Tree Decompositions Meet Induced Matchings: Beyond Max Weight Independent Set

Paloma T. Lima
IT University of Copenhagen, Denmark

Martin Milanič
FAMNIT and IAM, University of Primorska, Koper, Slovenia

Peter Muršič
FAMNIT, University of Primorska, Koper, Slovenia

Karolina Okrasa
Warsaw University of Technology, Poland

Paweł Rzążewski
Warsaw University of Technology, Poland
University of Warsaw, Poland

Kenny Štorgel
Faculty of Information Studies, Novo mesto, Slovenia
FAMNIT, University of Primorska, Koper, Slovenia

Abstract

For a tree decomposition $\mathcal{T}$ of a graph G, by $\mu(\mathcal{T})$ we denote the size of a largest induced matching in G all of whose edges intersect one bag of $\mathcal{T}$. The *induced matching treewidth* of a graph G is the minimum value of $\mu(\mathcal{T})$ over all tree decompositions $\mathcal{T}$ of G. Yolov [SODA 2018] proved that for graphs of bounded induced matching treewidth, tree decompositions with bounded $\mu(\mathcal{T})$ can be computed in polynomial time and Max Weight Independent Set can be solved in polynomial time.

In this paper we explore what other problems are tractable in such classes of graphs. As our main result, we give a polynomial-time algorithm for Min Weight Feedback Vertex Set. We also provide some positive results concerning packing induced subgraphs, which in particular imply a PTAS for the problem of finding a largest induced subgraph of bounded treewidth.

These results suggest that in graphs of bounded induced matching treewidth, one could find in polynomial time a maximum-weight induced subgraph of bounded treewidth satisfying a given CMSO_2 formula. We conjecture that such a result indeed holds and prove it for graphs of bounded tree-independence number, which form a rich and important family of subclasses of graphs of bounded induced matching treewidth.

We complement these algorithmic results with a number of complexity and structural results concerning induced matching treewidth, including a linear relation to treewidth for graphs with bounded degree.

2012 ACM Subject Classification Mathematics of computing → Graph algorithms; Theory of computation → Graph algorithms analysis

Keywords and phrases induced matching treewidth, tree-independence number, feedback vertex set, induced packing, algorithmic meta-theorem

Digital Object Identifier 10.4230/LIPIcs.ESA.2024.85

Related Version *Full Version*: https://arxiv.org/abs/2402.15834 [44]

Funding *Paloma T. Lima*: Supported by the Independent Research Fund Denmark grant agreement number 2098-00012B.

Martin Milanič: Supported in part by the Slovenian Research and Innovation Agency (I0-0035, research program P1-0285 and research projects J1-3001, J1-3002, J1-3003, J1-4008, J1-4084, and

© Paloma T. Lima, Martin Milanič, Peter Muršič, Karolina Okrasa, Paweł Rzążewski, and Kenny Štorgel;
licensed under Creative Commons License CC-BY 4.0

N1-0370) and by the research program CogniCom (0013103) at the University of Primorska.
Peter Muršič: Supported in part by the Slovenian Research and Innovation Agency (J1-4008).
Karolina Okrasa: Supported by the Polish National Science Centre grant no. 2021/41/N/ST6/01507.
Paweł Rzążewski: Supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme grant agreement number 948057.
Kenny Štorgel: Supported in part by the Slovenian Research and Innovation Agency (research program P1-0383 and research projects J1-3002 and J1-4008).

1 Introduction

Structured graph decompositions and the corresponding graph width parameters have become one of the central tools for dealing with algorithmically hard graph problems. One of the best known and well-studied graph width parameters is *treewidth*, which was introduced independently by several groups of authors with various motivations [7, 13, 36, 55]. Roughly speaking, treewidth measures how similar the graph is to a tree. A graph G with small treewidth can be represented by a tree T whose every node holds a small subset (called *bag*) of vertices of $V(G)$, such that the connectivity properties of G are reflected in the local structure of T (and the bags). A tree decomposition $\mathcal{T}$ consists of the tree T and a function that assigns a bag to each node of T, and the *width* of $\mathcal{T}$ is the size of its largest bag minus 1 (this is by convention). The *treewidth* $\mathrm{tw}(G)$ is the minimum width of a tree decomposition of G.

Having a tree decomposition of small width is excellent for algorithmic applications, as we can mimic the standard bottom-up dynamic programming on trees. For example, in order to solve Max Weight Independent Set (MWIS), the dynamic programming table for each node of T should be indexed by all possible independent sets in the corresponding bag. Consequently, if each bag is *small*, say, its size is bounded by a constant k, the running time we obtain is $2^k \cdot (n + |V(T)|)^{\mathcal{O}(1)}$ where $n \coloneqq |V(G)|$. As $|V(T)|$ can be assumed to be linear in n, this yields a polynomial-time algorithm for MWIS when restricted to graphs of bounded treewidth.

Such an approach has found numerous applications in solving classic NP-hard problems on graphs with bounded treewidth; see the handbook by Cygan et al. [24, Section 7]. The richness of the family of problems that can be solved by a dynamic programming over a tree decomposition is witnessed by the celebrated meta-theorem by Courcelle [23]. It asserts that each problem expressible in *Counting Monadic Second Order Logic* (CMSO_2)[1] can be solved in polynomial time (actually, in linear time) for graphs of bounded treewidth.

An astute reader might notice a caveat in the results mentioned above: they assume that the input graph is given with a corresponding decomposition. This might be a serious obstacle, as computing an optimal tree decomposition is NP-hard [7, 15]. However, if treewidth is bounded by a constant, an optimal decomposition can be found in polynomial time [14, 41]. Actually, in the context discussed above even an approximation of an optimal decomposition is sufficient. Luckily, such approximation algorithms are known not only for treewidth [16, 40, 56], but also for other parameters discussed in the paper [28, 60]. Thus, we will implicitly assume that the instance graphs are always provided with a corresponding tree decomposition.

[1] In this logic one can use vertex, edge, and (vertex or edge) set variables, check vertex-edge incidence, quantify over variables, and apply counting predicates modulo fixed integers. For a formal introduction, see the full version of the paper [44].

Let us go back to the MWIS problem, which serves as the starting point of our investigations. Notice that in the argument sketched above we do not really need the fact that the treewidth is bounded by a constant. Indeed, if the size of each bag is bounded by a *logarithmic function of* n, then the number of all independent sets in a single bag is bounded by $2^{\mathcal{O}(\log n)} = n^{\mathcal{O}(1)}$, i.e., by a polynomial, which still yields a polynomial-time algorithm for MWIS. This is one of the motivations to study classes of graphs with *logarithmic treewidth*, which has been a very active topic in structural graph theory in recent years [3, 18].

However, the crux of the algorithm above is not really the size of the bag, but the number of independent sets inside each bag. Consider, for example, the class of chordal graphs, i.e., graphs that do not contain an induced cycle with at least four vertices. Equivalently, these are graphs that admit a tree decomposition whose every bag is a clique. Even though these cliques can be arbitrarily large, they still contain very few (that is, polynomially many) independent sets and thus such a decomposition can be used to solve MWIS in polynomial time.

This observation is heavily extended by the framework of *potential maximal cliques* by Bouchitté and Todinca [19]. Essentially, there we work with an implicitly given tree decomposition whose every bag is promised to contain only a constant number of vertices from the optimal independent set. This approach allows us, for example, to solve MWIS and many related problems for graph classes with polynomially many minimal separators [19, 32]. With some care, it can also be used for other graph classes, e.g., graphs that exclude a fixed induced path [35, 45] or long induced cycles [4].

Another, more direct way of generalizing bounded-treewidth graphs and chordal graphs was suggested by Yolov [60] and later, independently, by Dallard, Milanič, and Štorgel [27]. We use the notation and terminology of Dallard et al. [27]. Given a tree decomposition $\mathcal{T}$ of a graph G, by $\alpha(\mathcal{T})$ we denote the size of a largest independent set contained in a bag of $\mathcal{T}$. The *tree-independence number* of G, denoted by $\mathsf{tree}\text{-}\alpha(G)$, is the minimum value of $\alpha(\mathcal{T})$ over all tree decompositions $\mathcal{T}$ of G. Note that given a tree decomposition $\mathcal{T}$ of G with $\alpha(\mathcal{T}) \leqslant k$, we can solve MWIS on G in time $n^{\mathcal{O}(k)}$, which is polynomial for constant k. Let us remark that we always have $\mathsf{tree}\text{-}\alpha(G) \leqslant \mathrm{tw}(G)+1$ and if G is chordal, then $\mathsf{tree}\text{-}\alpha(G) \leqslant 1$ (actually, the reverse implication also holds [27]). Tree-independence number proved to be a fruitful topic in recent years and attracted some attention, both with structural [1, 25] and algorithmic [27, 28] motivations.

In a recent work, Dallard, Fomin, Golovach, Korhonen, and Milanič [28] showed that tree-independence number is in a sense the most general "natural" parameter of a tree decomposition whose boundedness yields a polynomial-time algorithm for MWIS. More precisely, let γ be any graph invariant satisfying $\gamma(G-v) \leqslant \gamma(G)$ for any graph G and any $v \in V(G)$. For a tree decomposition $\mathcal{T}$ of a graph G, define $\gamma(\mathcal{T})$ as the maximum value of $\gamma(B)$ taken over all subgraphs B induced by a single bag of $\mathcal{T}$. By $\gamma_{\mathrm{tw}}(G)$ we denote the minimum value of $\gamma(\mathcal{T})$ over all tree decompositions $\mathcal{T}$ of G.[2] Note that if $\gamma(G) = |V(G)|-1$, then γ_{tw} is precisely the treewidth, and if $\gamma(G)$ is the size of a largest independent set in G, then γ_{tw} is the tree-independence number. Dallard et al. [28] showed that for every invariant γ as above, either $\gamma_{\mathrm{tw}}(G)$ is upper-bounded by a function of $\mathsf{tree}\text{-}\alpha(G)$ (i.e., boundedness of $\gamma_{\mathrm{tw}}(G)$ implies boundedness of $\mathsf{tree}\text{-}\alpha(G)$), or MWIS remains NP-hard even for graphs G with constant $\gamma_{\mathrm{tw}}(G)$.

[2] While not immediately relevant for our paper, it may be worth mentioning that such tree decomposition based parameters can also be defined and studied in the more general context of hypergraphs (see [5, 47, 60]).

Six years earlier, Yolov [60] (obviously unaware of the work of Dallard et al. [28]) had defined another, more general parameter of a tree decomposition that can still be used to solve MWIS. There is no contradiction here – Yolov's parameter does not fall into the category of "natural" parameters considered by Dallard et al. In the original paper of Yolov [60], the parameter in question is called *minor-matching hypertree width* and is defined in a much more general setting of hypergraphs. As such a general definition is not relevant to our work, let us focus on the case of graphs; here we will call this parameter *induced matching treewidth*. For a tree decomposition $\mathcal{T}$ of G, by $\mu(\mathcal{T})$ let us denote the size of a largest induced matching in G all of whose edges intersect a single bag of $\mathcal{T}$ (we say that such a matching *touches* this bag). Now, the induced matching treewidth of G, denoted by $\mathsf{tree}\text{-}\mu(G)$, is the minimum value of $\mu(\mathcal{T})$ over all tree decompositions $\mathcal{T}$ of G. Note that $\mu(\mathcal{T})$ does not depend only on subgraphs induced by single bags, as some edges of an induced matching defining $\mu(\mathcal{T})$ might have one endpoint outside the bag.

It follows immediately from the definitions that for every graph it holds that $\mathsf{tree}\text{-}\mu(G) \leqslant \mathsf{tree}\text{-}\alpha(G)$. On the other hand, boundedness of induced matching treewidth does not imply boundedness of tree-independence number: for a biclique $K_{n,n}$ with each part of size n we have $\mathsf{tree}\text{-}\mu(K_{n,n}) = 1$ but, as observed by Dallard et al. [27], $\mathsf{tree}\text{-}\alpha(K_{n,n}) = n$. Thus indeed, induced matching treewidth is a more general parameter.

It is also important to note that the families of graph classes with bounded induced matching treewidth and classes of graphs with polynomially many minimal separators are incomparable. For example, the class of all graphs G_k consisting of k internally disjoint paths of length three with the same endpoints has bounded treewidth and hence bounded induced matching treewidth, but graphs in the class have exponentially many minimal separators (it is not difficult to see that the graph G_k has at least 2^k minimal separators). On other other hand, P_4-free graphs have a polynomial number of minimal separators (see [50]) but unbounded induced matching treewidth. This follows from a construction that we present in the full version of the paper [44].

At first it is not clear why boundedness of induced matching treewidth is helpful in solving MWIS. However, the following structural result by Yolov [60] is the key to understanding the connection between these two notions.

▶ **Lemma 1** (Yolov [60]). *Let k be a fixed integer. For an n-vertex graph G, a tree decomposition $\mathcal{T}$ of G with $\mu(\mathcal{T}) \leqslant k$, and a node t of $\mathcal{T}$, in time $n^{\mathcal{O}(k)}$ we can enumerate a set $\mathbb{I}_t$ with the following property: For every maximal independent set I, its intersection with the bag associated with t is in $\mathbb{I}_t$.*

Since every *maximum* independent set is in particular *maximal*, this immediately implies that the number of ways an optimal solution might intersect each bag of $\mathcal{T}$ is polynomial in n, if k is a constant. Using also the fact that given a graph with induced matching treewidth at most k, in polynomial time we can compute a tree decomposition such that $\mu(\mathcal{T}) = \mathcal{O}(k)$ (see Theorem 13), we obtain the following algorithmic result as a consequence.

▶ **Theorem 2.** *For every fixed k, the MWIS problem on n-vertex graphs G with $\mathsf{tree}\text{-}\mu(G) \leqslant k$ can be solved in time $n^{\mathcal{O}(k)}$.*

The tools developed by Yolov [60] can also be used to solve other problems that boil down to finding a constant number of (maximal) independent sets, like r-Coloring or finding a homomorphism to a fixed target graph. Even though this was not stated explicitly in [60], one can also obtain a polynomial-time algorithm for the problem of finding a largest induced r-colorable subgraph. For $r = 2$ this problem is equivalent (by taking the complement of the solution) to Min Odd Cycle Transversal.

The main goal of this paper is to explore what other problems that do not fall into the above category can be solved in polynomial time in classes of bounded induced matching treewidth. Recall that for bounded-treewidth graphs, a rich family of tractable problems is provided by the meta-theorem of Courcelle [23]. For graphs G with polynomially many minimal separators, where the framework of potential maximal cliques can be applied, a somewhat similar general result is provided by Fomin, Todinca, and Villanger [32]. For fixed integer r and fixed CMSO_2 formula ψ, by (r, ψ)-MWIS we denote the following computational problem (here "MWIS" stands for MAX WEIGHT INDUCED SUBGRAPH).

(r, ψ)-MWIS

Input: A graph G equipped with a weight function $\mathfrak{w}\colon V(G) \to \mathbb{Q}_+$.

Task: Find a set $F \subseteq V(G)$, such that
- $G[F] \models \psi$,
- $\mathrm{tw}(G[F]) \leqslant r$,
- F is of maximum weight subject to the conditions above,

or conclude that no such set exists.

Fomin, Todinca, and Villanger [32] proved that for each fixed r and ψ, the (r, ψ)-MWIS problem can be solved in time polynomial in the size of the input graph G and the number of minimal separators in G. Thus the running time is polynomial when restricted to classes of graphs with polynomial number of minimal separators.

Algorithms for (r, ψ)-MWIS are also known for graphs excluding a fixed induced path, or graphs excluding long induced cycles [4, 22, 33]. We conjecture that (r, ψ)-MWIS can also be solved in polynomial time for graphs of bounded induced matching treewidth.

▶ **Conjecture 3.** *For every fixed k, r and a CMSO_2 formula ψ, the (r, ψ)-MWIS problem can be solved in polynomial time for graphs with induced matching treewidth at most k.*

Even though we are not (yet) able to prove Conjecture 3, we provide some substantial evidence by approaching it from three different directions.

2 Overview of our results

The paper contains three main algorithmic contributions, as well as an initial set of results regarding induced matching treewidth, including computational results and bounds. We give an overview of these four sets of results and their implications in the following subsections. Due to space limits, all proofs of results listed above can be found in the full version of the paper [44].

2.1 Solving Max Weight Induced Forest

As our first and main result, we show that in graphs of bounded induced matching treewidth, one can in polynomial time find an induced forest of maximum possible weight.

▶ **Theorem 4.** *For every fixed k, the MAX WEIGHT INDUCED FOREST problem on n-vertex graphs G with $\mathsf{tree}\text{-}\mu(G) \leqslant k$ can be solved in time $n^{\mathcal{O}(k)}$.*

Note that MAX WEIGHT INDUCED FOREST is equivalent to finding a maximum-weight induced subgraph of treewidth at most 1, i.e., is a special case of (r, ψ)-MWIS for $r = 1$ and ψ being any formula satisfied by all graphs. Furthermore, by complementation, MAX WEIGHT INDUCED FOREST is equivalent to the well-studied MIN WEIGHT FEEDBACK VERTEX SET problem.

Before sketching the proof of Theorem 4, let us recall the textbook algorithm for Max Weight Induced Forest for graphs of bounded treewidth (see [24]). Let $\mathcal{T}$ be a tree decomposition of the instance graph, t be a node of $\mathcal{T}$, and X_t be the bag corresponding to t. For each set $Z \subseteq X_t$ and each partition π of Z, we keep the maximum weight of an induced forest F contained in the subgraph of G induced by the bags of the subtree of $\mathcal{T}$ rooted at t, such that:

- the intersection of $V(F)$ with X_t is exactly Z, and
- the partition π corresponds to the connected components of F.

For any induced forest F in G, we call a pair (Z, π) as above the *signature of F at t*. Now, processing $\mathcal{T}$ in a bottom-up fashion, we can find a maximum-weight induced forest in G using the information stored for each node.

The key insight leading to the proof of Theorem 4 is the following structural lemma.

▶ **Lemma 5.** *Let k be a fixed integer. For an n-vertex graph G, a tree decomposition $\mathcal{T}$ of G with $\mu(\mathcal{T}) \leqslant k$, and a node t of $\mathcal{T}$, in time $n^{\mathcal{O}(k)}$ we can enumerate a set $\mathbb{F}_t$ that contains the signature at t of every maximal induced forest in G.*

Recall that an analogous result for MWIS was shown by Yolov [60] (see Lemma 1): for each bag X_t corresponding to a node t of a tree decomposition $\mathcal{T}$ of bounded $\mu(\mathcal{T})$, in polynomial time we can enumerate a family $\mathbb{I}_t$ that contains an intersection of each maximal independent set in G with X_t.

As every maximum induced forest is in particular inclusion-wise maximal, Lemma 5 allows us to solve Max Weight Induced Forest by essentially mimicking the textbook algorithm for bounded-treewidth graphs. We just need to find the set $\mathbb{F}_t$ for each node t of $\mathcal{T}$, and instead of indexing the dynamic programming table by *all* possible pairs (Z, π), we use just the pairs that are in $\mathbb{F}_t$. This yields the running time claimed in Theorem 4.

So let us sketch the proof of Lemma 5.

What lives in a forest. Let F be a maximal induced forest in G; we think of it as an (unknown) optimal solution. Consider a node t of $\mathcal{T}$ and its corresponding bag X_t. The difficulty that we need to face is that the intersection of F with X_t might be arbitrarily large; for example, a component of F might be a large induced star contained in X_t. Thus we need to find some compact way to encode such possible intersections.

Let $\skull(F)$ be the *skeleton of F*, i.e., the set of vertices of degree at least two in F together with an arbitrary but fixed vertex v_C of each two-vertex component C of F. By $L(F)$ and $T(F)$ we denote, respectively, the sets of vertices of degree 1 (*leaves*) and of degree 0 (*trivial* vertices) of F, with the exception that for each two-vertex component C of F we include in $L(F)$ exactly one vertex from C, namely the vertex of C different from v_C. Clearly, the sets $\skull(F), L(F), T(F)$ form a partition of the vertex set of F.

Finding the skeleton. The first important observation is that for every node t of $\mathcal{T}$, the set $S \coloneqq \skull(F) \cap X_t$ cannot be too large in terms of k. Let us sketch the argument. Note that S induces a forest in G. First, each vertex v of degree at most 1 in $G[S]$, since it belongs to $\skull(F)$, must be adjacent to some vertex v' in $L(F)$ (not necessarily inside X_t), and the edges of the form vv' induce a matching whose every edge intersects X_t. Thus there can be at most k vertices of degree at most 1 in $G[S]$. Consequently, $G[S]$ can have at most k vertices of degree at least 3. Finally, $G[S]$ cannot contain long induced paths, as again, an induced path with $3k - 1$ vertices contains an induced matching with k edges. Altogether this allows us to bound the number of vertices of $G[S]$ by a function of k. A more careful analysis shows that $|S| \leqslant 8k$. Thus we can directly enumerate all candidates for $\skull(F) \cap X_t$.

Finding the the remaining vertices and detecting impostors. So we are left with finding an encoding of the set $(L(F) \cup T(F)) \cap X_t$. Observe that $\dddot{\text{🍃}}(F) \coloneqq L(F) \cup T(F)$ is an independent set in G. However, this is not very useful, as $L(F) \cup T(F)$ does not have to be maximal and we only have some information about the intersections of maximal independent sets with X_t, namely, by Lemma 1, they have to be in $\mathbb{I}_t$, which is of polynomial size. Let $I^*(F)$ be some maximal independent set of G containing $\dddot{\text{🍃}}(F)$; we know that its intersection I with X_t can only be chosen in a polynomial number of ways. However, I contains some vertices that are not in the solution (called *impostors*) and we should not consider their weight when choosing an optimal solution. Why does a vertex $v \in I$ not belong to $\dddot{\text{🍃}}(F)$? As F is a maximal forest, there are two possible reasons why v is an impostor:

- $v \in \text{💀}(F)$, or
- adding v to F would create a cycle.

Note that the vertices of the first type can be easily recognized as $S = \text{💀}(F) \cap X_t$ is directly represented. On the other hand, each vertex v of the second type is adjacent to at least two vertices from some connected component of F. Furthermore, we know that these two vertices are in $\text{💀}(F)$, as $\dddot{\text{🍃}}(F) \cup \{v\} \subseteq I^*(F)$ and $I^*(F)$ is independent. If these two neighbors of v in $\text{💀}(F)$ are in X_t, we are done, as we can guess the set $S = \text{💀}(F) \cap X_t$. However, the neighbors may lie outside the bag X_t.

Summing up, the type of a vertex $v \in I \setminus S$ (which may or may not belong to $\dddot{\text{🍃}}(F)$) is determined by its neighborhood in $\widehat{S} \coloneqq \text{💀}(F) \cap N[X_t]$, where by $N[X_t]$ we denote the set consisting of the vertices from X_t and their neighbors. If v has no neighbors in $\widehat{S}$, then v is trivial, i.e., $v \in T(F)$. If v has one neighbor in $\widehat{S}$, then v is a leaf, i.e., $v \in L(F)$. Finally, if v has at least two neighbors in $\widehat{S}$, then v is an impostor, i.e., $v \in I \setminus \dddot{\text{🍃}}(F)$. If we could store the information about the set $\widehat{S}$ explicitly, we could easily distinguish between these three types of vertices. Unfortunately, the set $\widehat{S}$ might still be arbitrarily large. So we need to make one more step: we observe that an inclusion-wise minimal subset Q of $\widehat{S}$ that still can distinguish between the three types of vertices in $I \setminus S$ is of size bounded by a function of k. In particular, we show that $|Q| \leqslant 4k$. Thus we can guess every candidate for Q.

Enumerating possible signatures. From the argument above it follows that the triple (S, I, Q) can be exhaustively guessed in a polynomial number of ways. Furthermore, thanks to the properties of Q, we can uniquely extract $L(F) \cap X_t$ and $T(F) \cap X_t$ from I, distinguishing vertices by their number of neigbors in Q. This way we can compute Z, i.e., the intersection of F with X_t.

In order to obtain π, we can exhaustively guess the partition $\widehat{\pi}$ of $S \cup Q$ that corresponds to the components of F restricted to the subgraph of G induced by the bags in the subtree rooted at t. Here we use the fact that the size of $S \cup Q$ is bounded by a function of k, i.e., a constant. From $\widehat{\pi}$ we can uniquely reconstruct π, as each vertex in $L(F) \cap X_t$ has a neighbor in $S \cup Q$, and each vertex in $T(F)$ forms a separate component. This completes the sketch of proof of Lemma 5.

2.2 Independent packings of connected subgraphs

Let $\mathcal{H} = \{H_j\}_{j \in J}$ be a family of connected subgraphs of a graph G. By $G^\circ[\mathcal{H}]$ we denote the graph with vertex set J, where two distinct vertices $j, j' \in J$ are adjacent if H_j and $H_{j'}$ have a vertex in common or there is an edge with one endpoint in H_j and the other in

$H_{j'}$.[3] Such a construction was considered by Cameron and Hell in [21], who focused on the particular case when $\mathcal{H}$ is the set of all subgraphs of G isomorphic to a member of a fixed family $\mathcal{F}$ of connected graphs; they showed that if G is a chordal graph, then so is $G^\circ[\mathcal{H}]$. Dallard et al. [27] generalized this result by showing that $\mathsf{tree}\text{-}\alpha(G^\circ[\mathcal{H}]) \leqslant \mathsf{tree}\text{-}\alpha(G)$. Other graph classes closed under taking $G^\circ[\mathcal{H}]$ are classes of graphs excluding long induced paths or long induced cycles; see Gartland et al. [34]. We show that an analogous statement holds for graphs with bounded induced matching treewidth.

▶ **Lemma 6.** *Let G be a graph and let $\mathcal{H}$ be a set of connected non-null subgraphs of G. Then $\mathsf{tree}\text{-}\mu(G^\circ[\mathcal{H}]) \leqslant \mathsf{tree}\text{-}\mu(G)$.*

In order to prove Lemma 6, we start with a tree decomposition $\mathcal{T}$ of G of small induced matching treewidth. Now we modify $\mathcal{T}$ into a tree decomposition of $G^\circ[\mathcal{H}]$ as follows. The shape of the tree is not altered, and for every node t of $\mathcal{T}$, its associated bag becomes $\{j \in J \mid V(H_j) \cap X_t \neq \emptyset\}$, where X_t is the bag assiociated with t in $\mathcal{T}$. It is not hard to verify that any induced matching touching a single bag of the modified tree decomposition corresponds to an induced matching of the same size that touches the corresponding bag of $\mathcal{T}$. Let us remark that if every graph in $\mathcal{H}$ has at least two vertices, this even yields a stronger result that $\mathsf{tree}\text{-}\alpha(G^\circ[\mathcal{H}]) \leqslant \mathsf{tree}\text{-}\mu(G)$.

Lemma 6 yields several algorithmic corollaries.

Independent packings of small connected subgraphs. First, notice that the MWIS problem on $G^\circ[\mathcal{H}]$ (with some weight function defined on J) corresponds to the problem of packing pairwise disjoint and nonadjacent graphs from $\mathcal{H}$, called Max Weight Independent Packing. Thus, combining Lemma 6 with Theorem 2, we show that Max Weight Independent Packing in classes of bounded induced matching treewidth can be solved in time polynomial in the size of G and $\mathcal{H}$.

▶ **Theorem 7.** *Let k be a fixed constant. Given a graph G of induced matching treewidth at most k and a finite family $\mathcal{H} = \{H_j\}_{j\in J}$ of connected non-null subgraphs of G, and a weight function $\mathfrak{w}\colon J \to \mathbb{Q}_+$, the Max Weight Independent Packing problem can be solved in time polynomial in $|V(G)|$ and $|\mathcal{H}|$.*

When $\mathcal{H}$ is the set of all subgraphs of G isomorphic to a member of a fixed family $\mathcal{F}$ of connected graphs, then Max Weight Independent Packing coincides with the Max Weight Independent $\mathcal{F}$-Packing problem studied in Dallard et al. [27]. This latter problem in turn generalizes several problems studied in the literature:

- the MWIS problem, which corresponds to the case $\mathcal{F} = \{K_1\}$,
- the Max Weight Induced Matching problem (see, e.g., [6, 53]), which corresponds to the case $\mathcal{F} = \{K_2\}$,
- the Dissociation Set problem (see, e.g., [51, 59, 61]), which corresponds to the case when $\mathcal{F} = \{K_1, K_2\}$ and the weight function assigns to each subgraph H_j the weight equal to $|V(H_j)|$,
- the k-Separator problem (see, e.g., [9, 42, 43]), which corresponds to the case when $\mathcal{F}$ contains all connected graphs with at most k vertices, the graph G is equipped with a vertex weight function $\mathfrak{w} : V(G) \to \mathbb{Q}_+$, and the weight function on $\mathcal{H}$ assigns to each subgraph H_j the weight equal to $\sum_{x\in V(H_j)} \mathfrak{w}(x)$.

[3] This extends the notation G° used in [33] for the graph $G^\circ[\mathcal{H}]$ in the particular case when $\mathcal{H}$ is the set of all non-null connected subgraphs of G.

Independent packings of sparse connected subgraphs. Observe that Max Weight Induced Forest for a weighted graph G is equivalent to Max Weight Independent Packing for an instance G, $\mathcal{H}$ (and $\mathfrak{w}$), where $\mathcal{H}$ consists of all induced subtrees of G (and the weight of H_j is the total weight of the vertices in H_j). Despite that, we note that Theorem 7 cannot be used to solve Max Weight Induced Forest in polynomial time. Indeed, the size of $\mathcal{H}$ is not bounded by a polynomial function of $|V(G)|$.

However, Lemma 6 might still be of use in this context. As already observed by Gartland et al. [34], we can sacrifice a small part of the optimal solution to break it into constant-size parts, which can then be handled efficiently. Indeed, for each $\epsilon > 0$ and any forest F we can remove an ϵ-fraction of the vertices of F so that the remaining components are of constant size (where this constant depends on ϵ but not on F). This approach can be used to obtain a PTAS for Max Induced Forest (the unweighted variant of Max Weight Induced Forest).

Actually, this reasoning goes far beyond trees and can be used to pack large induced subgraphs from any *weakly hyperfinite* class. A class $\mathcal{C}$ of graphs is *weakly hyperfinite* if for every $\epsilon > 0$ there is $c(\epsilon) \in \mathbb{N}$ such that in every graph $G \in \mathcal{C}$ there is a subset $X \subseteq V(G)$ of at least $(1-\epsilon)|V(G)|$ vertices such that every connected component of $G[X]$ has at most $c(\epsilon)$ vertices [49, Section 16.2]. It turns out that every class that is closed under vertex and edge deletions and admits sublinear balanced separators is weakly hyperfinite. Many well-known classes of sparse graphs are weakly hyperfinite, e.g., graphs of bounded treewidth, planar graphs and, more generally, graphs of bounded genus. In fact, all proper minor-closed classes are weakly hyperfinite. As graphs of treewidth bounded by a fixed constant form a minor-closed class, we obtain the following result which provides further evidence for Conjecture 3.

▶ **Theorem 8.** *For every fixed $k, r \in \mathbb{N}$ and $\epsilon > 0$, given a graph G with induced matching treewidth at most k, in polynomial time we can find a set $F \subseteq V(G)$ such that:*

1. $\mathrm{tw}(G[F]) \leqslant r$*,*
2. *the size of F is at least $(1-\epsilon)\mathit{OPT}$, where OPT is the size of a largest set satisfying the first condition.*

The idea behind the proof is as follows. Let F be an (unknown) optimum solution. Since F belongs to a weakly hyperfinite class, for every $\epsilon > 0$, there exists a subset $X \subseteq F$ of size at least $(1-\epsilon)|F|$, such that every component of $G[X]$ is of *constant size* (where this constant depends on ϵ). Thus we can enumerate all candidate for the components of $G[X]$ and proceed using Theorem 7.

Note that the problem addressed in Theorem 8 is a special case of (r, ψ)-MWIS obtained by taking ψ to be any formula satisfied by all graphs and $\mathfrak{w}$ to be a uniform weight function.

Packings of small connected subgraphs at fixed even distance. Next, we generalize the polynomial-time solvability of the Max Weight Independent Packing problem for graphs of bounded induced matching treewidth to Max Weight Distance-d Packing: the problem of packing subgraphs at distance d, for all even positive integers d. The case $d = 2$ is precisely Max Weight Independent Packing. We remark that unless $\mathsf{P} = \mathsf{NP}$, this result cannot be generalized to odd values of d, since (as shown by Eto, Guo, and Miyano [31]), the distance-3 variant of Max Independent Set problem is NP-hard for chordal graphs, which have induced matching treewidth (and even tree-independence number) at most one.

This extension again follows from a structural observation that may be of independent interest. Given a graph G and a positive integer d, we denote by G^d the *d-th power* of G, that is, the graph obtained from the graph G by adding to it all edges between pairs of distinct vertices at distance at most d. Consequently, Max Weight Distance-d Packing in a graph G is equivalent to Max Weight Independent Packing in G^{d-1}. We show the following result.

▶ **Lemma 9.** *Let G be a graph with at least one edge and d a positive integer. Then*

$$\text{tree-}\mu(G^{d+2}) \leqslant \text{tree-}\alpha(G^{d+2}) \leqslant \text{tree-}\mu(G^{d}) \leqslant \text{tree-}\alpha(G^{d}) .$$

This is a significant generalization of a result of Duchet [30], who proved an analogous result for graphs with tree-independence number at most one: if G^d is chordal, then so is G^{d+2}. As a consequence, we obtain that for every positive integer d, the class of graphs with induced matching treewidth (resp. tree-independence number) at most d is closed under taking odd powers.

▶ **Lemma 10.** *Let G be a graph and d be a positive odd integer. Then* $\text{tree-}\alpha(G^d) \leqslant \text{tree-}\alpha(G)$ *and* $\text{tree-}\mu(G^d) \leqslant \text{tree-}\mu(G)$.

Lemma 10 generalizes a result due to Balakrishnan and Paulraja [8] stating that the class of chordal graphs is closed under taking odd powers.

Combining Lemma 10 with Theorem 7 yields the following algorithmic corollary.

▶ **Theorem 11.** *For every positive integer k and even positive integer d, given a graph G of induced matching treewidth at most k, a finite family $\mathcal{H} = \{H_j\}_{j\in J}$ of connected non-null subgraphs of G, and a weight function $\mathfrak{w}\colon J \to \mathbb{Q}_+$ on the subgraphs in $\mathcal{H}$, the Max Weight Distance-d Packing problem is solvable in time polynomial in $|V(G)|$ and $|\mathcal{H}|$.*

We complement these results by observing that the class of even powers of chordal graphs (for any fixed power) is not contained in any nontrivial hereditary graph class. This implies in particular that any such class has unbounded tree-independence number and induced matching treewidth.

2.3 Solving (r, ψ)-MWIS for bounded tree-independence number

Finally, we show that (r, ψ)-MWIS can be solved in polynomial time for graphs of bounded tree-independence number. As shown by Dallard et al. [25], many natural graph classes fall into this category. Furthermore, they all have bounded induced matching treewidth.

Actually, we show tractability of a more general problem, where instead of asking for a subgraph of bounded treewidth, we ask for a subgraph of clique number bounded by r; let us call this variant $(\omega \leqslant r, \psi)$-MWIS. We remark that $(\omega \leqslant r, \psi)$-MWIS is a generalization of $(r-1, \psi)$-MWIS, as every graph of treewidth at most $r-1$ has clique number at most r and the property of being of bounded treewidth can be expressed in CMSO_2 [34, Lemma 10]. On the other hand, there are natural classes of graphs of bounded clique number and unbounded treewidth, e.g., bipartite graphs or planar graphs. However, one can observe that for graphs of bounded tree-independence number, treewidth is upper-bounded by a function of the clique number (see [27]); this phenomenon is called (tw, ω)-*boundedness* in the literature (see, e.g., [26]). Thus we conclude that in classes of bounded tree-independence number, both $(\omega \leqslant r, \psi)$-MWIS and (r, ψ)-MWIS formalisms describe the same family of problems.

▶ **Theorem 12.** *For every fixed k, r and a CMSO_2 formula ψ, the $(\omega \leqslant r, \psi)$-MWIS for graphs with tree-independence number at most k can be solved in polynomial time.*

Some examples of problems that can be solved in polynomial time with this approach include:

- finding a largest induced planar subgraph (which is equivalent to Planarization [38,54]),
- finding a largest induced odd cactus (which is equivalent to Even Cycle Transversal [10,48,52]),
- finding a largest set of vertices inducing a subgraph of maximum degree at most k (see, e.g., [37]), and
- finding the maximum number of pairwise disjoint and non-adjacent cycles (for this problem we need a slightly stronger variant of Theorem 12 that we also prove).

In particular, these results solve an open problem regarding the complexity of finding a largest set of vertices inducing a subgraph of maximum degree at most k on interval graphs (see [37, Table 2]).

We point out that the bound on the clique number of the sought-for subgraph must be constant. Indeed, the property of being a clique is easily expressible in CMSO_2, but Max Clique is NP-hard for graphs with tree-independence number 2.

2.4 Computational and structural aspects of induced matching treewidth

As already announced in the introduction, given a graph of bounded induced matching treewidth, in polynomial time we can compute its tree decomposition of bounded induced matching treewidth (however, the bound here can be larger). Formally speaking, we have the following result (it can be seen as a constant-factor XP-time approximation algorithm for induced matching treewidth).

▶ **Theorem 13.** *Let k be a positive integer and let G be an n-vertex graph satisfying $\mathsf{tree}\text{-}\mu(G) \leqslant k$. Then, in time $n^{\mathcal{O}(k)}$ we can obtain a tree decomposition $\mathcal{T}$ of G such that $\mu(\mathcal{T}) \leqslant 8k$.*

This result follows by combining an analogous result of Dallard et al. [28] for tree-independence number with known structural observations concerning relationships between tree-independence number and induced matching treewidth [60].

It is a natural question whether this algorithm can be significantly improved, in particular:

1. Can induced matching treewidth be computed *exactly* in XP-time?
2. Can induced matching treewidth be *approximated* in FPT-time?

It turns out that, again, known relations between the induced matching treewidth the tree-independence number [60] together with known results on tree-independence number [27,28] and independence number [62] imply negative answers to these (and similar) questions. We summarize them below.

▶ **Theorem 14.** *The following lower bounds hold.*

1. *For every constant $k \geqslant 4$, it is NP-complete to decide whether $\mathsf{tree}\text{-}\mu(G) \leqslant k$ for a graph G.*
2. *For every $\varepsilon > 0$, there is no polynomial-time algorithm for approximating induced matching treewidth of an n-vertex graph to within a factor of $n^{1-\varepsilon}$ unless P = NP.*
3. *There is no constant-factor FPT-approximation algorithm for induced matching treewidth, unless FPT = W[1].*
4. *For any computable functions f, g, there is no $g(k)$-approximation algorithm for computing induced matching treewidth in time $f(k) \cdot n^{o(k)}$, unless the Gap-ETH fails.*[4]

[4] Gap-ETH states that for some constant $\epsilon > 0$, distinguishing between a satisfiable 3-SAT formula and one that is not even $(1 - \epsilon)$-satisfiable requires exponential time (see [29,46]).

While the complexity of recognizing graphs with induced matching treewidth k for $k \in \{2, 3\}$ remains open, we show that the problem is solvable in polynomial time for $k = 1$. To this end, we first prove that induced matching treewidth is monotone under induced minors, that is, it cannot decrease upon deleting a vertex or contracting an edge. Then, we use a characterization of graphs G such that the square of the line graph of G is chordal (see [57]) to characterize the class of graphs with induced matching treewidth at most 1 in terms of a finite family of forbidden induced minors. This characterization implies that the family of forbidden induced *subgraphs* for this class is very restricted, containing only finitely many graphs except for cycles of length at least 6, which immediately leads to a polynomial-time recognition algorithm.

We then consider the behavior of induced matching treewidth for graphs with bounded degree, in particular, its relation to tree-independence number and treewidth. While any graph class with unbounded induced matching treewidth is necessary of unbounded tree-independence number and hence of unbounded treewidth, it is known that the converse implications hold in the absence of a fixed complete bipartite graph as a subgraph (see the full version of the paper for details). In particular, for any class of graphs with bounded maximum degree, all the aforementioned parameters (induced matching treewidth, tree-independence number, and treewidth) are all equivalent to each other, in the sense that they are either all bounded or all unbounded. We strengthen this result by showing that for any class of graphs with bounded maximum degree, the three parameters are in fact linearly related to each other.

▶ **Theorem 15.** *For every graph G with at least one edge, it holds that*
1. $\text{tree-}\alpha(G) \leqslant 2\text{tree-}\mu(G) \cdot \Delta(G)^2$, *and*
2. $\text{tw}(G) \leqslant 2\text{tree-}\mu(G) \cdot \Delta(G)^2(\Delta(G) + 1)$.

Finally, we consider two families of graphs with unbounded maximum degree. We give a lower bound on the induced matching treewidth of hypercube graphs and identify a family of P_4-free graphs with unbounded induced matching treewidth.

3 Further research directions

The obvious direction for further research is to show Conjecture 3. A somewhat easier, but still very interesting problem, would be to show the following result, which significantly strengthens both Theorem 4 and Theorem 8.

▶ **Conjecture 16.** *Let $k, r \in \mathbb{N}$ be fixed. Let G be a graph with induced matching treewidth at most k, equipped with a weight function $\mathfrak{w}\colon V(G) \to \mathbb{Q}_+$. In polynomial time we can find a maximum-weight induced subgraph of G with treewidth at most r.*

We believe that the right way of approaching Conjecture 16 is to prove an analogue of Lemmas 1 and 5: for each node of a tree decomposition of bounded induced matching treewidth, there is only a polynomial number of states of the natural dynamic programming algorithm (developed for bounded-treewidth graphs [17]) that represent an inclusion-maximal induced subgraph of treewidth at most r.

Another interesting question is to try to find a "better parameter" than induced matching width. Recently, Bergougnoux et al. [12] introduced *one-sided mim-width* (or *o-mim-width*), which can be seen as a unification of tree-independence number and *mim-width*, another well-established graph width parameter useful in solving MWIS and many other problems [11].

Just like the definition of induced matching treewidth, the definitions of mim-width and one-sided mim-width are also based on the size of a maximum induced matching in certain subgraphs of the graph (however, instead of tree decompositions they are based on branch decompositions). While induced matching treewidth lower-bounds the tree-independence number, it is incomparable with one-sided mim-width (and mim-width). Bergougnoux et al. [12] proved a result analogous to Theorem 4 for graphs given with a decomposition with bounded one-sided mim-width.

An interesting parameter that is also based on induced matchings and that in fact captures all these notions (o-mim-width, mim-width, tree-independence number, and induced matching treewidth) is *sim-width*, a graph parameter introduced in 2017 by Kang, Kwon, Strømme, and Telle (see [39]). One of the main open problems related to sim-width, first asked in [39], is whether MWIS is solvable in polynomial time for graphs with bounded sim-width.

The relationships between all these parameters and some others are summarized in Figure 1.

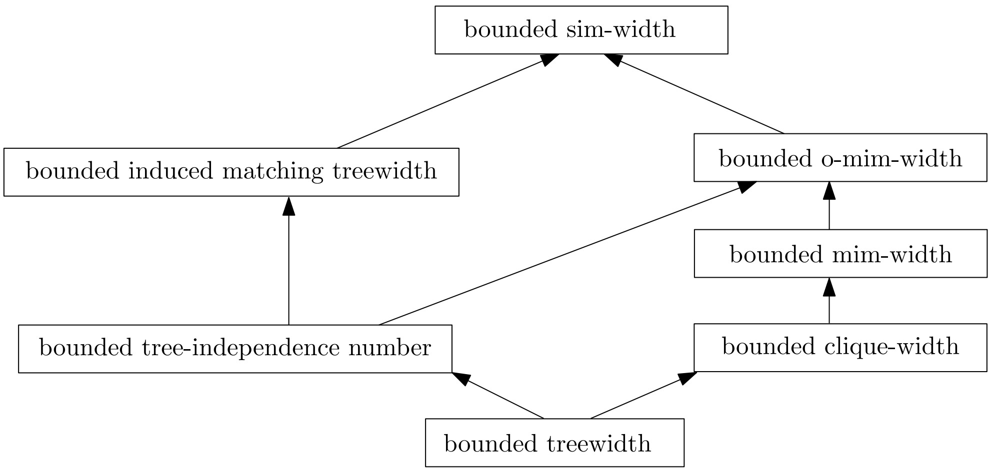

Figure 1 Relations between graph width parameters discussed in the paper.

Recall that classes of bounded tree-independence number are (tw, ω)-bounded, i.e., treewidth is bounded by a function of the clique number [26]. Combining some known results [12, 20] one can show that this is also the case for classes of bounded induced matching treewidth that additionally exclude some fixed biclique as an induced subgraph. Furthermore, excluding bicliques is necessary, as $\mathrm{tw}(K_{n,n}) = n$ and $\mathsf{tree}\text{-}\mu(K_{n,n}) = 1$. We believe that the following result could also hold.

▶ **Conjecture 17.** *For any two integers $k, t \in \mathbb{N}$ there exists an integer r such that each graph with induced matching treewidth at most k and no induced subgraph isomorphic to $K_{t,t}$ has tree-independence number at most r.*

The notion of (tw, ω)-boundedness can be seen as a variant of the extensively studied χ-*boundedness*, where we ask for which classes the chromatic number is upper-bounded by a function of the clique number [58]. We believe that graphs of bounded induced matching treewidth are χ-bounded.

▶ **Conjecture 18.** *For any two integers $k, c \in \mathbb{N}$ there exists an integer r such that each graph with induced matching treewidth at most k and clique number at most c has chromatic number at most r.*

Finally, we believe that it would be interesting to explore for what natural classes of graphs induced matching treewidth is bounded. We remark that such results are known for the closely related tree-independence number [1, 25].

Note. After the submission of the paper, Conjectures 17 and 18 were proved by Abrishami et al. [2].

References

1 Tara Abrishami, Bogdan Alecu, Maria Chudnovsky, Sepehr Hajebi, Sophie Spirkl, and Kristina Vušković. Tree independence number I. (Even hole, diamond, pyramid)-free graphs. *Journal of Graph Theory*, 106(4):923–943, 2024. doi:10.1002/jgt.23104.

2 Tara Abrishami, Marcin Briański, Jadwiga Czyżewska, Rose McCarty, Martin Milanič, Paweł Rzążewski, and Bartosz Walczak. Excluding a clique or a biclique in graphs of bounded induced matching treewidth. *CoRR*, abs/2405.04617, 2024. arXiv:2405.04617.

3 Tara Abrishami, Maria Chudnovsky, Sepehr Hajebi, and Sophie Spirkl. Induced subgraphs and tree decompositions III. Three-path-configurations and logarithmic treewidth. *Advances in Combinatorics*, 2022:6:1–29, 2022. doi:10.19086/aic.2022.6.

4 Tara Abrishami, Maria Chudnovsky, Marcin Pilipczuk, Paweł Rzążewski, and Paul D. Seymour. Induced subgraphs of bounded treewidth and the container method. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 1948–1964. SIAM, 2021. doi:10.1137/1.9781611976465.116.

5 Isolde Adler. *Width functions for hypertree decompositions.* PhD thesis, Albert-Ludwigs-Universität Freiburg im Breisgau, 2006.

6 Betül Ahat, Tınaz Ekim, and Z. Caner Taşkın. Integer programming formulations and benders decomposition for the maximum induced matching problem. *INFORMS J. Comput.*, 30(1):43–56, 2018. doi:10.1287/ijoc.2017.0764.

7 Stefan Arnborg, Derek G. Corneil, and Andrzej Proskurowski. Complexity of finding embeddings in a k-tree. *SIAM Journal on Algebraic Discrete Methods*, 8(2):277–284, April 1987. doi:10.1137/0608024.

8 Rangaswami Balakrishnan and P. Paulraja. Powers of chordal graphs. *J. Austral. Math. Soc. Ser. A*, 35(2):211–217, 1983.

9 Walid Ben-Ameur, Mohamed-Ahmed Mohamed-Sidi, and José Neto. The k-separator problem: polyhedra, complexity and approximation results. *J. Comb. Optim.*, 29(1):276–307, 2015. doi:10.1007/s10878-014-9753-x.

10 Benjamin Bergougnoux, Édouard Bonnet, Nick Brettell, and O-joung Kwon. Close relatives of feedback vertex set without single-exponential algorithms parameterized by treewidth. In Yixin Cao and Marcin Pilipczuk, editors, *15th International Symposium on Parameterized and Exact Computation, IPEC 2020, December 14-18, 2020, Hong Kong, China (Virtual Conference)*, volume 180 of *LIPIcs*, pages 3:1–3:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICS.IPEC.2020.3.

11 Benjamin Bergougnoux, Jan Dreier, and Lars Jaffke. A logic-based algorithmic meta-theorem for mim-width. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 3282–3304. SIAM, 2023. doi:10.1137/1.9781611977554.ch125.

12 Benjamin Bergougnoux, Tuukka Korhonen, and Igor Razgon. New width parameters for independent set: One-sided-mim-width and neighbor-depth. In Daniël Paulusma and Bernard Ries, editors, *Graph-Theoretic Concepts in Computer Science - 49th International Workshop, WG 2023, Fribourg, Switzerland, June 28-30, 2023, Revised Selected Papers*, volume 14093 of *Lecture Notes in Computer Science*, pages 72–85. Springer, 2023. doi:10.1007/978-3-031-43380-1_6.

13 Umberto Bertele and Francesco Brioschi. *Nonserial dynamic programming*. Academic Press, 1972.

14 Hans L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996. doi:10.1137/S0097539793251219.

15 Hans L. Bodlaender, Édouard Bonnet, Lars Jaffke, Dušan Knop, Paloma T. Lima, Martin Milanič, Sebastian Ordyniak, Sukanya Pandey, and Ondrej Suchý. Treewidth is NP-complete on cubic graphs. In Neeldhara Misra and Magnus Wahlström, editors, *18th International Symposium on Parameterized and Exact Computation, IPEC 2023, September 6-8, 2023, Amsterdam, The Netherlands*, volume 285 of *LIPIcs*, pages 7:1–7:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICS.IPEC.2023.7.

16 Hans L. Bodlaender, Pål Grønås Drange, Markus S. Dregi, Fedor V. Fomin, Daniel Lokshtanov, and Michal Pilipczuk. A $c^k n$ 5-approximation algorithm for treewidth. *SIAM J. Comput.*, 45(2):317–378, 2016. doi:10.1137/130947374.

17 Hans L. Bodlaender and Ton Kloks. Efficient and constructive algorithms for the pathwidth and treewidth of graphs. *J. Algorithms*, 21(2):358–402, 1996. doi:10.1006/jagm.1996.0049.

18 Marthe Bonamy, Edouard Bonnet, Hugues Déprés, Louis Esperet, Colin Geniet, Claire Hilaire, Stéphan Thomassé, and Alexandra Wesolek. Sparse graphs with bounded induced cycle packing number have logarithmic treewidth. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 3006–3028. SIAM, 2023. doi:10.1137/1.9781611977554.ch116.

19 Vincent Bouchitté and Ioan Todinca. Treewidth and minimum fill-in: Grouping the minimal separators. *SIAM J. Comput.*, 31(1):212–232, 2001. doi:10.1137/S0097539799359683.

20 Nick Brettell, Andrea Munaro, Daniël Paulusma, and Shizhou Yang. Comparing width parameters on graph classes. *CoRR*, abs/2308.05817, 2023. arXiv:2308.05817.

21 Kathie Cameron and Pavol Hell. Independent packings in structured graphs. *Math. Program.*, 105(2-3, Ser. B):201–213, 2006. doi:10.1007/s10107-005-0649-5.

22 Maria Chudnovsky, Rose McCarty, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzążewski. Sparse induced subgraphs in P_6-free graphs. In David P. Woodruff, editor, *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024, Alexandria, VA, USA, January 7-10, 2024*, pages 5291–5299. SIAM, 2024. doi:10.1137/1.9781611977912.190.

23 Bruno Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990. doi:10.1016/0890-5401(90)90043-H.

24 Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, MichałPilipczuk, and Saket Saurabh. *Parameterized algorithms*. Springer, Cham, 2015. doi:10.1007/978-3-319-21275-3.

25 Clément Dallard, Martin Milanič, and Kenny Štorgel. Treewidth versus clique number. III. Tree-independence number of graphs with a forbidden structure. *Journal of Combinatorial Theory, Series B*, 167:338–391, 2024. doi:10.1016/j.jctb.2024.03.005.

26 Clément Dallard, Martin Milanič, and Kenny Štorgel. Treewidth versus clique number. I. Graph classes with a forbidden structure. *SIAM J. Discrete Math.*, 35(4):2618–2646, 2021. doi:10.1137/20M1352119.

27 Clément Dallard, Martin Milanič, and Kenny Štorgel. Treewidth versus clique number. II. Tree-independence number. *J. Combin. Theory Ser. B*, 164:404–442, 2024. doi:10.1016/j.jctb.2023.10.006.

28 Clément Dallard, Fedor V. Fomin, Petr A. Golovach, Tuukka Korhonen, and Martin Milanič. Computing tree decompositions with small independence number. In Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson, editors, *51th International Colloquium on Automata, Languages, and Programming, ICALP 2024, July 8–12, 2024, Tallinn, Estonia*, volume 297 of *LIPIcs*, pages 51:1–51:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.ICALP.2024.51.

29 Irit Dinur. Mildly exponential reduction from gap 3SAT to polynomial-gap label-cover. *Electron. Colloquium Comput. Complex.*, TR16-128, 2016. URL: `https://eccc.weizmann.ac.il/report/2016/128`, `arXiv:TR16-128`.

30 Pierre Duchet. Classical perfect graphs: an introduction with emphasis on triangulated and interval graphs. In *Topics on perfect graphs*, volume 88 of *North-Holland Math. Stud.*, pages 67–96. North-Holland, Amsterdam, 1984. `doi:10.1016/S0304-0208(08)72924-4`.

31 Hiroshi Eto, Fengrui Guo, and Eiji Miyano. Distance-d independent set problems for bipartite and chordal graphs. *J. Comb. Optim.*, 27(1):88–99, 2014. `doi:10.1007/s10878-012-9594-4`.

32 Fedor V. Fomin, Ioan Todinca, and Yngve Villanger. Large induced subgraphs via triangulations and CMSO. *SIAM J. Comput.*, 44(1):54–87, 2015. `doi:10.1137/140964801`.

33 Peter Gartland, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzążewski. Finding large induced sparse subgraphs in $C_{>t}$-free graphs in quasipolynomial time. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 330–341. ACM, 2021. `doi:10.1145/3406325.3451034`.

34 Peter Gartland, Daniel Lokshtanov, Marcin Pilipczuk, Michał Pilipczuk, and Paweł Rzążewski. Finding large induced sparse subgraphs in $C_{>t}$-free graphs in quasipolynomial time. *CoRR*, abs/2007.11402, 2021. `arXiv:2007.11402`.

35 Andrzej Grzesik, Tereza Klimošová, Marcin Pilipczuk, and Michał Pilipczuk. Polynomial-time algorithm for Maximum Weight Independent Set on P_6-free graphs. *ACM Trans. Algorithms*, 18(1):4:1–4:57, 2022. `doi:10.1145/3414473`.

36 Rudolf Halin. S-functions for graphs. *Journal of Geometry*, 8(1):171–186, 1976. `doi:10.1007/BF01917434`.

37 Frédéric Havet, Ross J. Kang, and Jean-Sébastien Sereni. Improper coloring of unit disk graphs. *Networks*, 54(3):150–164, 2009. `doi:10.1002/net.20318`.

38 Bart M. P. Jansen, Daniel Lokshtanov, and Saket Saurabh. A near-optimal planarization algorithm. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1802–1811. SIAM, 2014. `doi:10.1137/1.9781611973402.130`.

39 Dong Yeap Kang, O-joung Kwon, Torstein J. F. Strømme, and Jan Arne Telle. A width parameter useful for chordal and co-comparability graphs. *Theoret. Comput. Sci.*, 704:1–17, 2017. `doi:10.1016/j.tcs.2017.09.006`.

40 Tuukka Korhonen. A single-exponential time 2-approximation algorithm for treewidth. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 184–192. IEEE, 2021. `doi:10.1109/FOCS52979.2021.00026`.

41 Tuukka Korhonen and Daniel Lokshtanov. An improved parameterized algorithm for treewidth. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 528–541. ACM, 2023. `doi:10.1145/3564246.3585245`.

42 Euiwoong Lee. Partitioning a graph into small pieces with applications to path transversal. In Philip N. Klein, editor, *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1546–1558. SIAM, 2017. `doi:10.1137/1.9781611974782.101`.

43 Euiwoong Lee. Partitioning a graph into small pieces with applications to path transversal. *Math. Program.*, 177(1-2, Ser. A):1–19, 2019. `doi:10.1007/s10107-018-1255-7`.

44 Paloma T. Lima, Martin Milanič, Peter Muršič, Karolina Okrasa, Paweł Rzążewski, and Kenny Štorgel. Tree decompositions meet induced matchings: beyond max weight independent set. *CoRR*, abs/2402.15834, 2024. `arXiv:2402.15834`.

45 Daniel Lokshantov, Martin Vatshelle, and Yngve Villanger. Independent set in P_5-free graphs in polynomial time. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 570–581. SIAM, 2014. `doi:10.1137/1.9781611973402.43`.

46 Pasin Manurangsi and Prasad Raghavendra. A birthday repetition theorem and complexity of approximating dense CSPs. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, volume 80 of *LIPIcs*, pages 78:1–78:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICS.ICALP.2017.78.

47 Dániel Marx. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. *J. ACM*, 60(6):42:1–42:51, 2013. doi:10.1145/2535926.

48 Pranabendu Misra, Venkatesh Raman, M. S. Ramanujan, and Saket Saurabh. Parameterized algorithms for even cycle transversal. In Martin Charles Golumbic, Michal Stern, Avivit Levy, and Gila Morgenstern, editors, *Graph-Theoretic Concepts in Computer Science - 38th International Workshop, WG 2012, Jerusalem, Israel, June 26-28, 2012, Revised Selcted Papers*, volume 7551 of *Lecture Notes in Computer Science*, pages 172–183. Springer, 2012. doi:10.1007/978-3-642-34611-8_19.

49 Jaroslav Nešetřil and Patrice Ossona de Mendez. *Sparsity - Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and combinatorics*. Springer, 2012. doi:10.1007/978-3-642-27875-4.

50 Stavros D. Nikolopoulos and Leonidas Palios. Minimal separators in P_4-sparse graphs. *Discrete Math.*, 306(3):381–392, 2006. doi:10.1016/j.disc.2005.12.008.

51 Yury Orlovich, Alexandre Dolgui, Gerd Finke, Valery Gordon, and Frank Werner. The complexity of dissociation set problems in graphs. *Discrete Appl. Math.*, 159(13):1352–1366, 2011. doi:10.1016/j.dam.2011.04.023.

52 Giacomo Paesani, Daniël Paulusma, and Paweł Rzążewski. Feedback Vertex Set and Even Cycle Transversal for H-free graphs: Finding large block graphs. *SIAM J. Discret. Math.*, 36(4):2453–2472, 2022. doi:10.1137/22m1468864.

53 B. S. Panda, Arti Pandey, Juhi Chaudhary, Piyush Dane, and Manav Kashyap. Maximum weight induced matching in some subclasses of bipartite graphs. *J. Comb. Optim.*, 40(3):713–732, 2020. doi:10.1007/s10878-020-00611-2.

54 Marcin Pilipczuk. A tight lower bound for Vertex Planarization on graphs of bounded treewidth. *Discret. Appl. Math.*, 231:211–216, 2017. doi:10.1016/j.dam.2016.05.019.

55 Neil Robertson and Paul D. Seymour. Graph minors. III. Planar tree-width. *J. Comb. Theory, Ser. B*, 36(1):49–64, 1984. doi:10.1016/0095-8956(84)90013-3.

56 Neil Robertson and Paul D. Seymour. Graph Minors. XIII. The Disjoint Paths Problem. *J. Comb. Theory, Ser. B*, 63(1):65–110, 1995. doi:10.1006/jctb.1995.1006.

57 Robert Scheidweiler and Sebastian Wiederrecht. On chordal graph and line graph squares. *Discrete Appl. Math.*, 243:239–247, 2018. doi:10.1016/j.dam.2018.02.013.

58 Alex Scott and Paul D. Seymour. A survey of χ-boundedness. *J. Graph Theory*, 95(3):473–504, 2020. doi:10.1002/jgt.22601.

59 Michalis Yannakakis. Node-deletion problems on bipartite graphs. *SIAM J. Comput.*, 10(2):310–327, 1981. doi:10.1137/0210022.

60 Nikola Yolov. Minor-matching hypertree width. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 219–233. SIAM, 2018. doi:10.1137/1.9781611975031.16.

61 Jie You, Jianxin Wang, and Yixin Cao. Approximate association via dissociation. *Discret. Appl. Math.*, 219:202–209, 2017. doi:10.1016/J.DAM.2016.11.007.

62 David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory Comput.*, 3:103–128, 2007. doi:10.4086/toc.2007.v003a006.