

Related Version *Full Version*: `http://arxiv.org/abs/2106.15337`

Funding Supported by the Czech Science Foundation, project no. 20-04567S.

1 Introduction

The new notion of twin-width (of graphs, digraphs, or matrices) was introduced just very recently, in 2020, by Bonnet, Kim, Thomassé and Watrigant [4], and yet has already found many very interesting applications. These applications span from efficient parameterized algorithms and algorithmic metatheorems, through finite model theory, to classical combinatorial questions. See also the series of follow-up papers [2, 1, 3, 5].

We leave formal definitions for the next section. Informally, in simple graphs, twin-width measures how diverse are the neighbourhoods of the graph vertices. E.g., cographs[1] have the lowest possible value of twin-width, 0, which means that the graph can be brought down to a single vertex by successively identifying twin vertices (hence the name, *twin-width*). Two vertices x and y are *twins* if they have the same neighbours in the graph, precisely $N(x) \setminus \{y\} = N(y) \setminus \{x\}$ (the concept of twin-width of graphs does not care about mutual adjacency of the identified vertices).

More generally, imagine we identify arbitrary two vertices x_1, x_2 in a graph G into a new vertex x; then the ordinary neighbours of new x will capture what the former neighbourhoods of x_1 and of x_2 in G have had in common (except each other of x_1, x_2), and we will additionally create new *red edges* from x to those vertices on which the former neighbourhoods of x_1 and

[1] Cographs are the graphs which can be built from singleton vertices by repeated operations of a disjoint union and taking the complement.

16th International Symposium on Parameterized and Exact Computation (IPEC 2021).
Editors: Petr A. Golovach and Meirav Zehavi; Article No. 6; pp. 6:1–6:13
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of x_2 disagreed in G. Moreover, all other previously created red edges at x_1 or x_2 will stay red and incident to x after the identification. Note that the former vertices x_1, x_2 are removed and no loop is created on x. Precisely, denote by $N(x_i)$ the ordinary ("black") neighbours of x_i in G and by $N_r(x_i)$ the red neighbours of x_i. After the identification of x_1 and x_2 into x, the ordinary neighbours of x will be the vertices of $\big(N(x_1) \cap N(x_2)\big) \setminus \big(\{x_1, x_2\} \cup N_r(x_1) \cup N_r(x_2)\big)$, and the red edges of x will go to the vertices of $\big((N(x_1) \Delta N(x_2)) \cup N_r(x_1) \cup N_r(x_2)\big) \setminus \{x_1, x_2\}$. With respect to this, a graph G has twin-width $\leq d$ if one can reduce G down into a single vertex by successively identifying pairs of its vertices such that the maximum degree of the subgraph on the red edges at every step of this reduction process is at most d.

We are, in particular, interested in the algorithmic metatheorem area. Namely, Bonnet et al. [4] have proved that classes of binary relational structures (such as simple graphs and digraphs) of bounded twin-width have efficient FO model checking algorithms. In one of previous studies on algorithmic metatheorems for dense structures, Gajarský et al. [7] proved that posets (which present a special case of simple digraphs) of bounded width admit efficient FO model checking algorithms. The *width of a poset* is the maximum size of an antichain in it. Since [4] have also proved that posets of bounded width have bounded twin-width, a combination of this finding with the FO model checking algorithm of [4] directly generalizes the algorithmic metatheorem of [7].

The proof of bounded width of posets in [4] is, however, indirect (it uses a characterization by so-called mixed minors of the adjacency matrix) and gives, for posets of width d, only a very loose upper bound of $2^{2^{\mathcal{O}(d)}}$ for the twin-width. Although the proof in [4] is, in principle, constructive, its intricacy makes it really hard to understand why posets of bounded width should have bounded twin-width, and which vertices to identify in the reduction process. In fact, as we will see in this paper, already for posets of width 2 there is no immediate way to optimally choose the pairs for identification.

The main contribution of our paper is in giving direct and tighter constructive linear lower and upper bounds for the twin-width of posets of width d. Precisely, the twin-width of such a poset in the worst case is at least $d-1$ and at most $8d-9$ (Proposition 2.4 and Theorem 3.1). Specifically for posets of width 2, we prove that their twin-width is also at most 2 and this bound cannot be further improved (Theorem 4.1). These results are accompanied by simple and fast algorithms to compute the corresponding contraction sequence.

Our refined results on twin-width of posets provide, in turn, better runtime bounds for FO model checking algorithms of posets [7] and of other classes which have previously been reduced to posets of bounded width, such as [8, 6, 9].

2 Preliminaries and formal definitions

We consider only finite graphs. Our graphs and digraphs are simple, meaning that they do not have parallel edges or loops, except that a simple digraph may have up to one oriented loop per vertex. Formally, a graph is a pair $G = (V, E)$ such that $E \subseteq \binom{V}{2}$, and a digraph is $G = (V, E)$ such that $E \subseteq V \times V$. We deal with (finite) partially ordered sets, shortly *posets*, which we represent as reflexive, antisymmetric and transitive digraphs. Let the *width of a poset P* be the maximum size of an antichain in P, that is the maximum independent set size in the digraph P.

We formally define twin-width using the "matrix-partitioning" view of [4, Section 5], and we restrict ourselves only to the symmetric twin-width which is relevant to graphs and posets. Let $\boldsymbol{A}$ be a square matrix with entries from a finite set L (e.g., $L = \{0, 1\}$ for undirected graphs and $L = \{0, 1, -1, 2\}$ for digraphs, as explained below), and assume that both the rows and the columns of $\boldsymbol{A}$ are indexed by the same ground set X. Let $\mathcal{R}$ denote any partition of X into nonempty sets. For two parts $R, Q \in \mathcal{R}$, the submatrix of $\boldsymbol{A}$ formed by the rows indexed by R and the columns indexed by Q is called the *$(R \times Q)$ zone* of $\boldsymbol{A}$. Naturally, a zone of $\boldsymbol{A}$ is *constant* if all entries in the zone are equal. For $P \in \mathcal{R}$, the *error value* (the "red degree") of a row part P (column part P) in $\boldsymbol{A}$ is the number of non-constant zones $P \times Q$ (zones $Q \times P$, respectively) in $\boldsymbol{A}$ over all $Q \in \mathcal{R}$ (including $Q = P$).

▶ **Definition 2.1.** *Let $\boldsymbol{A}$ be a square matrix with the rows and columns indexed by a ground set X, $|X| = n$. We say that the* symmetric twin-width *of $\boldsymbol{A}$* is at most *d if there exists a sequence of partitions $\mathcal{R}^1, \ldots, \mathcal{R}^n$ of X (a* contraction sequence*) such that;*

- *$\mathcal{R}^1$ is the finest partition of X ($|\mathcal{R}^1| = n$) and $\mathcal{R}^n$ is the coarsest partition of X ($|\mathcal{R}^n| = 1$),*
- *for each $i = 1, \ldots, n-1$, the partition $\mathcal{R}^{i+1}$ results by merging ("contraction" of) some two parts of $\mathcal{R}^i$, and*
- *for each $i = 1, \ldots, n$ and every $P \in \mathcal{R}^i$, the error value of the row P and the column P in $\boldsymbol{A}$ is at most d.*

For a quick illustration, consider Definition 2.1 applied to the adjacency matrix $\boldsymbol{A}_G$ of a graph G. Then the definition of symmetric twin-width of $\boldsymbol{A}_G$ coincides with the twin-width of G as stated in Section 1, except that the diagonal zones ($Q \times Q$ for $Q \in \mathcal{R}^i$ at step i) of $\boldsymbol{A}_G$ are often non-constant, and hence the symmetric twin-width of $\boldsymbol{A}_G$ may be equal or by one higher than the twin-width of G (this difference is neglected in [4]).

For a poset $P = (X, \leq_P)$, viewed as a digraph on the ground set X, we consider the matrix $\boldsymbol{A}_P$ defined as follows (according to [4]); $a_{u,v} = 1$ iff $u \leq_P v$, $a_{u,v} = -1$ iff $v \lneq_P u$, and $a_{u,v} = 0$ otherwise[2]. The *symmetric twin-width of P* is the symmetric twin-width of $\boldsymbol{A}_P$.

For our purpose of giving a closer relation between the width and the twin-width of posets it is, though, much more convenient to use a specialized definition which we call a *natural twin-width* of posets (for a distinction). We shortly write $a \sim_P b$ iff $a \leq_P b$ or $b \leq_P a$ (i.e., the vertices a, b are comparable in P). Our definition reads:

▶ **Definition 2.2** (Natural twin-width of a poset)**.** *A triple $P = (X, \leq_P, R)$ is a* red poset *if $P_0 = (X, \leq_P)$ is a poset and R is a set of unordered pairs of incomparable elements of P_0. The* red degree *of P is the maximum degree of the "red" graph (X, R).*

A contraction *of two vertices $x_1, x_2 \in X$ of P (into a new vertex x) creates the red poset $P' = (X', \leq'_P, R')$ where $X' = (X \setminus \{x_1, x_2\}) \cup \{x\}$ and*

- *$a \leq'_P b$ iff $a \leq_P b$ for all $a, b \in X \setminus \{x_1, x_2\}$, and $x \leq'_P x$,*
- *$a \leq'_P x$ (resp. $x \leq'_P a$) iff $a \leq_P x_1$ and $a \leq_P x_2$ (resp. $x_1 \leq_P a$ and $x_2 \leq_P a$),*
- *$R' = (R \restriction X') \cup R_1 \cup R_2$ where $R_1 = \{\{a, x\} : \{a, x_1\} \in R \vee \{a, x_2\} \in R\}$ and $R_2 = \{\{a, x\} : a \not\sim'_P x \wedge (a \sim_P x_1 \vee a \sim_P x_2)\}$ ($R \restriction X'$ stands for the restriction of R to X').*

In other words, the red edges R' of P' are; (i) those inherited from P which compose of the restriction of R to X' and the red edges R_1 formerly incident to x_1 or x_2 in P, and (ii) the new ones in R_2 between x and those vertices of $X \setminus \{x_1, x_2\}$ which compared in P to x_1 and to x_2 in different ways.

An (ordinary) poset $P_0 = (X, \leq_{P_0})$ has natural twin-width at most *d if the red poset $(X, \leq_{P_0}, \emptyset)$ can be reduced down to a single vertex by a sequence of contractions such that, at each step, the red degree is at most d. So, the natural twin-width of P_0 equals the minimum integer d such that P_0 has natural twin-width at most d.*

[2] If we considered general digraphs (which are not always antisymmetric), we would also consider value $a_{u,v} = 2$ iff both $(u, v), (v, u)$ were edges of the digraph.

With respect to the definition of a contraction (in red posets), the following is a useful convention: If a red poset $P = (X, \leq_P, R)$ resulted by a sequence of contractions from a poset $P_1 = (X_1, \leq_{P_1})$, then a vertex $y \in X$ uniquely corresponds to a set $Y \subseteq X_1$ of those vertices of P_1 which were contracted down to y, and hence we will chiefly refer to Y with the name $y \subseteq X_1$. Consequently, the vertices of such X at the same time form a partition of X_1 in P_1 (and, with negligible abuse of notation, X_1 itself is viewed as the partition of X_1 into singletons) which brings us very close to Definition 2.1 of symmetric twin-width.

In accordance with this convention, we will sometimes denote the vertex resulting from a contraction of the vertices x_1 and x_2 shortly by x_1x_2. See an example in Figure 1.

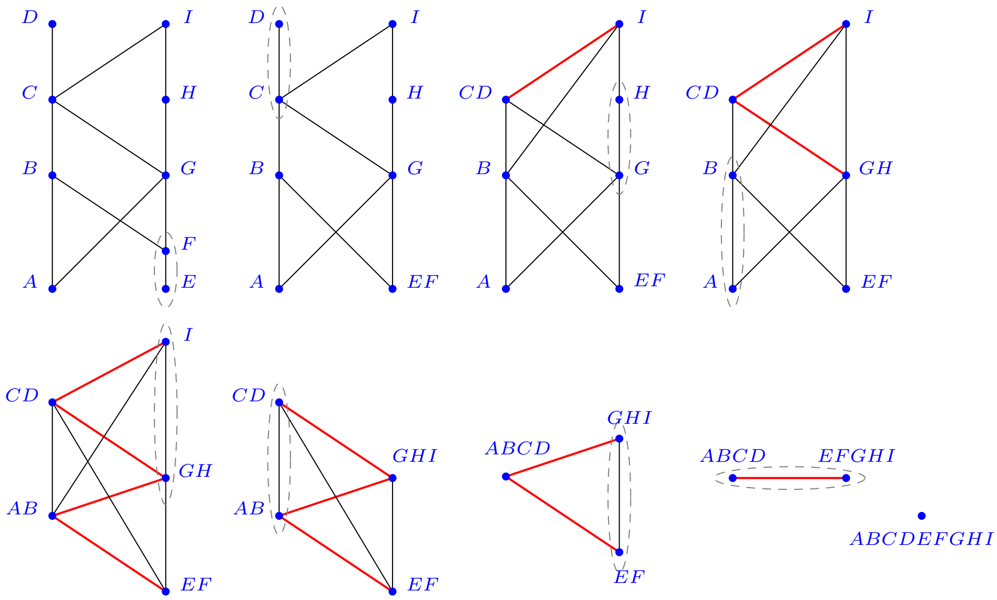

Figure 1 An example of a contraction sequence for the top-left poset (each step contracts the encircled pair), having red degree at most 2. As in Hasse diagrams, the edges (black) of the posets are oriented up, and we skip drawing edges which are implied by reflexivity and transitivity.

▶ **Proposition 2.3.** *If the symmetric twin-width of a poset P is d_s and the natural twin-width of P is d, then $d \leq d_s \leq d + 1$.*

Proof (sketch). We compare Definitions 2.1 and 2.2 for the same contraction sequence; a non-constant zone corresponds to a created red edge, and vice versa. The only difference is at the diagonal zones which may be non-constant while natural twin-width does not consider red loops, and so the symmetric twin-width may be by one higher than the natural one. ◀

2.1 Simple lower bound

▶ **Proposition 2.4.** *There exists a poset of width d and natural twin-width at least $d - 1$.*

Proof. For $d = 2$, we simply take a poset which has no pair of twin vertices, and so any contraction in it creates a red edge, that is, red degree $d - 1 = 1$. E.g., we take the poset formed by the divisibility relation on the set $\{2, 3, 4, 6\}$.

For $d \geq 3$ and $k \geq 4d - 8$, we construct a poset P on a ground set of $n = d(k + 1)$ vertices $C_1 \cup C_2 \cup \ldots \cup C_d$ where each $C_i = \{c_i^0, c_i^1, \ldots, c_i^k\}$ is a chain; $c_i^0 \leq_P c_i^1 \leq_P \ldots \leq_P c_i^k$. In the description of P, we consider indices i "modulo d", formally, we set $c_{d+1}^j = c_1^j, \ldots, c_{d+d}^j = c_d^j$.

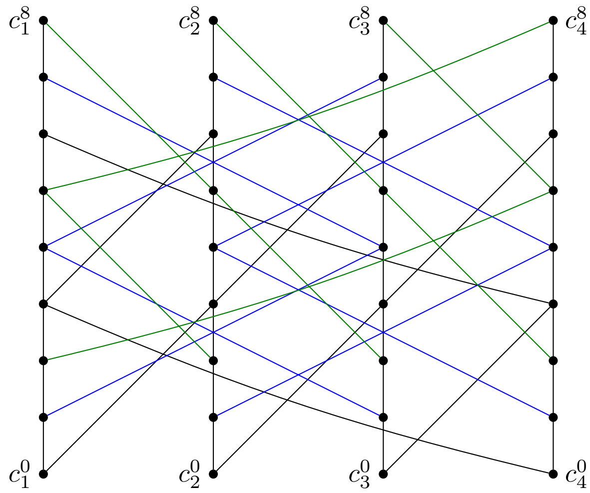

Figure 2 The poset from the proof of Proposition 2.4 for $d = 4$ and $k = 8$, depicted by its Hasse diagram.

Furthermore, for $i = 1, \dots, d$ and $j = 0, \dots, k-d+1$, with $a = 1 + (j \mod (d-1))$, we declare $c_i^j \leq_P c_{i+a}^{j+d-1}$. The rest of $\leq_P$ follows by the reflexive and transitive closure. See an example of this construction in Figure 2, and note that the inverse of P is isomorphic to P (informally, turning P "upside down" gives the same poset) which will be used to reduce the number of cases in the coming arguments by symmetry.

Our aim is to prove that a contraction of any pair of vertices of P already gives red degree $\geq d-1$. Suppose first that the contracted pair is from the same chain C_i, e.g., c_i^j and c_i^h for $0 \leq j < h \leq k$. We may assume $j \leq k - 2d + 3$, since otherwise we would have $h \geq j+1 \geq k-2d+5 \geq 2d-3$ and could apply the symmetric argument. Then, for $a = 1 + j \mod (d-1)$, $c_i^j \leq_P c_{i+a}^{j+d-1}$ but $c_i^h \not\leq_P c_{i+a}^{j+2d-3}$. Therefore, the contracted vertex $c_i^j c_i^h$ has at least $(j+2d-3) + 1 - (j+d-1) = d-1$ incident red edges to the chain C_{i+a}.

Suppose now that the contracted pair is c_i^j and c_y^h where $i \neq y$ and $0 \leq j \leq h \leq k$. Again, we may assume by symmetry that $j \leq k-d+1$. If $h > j$, then $c_y^h \not\leq_P c_i^{j+d-1}$ and the contracted vertex $c_y^h c_i^j$ has at least $d-1$ incident red edges to the chain C_i. If $h = j$, then $c_y^h c_i^j$ similarly has at least $2(d-2) \geq d-1$ incident red edges to the chains C_i and C_y. ◀

3 Upper bound for posets of width d

Complementing Proposition 2.4, we give the core upper estimate followed by its proof:

▶ **Theorem 3.1.** *A poset of width $d \geq 2$ has natural twin-width at most $8d-9$, and hence the symmetric (matrix) twin-width at most $8d-8$. The corresponding contraction sequence can be found in time $\mathcal{O}(dn^2)$ where n is the number of vertices of the poset.*

By Dilworth's theorem, a poset $P = (X, \leq_P)$ is of width d if and only if the ground set X can be partitioned into at most d *chains* (a chain is linearly ordered by $\leq_P$). Hence, from now on, we will consider a poset of width d with a fixed partition π of X into d (nonempty) chains, formally as a triple $P = (X, \leq_P, \pi)$ where $\pi = \{U_1, \dots, U_d\}$.

Our upper bound in Theorem 3.1 will use only a special type of contractions – of two consecutive vertices of the same chain of P. Since contractions inside a chain essentially preserve the chain partition of P, we will for simplicity refer to the new chain partition as to

π again. We shall thus work with the following special kind of red posets, which result from a chain-partitioned poset by our special contractions:

▶ **Definition 3.2** (red d-neighbourly poset). *Let $P_0 = (X_0, \leq_{P_0}, \pi)$ be a poset partitioned by π into d chains. A* neighbourly contraction *is a contraction of a vertex pair x_1, x_2 such that $x_1 \neq x_2$ belong to the same chain U of π and they are consecutive in this chain (i.e., no element of U is strictly between x_1 and x_2). Note, however, that for such a pair x_1, x_2 there could exist a vertex y in another chain of π such that $x_1 \lneq y \lneq x_2$.*

A tuple $P = (X, \leq_P, \pi, R)$ is called a red d-neighbourly poset *(shortly a* neighbourly poset*) if the red poset $(X, \leq_P, R)$ is obtained from P_0 by an arbitrary sequence of neighbourly contractions (we shortly say that P is a* contraction of P_0*).*

Roughly speaking, our proof of Theorem 3.1 is going to argue that, although some neighbourly contractions can create many new red edges, overall the number of red edges that can potentially be created by every neighbourly contraction is only proportional to the size of the poset. Therefore, one can always find a "good" neighbourly contraction. Later on in the contraction sequence, we also have to watch the number of previously created red edges which is straightforward. Of course, to fulfill Definition 2.2, we will have to eventually contract the remaining d single-vertex chains into one, but that part will be a trivial conclusion at the end of our proof.

When dealing with neighbourly posets such as $P = (X, \leq_P, \pi, R)$, we adopt some special notation. Consider a chain $U = (u_1, u_2, \ldots, u_m)$ ordered as $u_1 \leq_P \ldots, \leq_P u_m$. For a vertex $x \in X$ such that $x = u_a$ of U, we shall write x^{+i} for the vertex u_{a+i} and x^{-i} for the vertex u_{a-i} (of course, assuming $a+i \leq m$ or $a-i \geq 1$, respectively). Let x^+ and x^- be a shorthand for x^{+1} and x^{-1}.

We give a unified way of picturing neighbourly posets – a *chain diagram* (already seen in Figure 1), which is close to the traditional Hasse diagram of a poset, but not exactly the same.

▶ **Definition 3.3** (chain diagram). *Let P be a red d-neighbourly poset. Every chain of P is drawn as a vertical line, the red edges of R are drawn as* red bars *between pairs of the chains, and there is a* black bar *from a vertex u of a chain U to a vertex v of a chain $V \neq U$, if and only if v is the least vertex of V greater than u and u is the greatest vertex of U smaller than v.*[3] *A black bar is never drawn as horizontal and is implicitly directed up in the picture.*

3.1 Structure of black and red bars

While black bars of a neighbourly poset P are directed by $\leq_P$, that is we have a black bar (u, v) if $u \leq_P v$ as in Definition 3.3, red bars are by Definition 2.2 undirected. Nevertheless, we can assign a direction to a red bar $\{u, v\} \in R$ as follows.

▶ **Definition 3.4** (orienting the red bars). *Let the neighbourly poset $P = (X, \leq_P, \pi, R)$ be a contraction of an ordinary poset $P_0 = (X_0, \leq_{P_0}, \pi)$, and recall that we view $u \in X$ as the subset $u \subseteq X_0$ of the respective contracted vertices of P_0. Since $u \subseteq X_0$ belongs to one chain of P_0 by Definition 3.2, the minimum $min(u)$ of u is well-defined.*

We orient the red bar $\{u, v\} \in R$ as (u, v), from u to v, if $min(u) \leq_{P_0} x$ for some $x \in v$, but $min(v) \not\leq_{P_0} min(u)$. Though, if both directions (u, v) and (v, u) are assigned by this criterion (which is possible, e.g., when the minima of u and v are incomparable), then we choose (u, v) if the last contraction into u happened later than that into v.[4]

[3] Notice that since R contains only incomparable pairs, there cannot be a red and a black bar together between the same pair.

[4] The latter criterion of choosing between (u, v) and (v, u) is not really important; we introduce it only to "break the tie" in a deterministic way.

Observe that Definition 3.4 assigns exactly one orientation to each $\{u, v\} \in R$. As an informal explanation, a red bar $\{u, v\} \in R$ in P means that between the sets $u, v \subseteq X_0$ in P_0, the edges (and non-edges) are not uniform (not all in one direction), and then we choose the "prevailing direction" for the orientation of $\{u, v\}$. We shall write a red bar as $\{u, v\} \in R$ if we do not care about the orientation of it, and as $(u, v) \in R$ if we do care.

Now we summarize some (basically trivial) technical properties used in further proofs.

▶ **Lemma 3.5.** *Let $P = (X, \leq_P, \pi, R)$ be a d-neighbourly poset and U and V be two distinct chains of P determined by π.*

a) *If $u \leq_P v$ where $u \in U$ and $v \in V$, then there is no $i > 0$ such that $(u, v^{+i}) \in R$. Analogously, if $u \geq_P v$, then there is no $i > 0$ such that $(v^{-i}, u) \in R$.*

b) *If $(u, v) \in R$ where $u \in U$ and $v \in V$ such that $(u, v^+) \notin R$, then $u \leq_P v^+$ (informally, red bars oriented from u to vertices of V come in a consecutive strip "capped" by a black bar). An analogous claim symmetrically holds for red bars oriented towards u from V.*

c) *If $(u, v) \in R$ where $u \in U$ and $v \in V$, then there are no $i, j > 0$ such that $(u^{+i}, v^{-j}) \in R$ or $u^{+i} \leq_P v^{-j}$ (informally, no two red bars of the same orientation from U to V may "cross", and no black bar from U to V may be "crossed" by a red bar starting below it in U).*

d) *There are together at most $|U| + |V| - 1$ red bars from a vertex of U to a vertex of V.*

e) *Assume $u \in U$, $v \in V$ such that $(u, v) \in R$ is a red bar that has been created by a neighbourly contraction into u, and that no contraction into v has happened after the creation of red (u, v). Then no neighbourly contraction in P in the chain U can create another red bar oriented towards v.*

An analogous claim holds for $(v, u) \in R$ and creation of red bars oriented from v.

Proof. Let P be a contraction of the ordinary poset $P_0 = (X_0, \leq_{P_0}, \pi)$.

a) Trivially by transitivity, $u \leq_P v \leq_P v^{+i}$ contradicts that pairs in R are incomparable.

b) By $(u, v) \in R$ and Definition 3.4, for some $x \in v$ of P_0 we have $min(u) \leq_{P_0} x \leq_{P_0} min(v^+)$. Then, if $max(u) \not\leq_{P_0} min(v^+)$, we would have forbidden $(u, v^+) \in R$. Therefore, by homogeneity of the edges from u to v^+ in P_0, we get desired $u \leq_P v^+$.

c) Assume the contrary, that $(u^{+i}, v^{-j}) \in R$. Then, by Definition 3.4 and transitivity in P_0, $max(u) \leq_{P_0} min(u^{+i}) \leq_{P_0} x \leq_{P_0} min(v)$ where $x \in v^{-j}$, which contradicts the assumption $\{u, v\} \in R$. The same argument goes through if $u^{+i} \leq_P v^{-j}$ with $x = min(v^{-j})$.

d) We ignore the chains other than U or V, and we prove the statement by induction. The base case is $|U| = 1$; clearly, there may be at most $|V| = |U| + |V| - 1$ red bars going from U to V. Now, let u_1 be the minimum of the chain U, and assume that there are $q \geq 0$ red bars oriented from u_1 to V. Let $V_1 \subseteq V$ be the lowest $q - 1$ vertices of the chain V. By (c), only red bars from u_1 may end in V_1. So, we remove the vertices $\{u_1\} \cup V_1$ and, by induction, there are at most $|U| - 1 + |V| - (q - 1) - 1 = |U| + |V| - q - 1$ red bars from $U \setminus \{u_1\}$ to $V \setminus V_1$. With the q red bars starting in u_1, we get the desired bound.

e) If $(u, v) \in R$ has been created by a contraction into u, and not by a prior contraction into v, then $min(u) \leq_{P_0} min(v)$ using Definition 2.2. Hence a further neighbourly contraction in the chain U below u cannot at all create a red bar incident to v, and a neighbourly contraction in the chain U above u can only create a new red bar oriented from v, according to Definition 3.4. ◀

3.2 Minimizing the red potential

Now comes the core of the proof of Theorem 3.1, estimating how many red edges can potentially result from all possible neighbourly contractions in P. Let $u \in U$ be a vertex of a chain U which is not maximal, u^0 be the vertex created by the contraction of u and u^+, and define the *red potential of* u as the number of red edges incident to u^0 after the contraction (so, previous red edges incident to u or u^+ are also counted here). The *red potential of the chain* U is simply the sum of red potentials over the non-max vertices of U, and the *red potential of* P is the sum over all chains of P.

▶ **Lemma 3.6.**
a) *If $P = (X, \leq_P, \pi, R)$ is a red d-neighbourly poset with $m = |X|$ elements, then the red potential of P is at most $2(d-1)(m-d) + 2|R|$.*
b) *There are at most $|R| \leq 2(d-1)(m-d)$ red bars in P.*

Proof. In both parts of this proof, let U and V be any two distinct chains of P.

a) For any vertex $x \in U$ (resp. $x \in V$), let the *red supplement of* x, denoted $rs(x)$, be the number of neighbourly contractions in V (resp. in U) that create a red bar incident to x in the ordinary poset $P_0 = (X, \leq_P, \pi)$, i.e., if we ignore the red bars of P. Let $\rho_{U,V}$ be the sum of $rs(x)$ over all $x \in U \cup V$. Notice that $\rho_{U,V}$ is the red potential of P_0 restricted to $U \cup V$.

By Definition 3.4, $rs(x) \leq 2$, since at most one neighbourly contraction creates a new red bar oriented to x and at most one creates a new red bar oriented from x, and hence $\rho_{U,V} \leq 2(|U| + |V|)$. We slightly improve this immediate bound as follows. Let u_1 and v_1 be the minima of the chains U, V, and let u' and v' be their maxima. If $rs(u_1) + rs(v_1) > 2$ then, up to symmetry, $rs(u_1) = 2$. This would mean that there is a vertex $v \in V$ such that $v \leq u_1$. Since $v_1 \leq v$, we get $v_1 \leq u_1 \leq u$ for all $u \in U$, which implies that $rs(v_1) = 0$. Consequently, $rs(u_1) + rs(v_1) \leq 2$, and $rs(u') + rs(v') \leq 2$ by symmetry. Summing with the remaining vertices of $U \cup V$ we obtain an estimate $\rho_{U,V} \leq 2(|U| + |V|) - 4$.

Let the cardinalities of the d chains of P be $m_1, m_2, \ldots, m_d$. Summing the latter estimate over all $\binom{d}{2}$ unordered pairs of chains of P, we get

$$\sum_{1 \leq i < j \leq d} \rho_{U_i,U_j} = \sum_{1 \leq i < j \leq d} \left[2 \cdot (m_i + m_j) - 4 \right] = 2(d-1)m - 4 \cdot \binom{d}{2} = 2(d-1)(m-d)\,. \quad (1)$$

Now it remains to add the contribution of the red bars of R in P. Each red bar $(u, v) \in R$, where $u \in U$, contributes to the red potentials of at most two neighbourly contractions in U; namely to those of the pairs u^-, u and u, u^+. However, we are over-counting this way, and we now show that it is enough to count a "+1" contribution towards one of the two contractions. Precisely, we contribute a red bar (u, v) to the contraction of the pair u, u^+ since the following holds: If $(u^-, v) \in R$, then the contraction of u^-, u into u^0 anyway makes only one inherited red bar (u^0, v) (out of the two red bars (u, v) and (u^-, v) of P) and our rule contributes (u^-, v) to the pair u^-, u. Otherwise, by Lemma 3.5(b), we have $u^- \leq_P v$ and the inherited red bar (u^0, v) has already been counted as a new one in $rs(v)$ above.

Since each red bar in R increases the red potential in two chains, we in total get that the red potential of P is at most $2(d-1)(m-d) + 2|R|$.

b) A straightforward application of Lemma 3.5(d) gives an upper bound of $2(|U|+|V|)-2$ on the number of red bars between chains U and V (in both directions). However, if u_1, v_1 denote the minima of U, V, then $(u_1, v_1) \notin R$ up to symmetry. So, an application of Lemma 3.5(d) to $U \setminus \{u_1\}$ and V gives at most $|U| + |V| - 2$ red bars from U to V. A symmetric "saving" can be obtained for the maxima of U, V, and hence we can have at most $2(|U| + |V|) - 4$ red bars between chains U and V in both directions.

Finally, summing the latter bound over all pairs of chains as in (1), we get the desired bound $|R| \leq 2(d-1)(m-d)$. ◀

Proof of Theorem 3.1. Let $P_0 = (X_0, \leq_{P_0}, \pi)$ be an ordinary poset partitioned into d chains. We are now ready to finish the main proof; to find a desired contraction sequence of P_0 of bounded red degree. The natural idea at each step (with a neighbourly poset P obtained from P_0 so far) is to exhaustively find a neighbourly contraction in P of the smallest red potential, which is upper-bounded independently of the size of P based on Lemma 3.6.

Let $P = (X, \leq_P, \pi, R)$ be a contraction of P_0, as above, and $m = |X| \leq |X_0| = n$. The red potential of whole P is at most $2(d-1)(m-d) + 4(d-1)(m-d) = 6(d-1)(m-d)$ by Lemma 3.6. Since there are $m - d$ possible neighbourly contractions in P, one of them has the red potential at most $6(d-1)$. This of course makes sense only for $m > d$ but for $m \leq d$, we can finish the contraction sequence arbitrarily, without getting a high red degree.

We are nearly done, but there is a small catch. For a vertex $x \in X$ of P, call a red edge $\{x, y\}$ incident to x *domestic* (to x) if it has been there already the last time we have contracted into x along our contraction sequence; otherwise, call red $\{x, y\}$ *foreign*. While the argument in the previous paragraph bounded the number of domestic red edges incident to any vertex along the whole sequence, we have not yet bounded the number of potential foreign red edges (that is those which have been created by contraction to other vertices later on). Using Lemma 3.5(e), we claim that there can be at most one foreign edge incident to x oriented towards x, and one oriented from x, per each other chain of P. Hence the number of foreign red edges incident to any x is at most $2(d-1)$, and at most $2d-3$ if x has also some domestic red edges (again by Lemma 3.5(e)). Altogether, the maximum red degree along our contraction sequence is at most $6(d-1) + 2d - 3 = 8d - 9$.

The above proof straightforwardly translates into a simple and efficient algorithm. The red potential of one vertex of P can be found in time proportional to d and the value of this red potential, and hence the minimum red potential of the poset P in the current step of a contraction sequence is determined in time $\mathcal{O}(dm) \leq \mathcal{O}(dn)$. The same time is sufficient to update P for the next step. Since we need $n-1$ steps of the contraction sequence for P_0, this computation is finished in time $\mathcal{O}(dn^2)$. ◀

4 Tight estimate for posets of width 2

From Section 3 we get that in the worst case a poset of width 2 has twin-width at least 1 and at most 7 (where the upper bound of Theorem 3.1 can likely be improved a bit in this special case of $d = 2$). However, in order to get an exact worst-case value of twin-width, namely value 2, we had to employ a very different approach – one which is special only for posets of width 2 (unfortunately, this argument does not generalize even to width 3).

We start with the upper estimate:

▶ **Theorem 4.1.** *A poset of width 2 has natural twin-width at most 2, and the corresponding contraction sequence can be computed in linear time.*

We prove the statement by providing the claimed algorithm and proving its correctness. On a high level, our algorithm performs a depth-first search for a "safe" possibility of a neighbourly contraction in one of the two chains of a poset P, starting from a minimal vertex u_1 of the poset. By a safe neighbourly contraction we mean one in which we have or create at most one incoming and at most one outgoing red bar in the contracted vertex. We also stay in firm control of all red bars in intermediate contracted red posets.

To control the search for neighbourly contraction pairs and the occurrence of red bars, we introduce the notion of a *directed bar path* (recall also Definition 3.3 of a chain diagram with bars). A bar path in a red 2-neighbourly poset P is a directed path represented as a vertex sequence $(x_1, x_2, \ldots, x_k)$ in P such that

- x_1 is a minimal vertex of P (either of the possible two),
- for $i = 1, \ldots, k-1$, (x_i, x_{i+1}) is a black or red bar in P oriented this way, and
- if both x_i^+, x_{i+1}^+ exist in P, then $x_i^+ \not\leq_P x_{i+1}^+$.

Notice that a bar path is "zig-zag" switching between the two chains of P. Our bar-path controlled algorithm is then formalized in Algorithm 1.

Commenting on this algorithm, we remark that the main part (the one searching for a "safe" neighbourly contraction along bar path B) is presented on lines 15 to 26. The preceding supplementary part on lines 9 to 13 is there to prepare for a possible contraction at the root u_1 of bar path B; if u_1 is not the global minimum of P, then the lower vertices of the other chain V of P could create many red bars oriented towards u_1 after neighbourly contraction to u_1. This is eliminated by safe contractions of the problematic vertices of V to v_1 on line 13, which make future neighbourly contractions to u_1 safe as well.

Actually, the course of Algorithm 1 is illustrated in previous Figure 1. It is $u_1 = A$, $v_1 = E$ (after the supplementary first step, $v_1 = EF$), and the bar path leading to the first contraction (of C, D) on line 22 is $B = (A, G, C)$. Further on, for example, in the fourth picture (the top-right red poset) we have $B = (A)$ and $B^+ = (A, GH, CD, I)$. In the fifth picture (bottom-left), we get $B = (AB, GH)$ and $B^+ = (EF, AB, GH, CD, I)$, and so on.

Proof of Theorem 4.1. We refer to Algorithm 1 and the definition of bar path B. Let P_0 be the input ordinary poset and P the current 2-neighbourly poset, as in the algorithm. Let U_0, V_0 denote the two chains of P_0 and $u_1 = min(U_0)$ be the start (root) of bar path B which stays fixed during the course of computation. Let $v_1 = min(V_0)$. Note that we slightly abuse notation by referring to these vertices as to u_1 and v_1 also in the poset P, after possible contractions into u_1 or v_1.

For the purpose of analysis of Algorithm 1, we define an *extended bar path* $B^+ \supseteq B$ of the current bar path B in P as follows: B^+ starts with (v_1, u_1) if $(v_1, u_1) \in R$, and B^+ starts in u_1 otherwise. Then B^+ contains all bars of B in order, then possibly one black bar starting in the last vertex of B, and finally, B^+ ends as a directed path using only red bars of P. We claim the following *invariant* at the beginning and after every iteration of the loop from line 6:

(I) B conforms to the conditions of a bar path.

(II) There exists an extended bar path B^+ of B in P containing all red bars of P.

To better understand the role of an extended bar path, observe that, modulo renaming of contracted vertices, B^+ coincides with the former bar path in the iteration in which the upper-most red bar of B^+ has been created by a contraction.

Since an extended bar path is acyclic and does not repeat vertices – this is not trivial but follows from Definition 3.4 – condition (II) then implies that the red degree of P after every iteration is at most 2, thus proving the conclusion of Theorem 4.1. Our aim hence is to prove the invariant, by induction on the iterations of the main cycle.

At the beginning, $B = (u_1)$ satisfies (I) and (II) is trivial since $R = \emptyset$. We now assume that these hold when an iteration of the main cycle (line 6) starts. Possible contractions on line 13 do not create red edges or change B, except that when $|U| = 1$ they could create the red bar (v_1, u_1) which will be included in our extended bar path there.

For the next arguments, note that line 15 always selects a red bar (u, v) if there is one starting from u. So, if there is no red bar from u and $u^+ \leq_P v$, then (u, v) is *not* a black bar either (Definition 3.3) and the extended bar path B^+ from (II) before this iteration cannot reach v. Consequently, the contraction of u, u^+ on line 22 of this iteration does not make

Algorithm 1 Constructing a contraction sequence of red degree 2.

Input: Given a poset $P_0 = (X_0, \leq_{P_0}, \pi)$ partitioned into 2 chains.
Output: A record of the constructed contraction sequence of P_0 of red degree ≤ 2.

1: **declare** $P = (X, \leq_P, \pi, R)$ a red 2-neighbourly poset
2: **declare** B a directed bar path (of black and red bars of P)
3: $P \leftarrow$ red neighbourly poset $P = (X_0, \leq_{P_0}, \pi, \emptyset)$;
4: $B \leftarrow$ bar path (u_1), where u_1 is a minimal vertex of P_0 (any of possible two);
5: *// bar path B directs the course of the algorithm (kind of similarly to classical DFS); note that B stays rooted in $u_1 = min(U)$ till the end of the main loop*
6: **while** $B \neq \emptyset$ **do**
7: $\quad u \leftarrow$ the (upper) end vertex of bar path B, i.e., $B = (x_1, \dots, x_k, u)$;
8: $\quad \{U, V\} \leftarrow$ the two chains of P (by π) such that $u \in U$;
9: $\quad$ **if** $B = (u)$, which is equivalent to $u = min(U)$ **then**
10: $\quad\quad$ *// here we contract the lower section of V which is homogeneous towards U:*
11: $\quad\quad$ $v_1 \leftarrow min(V)$; $u_2 \leftarrow$ smallest $u_2 \in U$ such that $v_1 \leq_P u_2$, or u_2 nonexistent;
12: $\quad\quad$ **while** $|V| > 1$ and $\big(|U| = 1$ or (u_2 exists and $v_1^+ \leq_P u_2)\big)$ **do**
13: $\quad\quad\quad$ $P \leftarrow$ **contraction** of the pair v_1, v_1^+ in P; $v_1 \leftarrow v_1 v_1^+$;
14: $\quad$ *// the main part; prolonging B, or possibly contracting at the end and shortening B:*
15: $\quad$ $v \leftarrow$ the smallest vertex $v \in V$ such that $u \leq_P v$ or $(u, v) \in R$, or v nonexistent;
16: $\quad$ **if** v, v^+ and u^+ exist in P, and $u^+ \not\leq_P v^+$ **then**
17: $\quad\quad\quad$ *// note that since $u^+ \not\leq_P v^+$, we get that $(u, v) \in R$ or (u, v) is a black bar*
18: $\quad\quad$ $B \leftarrow (B, v)$, i.e., prolongation of bar path B by the bar (u, v);
19: $\quad$ **else**
20: $\quad\quad$ **if** u^+ exists in P **then**
21: $\quad\quad\quad$ *// the contraction here is safe; even if v, v^+ exist, we have $u^+ \leq_P v^+$ and no red bar towards v^+ (or further up) is created*
22: $\quad\quad\quad$ $P \leftarrow$ **contraction** of the pair u, u^+ in P; $u \leftarrow uu^+$;
23: $\quad\quad$ **if** $B = (u)$ **then**
24: $\quad\quad\quad$ **if** u^+ nonexistent in P **then** $B \leftarrow \emptyset$;
25: $\quad\quad$ **else**
26: $\quad\quad\quad$ $B \leftarrow B \setminus (u)$, i.e., shortening of bar path B by the last bar towards u;
27: $P \leftarrow$ **contraction** of the remaining pair $min(U), min(V)$ in P;

(u, v) red, and (I) and (II) are satisfied after the iteration with the same B^+. Hence we can assume that $(u, v) \in R$ or $u^+ \not\leq_P v$ hold in the coming case analysis.

We now analyze the remaining cases according to the “if” statements from line 16 onward:

- If v, v^+ and u^+ exist and $u^+ \not\leq_P v^+$, then (u, v) is a red or black bar in P, and so we explicitly satisfy all three conditions of a bar path for the prolongation (B, v) on line 18. No new red bars are created in this iteration, and we claim that an extended bar path B^+ from the previous iteration contains or will contain the new bar (u, v); this is trivial if (u, v) is red, and for a black bar (u, v) it is the only bar of P starting in u anyway.
- If u^+ and v exist, and v^+ is nonexistent or $u^+ \leq_P v^+$, then the contraction of u, u^+ on line 22 makes (u, v) red if it was not such before. No other red bar exists or is created from u (although, the bar of B towards u may already be red). Altogether, we inherit

B^+ from the previous step, and with (now) red (u, v) this will be (II) a valid extended bar path of the shortened bar path $B \setminus (u)$ at the end of the iteration. Not to forget (I), $B \setminus (u)$ will be a valid bar path as well.

- If $B = (u)$ in the previous case, we do not shorten B (since u^+ exists), but the arguments stay the same.
- if u^+ exists but v does not, then an extended bar path cannot reach beyond u. So, after the contraction and shortening B by u we will trivially satisfy (I) and (II) again.
- Lastly, if u^+ is nonexistent, then we only shorten B by u, or we are at the end of the algorithm if u is the root of B (note that here the procedure on lines 9 to 13 also takes part and shortens V to one vertex before we stop the cycle from line 6).

It remains to argue why the algorithm stops, and what is the runtime. The first part is clear since every iteration of the main loop either prolongs the current bar path (which is bounded), or eventually finds a next contraction pair.

As for the runtime, we use the following special representation of the working poset P, which extends the chain diagram of Definition 3.3: For every $u \in U$ we record, besides the red bars of u, the least $v \in V$ such that $u \leq_P v$. We input the poset P_0 as a traditional Hasse diagram, find its chain cover by standard means and prepare our special representation of it in linear time with respect to $|X_0|$. The total number of iterations of the main cycle is linearly proportional to the number of performed contractions in the main part. Then, at every iteration of the main cycle, we can perform the computation, and the update of the structure representing P, in constant time. The only exception is a possible iteration of the inner cycle on line 12, which is counted amortized towards the total number of contractions. ◀

The lower estimate matching Theorem 4.1 is as follows.

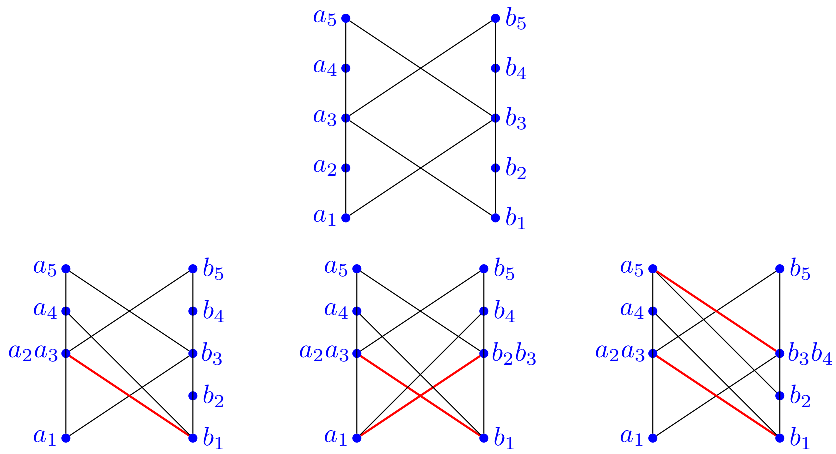

Figure 3 [top] A poset of width 2 which has natural twin-width (at least) 2. [bottom] All possible contractions of this poset, up to symmetry, which have red degree 1.

▶ **Proposition 4.2.** *The poset depicted in Figure 3 has natural twin-width at least* 2.

Proof. [5] Thanks to symmetries in the depicted poset P, it is routine and easy to verify that every contraction of a pair in P results in two red edges incident to the contracted vertex, except the contraction of the pair a_2, a_3 (or the symmetric pairs b_2, b_3 or b_3, b_4 or a_3, a_4). The result of this contraction, a red poset P_1, is on bottom left of Figure 3.

[5] To be safe, we have independently verified this whole proof by an exhaustive computer check.

Now, in P_1, the contraction of (now red) pair a_2a_3, b_1 creates three incident red edges. For every other contracted pair in P_1, we either get the same two red edges as if it was contracted in P, or another red edge incident to a_2a_3 or to b_1, or one of the further two possibilities of isolated red edges depicted at the bottom of Figure 3. In those cases, again by a boring but routinely easy case analysis, one can see that every contraction creates red degree at least two. ◀

5 Conclusions

We have proved a tight relation up to a multiplicative constant between the width of a poset and its twin-width. The constants in this linear relation are very reasonable, but they can likely still be improved. The result also gives a simple and practically usable approximation algorithm for the twin-width of posets of small width, while the same is not known in the more general cases of all graphs or all posets.

References

1. Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width III: max independent set and coloring. *CoRR*, abs/2007.14161, 2020. `arXiv:2007.14161`.
2. Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width II: small classes. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 1977–1996. SIAM, 2021. `doi:10.1137/1.9781611976465.118`.
3. Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, and Stéphan Thomassé. Twin-width IV: low complexity matrices. *CoRR*, abs/2102.03117, 2021. `arXiv:2102.03117`.
4. Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 601–612. IEEE, 2020. `doi:10.1109/FOCS46700.2020.00062`.
5. Édouard Bonnet, Jaroslav Nešetřil, Patrice Ossona de Mendez, Sebastian Siebertz, and Stéphan Thomassé. Twin-width and permutations. *CoRR*, abs/2102.06880, 2021. `arXiv:2102.06880`.
6. Simone Bova, Robert Ganian, and Stefan Szeider. Model checking existential logic on partially ordered sets. *ACM Trans. Comput. Log.*, 17(2):10:1–10:35, 2016. `doi:10.1145/2814937`.
7. Jakub Gajarský, Petr Hliněný, Daniel Lokshtanov, Jan Obdržálek, Sebastian Ordyniak, M. S. Ramanujan, and Saket Saurabh. FO model checking on posets of bounded width. In *FOCS*, pages 963–974. IEEE Computer Society, 2015. `doi:10.1109/FOCS.2015.63`.
8. Robert Ganian, Petr Hliněný, Daniel Král, Jan Obdržálek, Jarett Schwartz, and Jakub Teska. FO model checking of interval graphs. *Log. Methods Comput. Sci.*, 11(4), 2015. `doi:10.2168/LMCS-11(4:11)2015`.
9. Petr Hliněný, Filip Pokrývka, and Bodhayan Roy. FO model checking on geometric graphs. *Comput. Geom.*, 78:1–19, 2019. `doi:10.1016/j.comgeo.2018.10.001`.