Every Graph is Essential to Large Treewidth

Bogdan Alecu[1], Édouard Bonnet[2],
Pedro Bureo Villafana[1], and Nicolas Trotignon[2]

[1]School of Computer Science, University of Leeds, Leeds, LS2 9JT, UK
[2]CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP, UMR 5668, Lyon, France

Abstract

We show that for every graph H, there is a hereditary weakly sparse graph class $\mathcal{C}_H$ of unbounded treewidth such that the H-free (i.e., excluding H as an induced subgraph) graphs of $\mathcal{C}_H$ have bounded treewidth. This refutes several conjectures and critically thwarts the quest for the unavoidable induced subgraphs in classes of unbounded treewidth, a wished-for counterpart of the Grid Minor theorem. We actually show a stronger result: For every positive integer t, there is a hereditary graph class $\mathcal{C}_t$ of unbounded treewidth such that for any graph H of treewidth at most t, the H-free graphs of $\mathcal{C}_t$ have bounded treewidth. Our construction is a variant of so-called *layered wheels*.

We also introduce a framework of abstract layered wheels, based on their most salient properties. In particular, we streamline and extend key lemmas previously shown on individual layered wheels. We believe that this should greatly help develop this topic, which appears to be a very strong yet underexploited source of counterexamples.

1 Introduction

A possible reading of the Grid Minor theorem of Robertson and Seymour [22] is that for every class $\mathcal{C}$ of unbounded treewidth there is a family $\mathcal{F}$ of subdivided walls of unbounded treewidth, such that every graph of $\mathcal{F}$ is a subgraph of some graph in $\mathcal{C}$. Said informally, this identifies the subdivided walls as the unavoidable subgraphs witnessing large treewidth. Recently, some effort has been put into unraveling the unavoidable induced subgraphs of classes of unbounded treewidth; see for instance the series "*Induced subgraphs and tree decompositions*" [2].

In this line of research, the holy grail would be a hereditary family $\mathcal{F}^*$ (analogous to the class of all subdivided walls in the subgraph case), ideally comprising a few canonical families of unbounded treewidth, such that for every class $\mathcal{C}$ of unbounded treewidth there is a subfamily $\mathcal{F} \subseteq \mathcal{F}^*$ of unbounded treewidth with the property that every member of $\mathcal{F}$ is an induced subgraph of some graph in $\mathcal{C}$. A first inspection reveals that $\mathcal{F}^*$ has to contain (an infinite subfamily of) all complete graphs, all complete bipartite graphs, all subdivided walls, and all the line graphs of subdivided walls. This however is not enough. Quickly, several incomparable families were discovered that had to be added to $\mathcal{F}^*$: the *layered wheel* of Sintiari and Trotignon [25], the *Pohoata–Davies grid* [21, 15], the *death star* of Bonamy et al. [4]. One can still entertain hopes that all these constructions can be unified in a common, and relatively descriptive family.

By generalizing the layered wheel construction, we put an end to this endeavor by showing that no family $\mathcal{F}^*$ can actually work, except the class of all graphs. A class is *weakly sparse* if it excludes, for some finite integer t, the biclique $K_{t,t}$ as a subgraph.[1]

Theorem 1.1. *For every graph H, there is a hereditary weakly sparse class $\mathcal{C}_H$ of unbounded treewidth such that the subclass of H-free graphs of $\mathcal{C}_H$ has bounded treewidth.*

Hence, if there was a graph H not in $\mathcal{F}^*$, the class $\mathcal{C}_H$ of Theorem 1.1 would not admit any desirable subfamily $\mathcal{F} \subseteq \mathcal{F}^*$. We actually show a significantly stronger result than Theorem 1.1.

Theorem 1.2. *For every positive integer t, there is a hereditary (weakly sparse) class $\mathcal{C}_t$ of unbounded treewidth such that for any graph H of treewidth at most t, the subclass of H-free graphs of $\mathcal{C}_t$ has bounded treewidth.*

One can observe that a hereditary class satisfying Theorem 1.2 has to be weakly sparse (in contrast with Theorem 1.1 when H is a complete graph or a complete bipartite graph).

In one sweep, Theorem 1.2 refutes several conjectures stating that every hereditary (weakly sparse) class of unbounded treewidth contains induced subgraphs of unbounded treewidth with some additional property. Here are some examples.

Conjecture 1 (Hajebi, Conjecture 1.15 of [19], previously refuted in [13])**.** *Every hereditary weakly sparse class of unbounded treewidth contains a subclass of 2-degenerate graphs of unbounded treewidth.*

Conjecture 2 (Hajebi, Conjecture 1.14 of [19], previously refuted in [13])**.** *Every hereditary weakly sparse class of unbounded treewidth contains a subclass of clique number at most 4 (resp. at most c for some fixed integer $c \geqslant 4$) and unbounded treewidth.*

Conjecture 3 (Trotignon)**.** *Every hereditary class of unbounded treewidth contains a subclass of string graphs of unbounded treewidth or a $K_{\ell,\ell}$ induced minor for every ℓ.*

Applying Theorem 1.1 with H being the 4-vertex clique, and 5-vertex clique (resp. $(c+1)$-vertex clique) refutes Conjectures 1 and 2, respectively. Applying Theorem 1.2 with $t = 4$ refutes Conjecture 3 as graphs of treewidth at most 4 includes the 1-subdivision of the 5-vertex clique and K_4. The former graph ensures that there is no subclass of $\mathcal{C}_t$ both of string graphs and of unbounded treewidth. Besides, $\mathcal{C}_t$ does not contain as induced subgraphs (line graphs of) subdivided walls of arbitrarily large treewidth (as these graphs are K_4-free), or equivalently $\mathcal{C}_t$ does not contain arbitrarily large walls (or grids) as induced minors (see for instance [1]). Thus, $\mathcal{C}_t$ cannot have $K_{\ell,\ell}$ induced minors for arbitrarily large ℓ, since a result

[1] All the other notions used without a definition in this introduction are defined in Section 2.

of Chudnovsky et al. [12] implies that such classes have arbitrarily large walls as induced minors or K_4-free induced subgraphs of arbitrarily large treewidth.

Chudnovsky and Trotignon [13] note that the class they use to refute Conjecture 2 with $c \geqslant 4$ has clique number $c+2$, and ask if this can be done with a class of clique number $c+1$; see Question 1.8. Our new construction achieves this as a byproduct, since the clique number of the class $\mathcal{C}_t$ of Theorem 1.2 is $t+1$ (see Observation 4).

Rose McCarty and the fourth author first observed that the then-existing layered wheels as well as the Pohoata-Davies grid are string graphs, while the death star admits $K_{\ell,\ell}$ induced minors for arbitrarily large ℓ, which motivated Conjecture 3. We also note that Chudnovsky, Fischer, Hajebi, Spirkl, and Walczak disprove the strengthening of Conjecture 3 asking for a subclass of outerstring graphs instead [10].

Let us unify the previous conjectures under (two variants of) a generalizing meta-conjecture.

Meta-conjecture 1 (Meta-conjecture(Π))**.** *Every hereditary class $\mathcal{C}$ of unbounded treewidth contains a subclass $\mathcal{C}' \subseteq \mathcal{C}$ of unbounded treewidth with property Π.*

Meta-conjecture 2 (Meta-conjecture-ws(Π))**.** *Every hereditary weakly sparse class $\mathcal{C}$ of unbounded treewidth contains a subclass $\mathcal{C}' \subseteq \mathcal{C}$ of unbounded treewidth with property Π.*

One can see that Conjectures 1 and 2 are particular cases of Meta-conjecture 2, and Conjecture 3, of Meta-conjecture 1. We observe that Meta-conjecture(Π) straightforwardly implies Meta-conjecture-ws(Π). Note also that we do not require $\mathcal{C}'$ to be itself hereditary, although for *some* properties Π this could be ensured for free.

Let us say that a property Π is *finitely-hereditary* if there is a (finite) graph H such that if H is an induced subgraph of some $G \in \mathcal{C}'$, then $\mathcal{C}'$ does not satisfy Π. For instance, the properties of being 2-degenerate, having clique number at most 4, or being a string graph are all finitely-hereditary (due to the abovementioned respective graphs H). A weaker assumption on Π is that it is *treewidth-hereditary*, i.e., that there is a (finite) integer t such that if $\mathcal{C}'$ contains as (induced) subgraphs every graph of treewidth at most t then $\mathcal{C}'$ does not satisfy Π.

One can then reformulate Theorem 1.1, and its strengthening Theorem 1.2, this way.

Theorem 1.1. *For every finitely-hereditary property Π, Meta-conjecture-ws(Π) and thus Meta-conjecture(Π) are refuted.*

Theorem 1.2. *For every treewidth-hereditary property Π, Meta-conjecture-ws(Π) and thus Meta-conjecture(Π) are refuted.*

Thus in any further attempt to find properties forced by large treewidth, one has to consider properties Π that are not finitely-hereditary nor, more generally, treewidth-hereditary. A result of Bonnet [5] overcomes the very barrier of Theorems 1.1 and 1.2 by relying on the non-hereditary parameter of average degree (or edge density): for every $\varepsilon > 0$, every hereditary weakly sparse class of unbounded treewidth admits a (non-hereditary) subclass of unbounded treewidth and average degree at most $2+\varepsilon$. Recall indeed that we do not require the subfamily $\mathcal{C}'$ to be hereditary.

Another property that is not treewidth-hereditary is that of having bounded twin-width.[2] We thus give the following special case of Meta-conjecture 1, not refuted by our present work, and motivated by the fact that the weakly sparse layered wheels, the Pohoata–Davies grids (and their extensions), and the death star all have bounded twin-width.

Conjecture 4. *Every hereditary class of unbounded treewidth admits a subclass of unbounded treewidth and bounded twin-width.*

We will come back to Conjecture 4.

Say that a graph H is *essential* if there is a hereditary class $\mathcal{C}_H$ of unbounded treewidth such that the class $\{G \in \mathcal{C}_H \ : \ G \text{ is } H\text{-free}\}$ has bounded treewidth. Call a family $\mathcal{H}$ of graphs *essential* if there is a hereditary class $\mathcal{C}_\mathcal{H}$ of unbounded treewidth such that, for every $H \in \mathcal{H}$, the class $\{G \in \mathcal{C}_\mathcal{H} \ : \ G \text{ is } H\text{-free}\}$ has bounded treewidth. In this terminology, we can restate a slightly weaker form of Theorem 1.1 and equivalent form of Theorem 1.2.

Theorem 1.3. *Every graph is essential.*

Theorem 1.4. *For every $t \in \mathbb{N}$, the family of graphs of treewidth at most t is essential.*

Our proof of Theorem 1.2 (or Theorem 1.4) revisits a variant of the layered wheels by Chudnovsky and Trotignon [13], makes it more general, and abstracts out its properties.

1.1 Layered-wheelology

Several constructions have been called *layered wheels*: two constructions in [25] and several variants in [13]. In order to facilitate further research, such as refuting Meta-conjecture(Π) and Meta-conjecture-ws(Π) for some properties Π that are not treewidth-hereditary, we abstract out mandatory and optional properties of the layered wheels.

A *layered wheel* is a (countably infinite) graph G on the same vertex set as a countably infinite *locally-finite* (i.e., every node has finite degree) planarly-embedded rooted tree T such that the following conditions hold.

(1) For every natural number n, the set of nodes at distance n from the root in T, L_n, induces in G a path[3] that goes "left-to-right" in the planar embedding of T. Each set L_n is called a *layer*, the edges induced by a layer are called *layer edges*, and the set of all layer edges is denoted by E_L.
(2) Every edge in $E(G) \setminus E_L$ is between a pair of ancestor–descendant of T.
(3) T has no arbitrarily long paths of vertices of degree 2 (in T).

We may refer to Condition (1) as the *layer condition* and to Condition (2) as the *treedepth condition* (since this is exactly the requirement of treedepth decompositions). Condition (3) is very mild, but forces T to branch. We note that the so-called *(K_4, even hole)-free layered wheel* from [25] does not satisfy the treedepth condition. Thus we propose a weakening of Condition (2).

[2]Actually *bounded treewidth* is also not treewidth-hereditary, but the associated meta-conjectures are trivially false.

[3]In previous layered wheels, L_n sometimes induced a cycle rather than a path, but this came without any functional differences.

For any non-negative integer t, a *t-stroll* between $u \in V(T)$ and $v \in V(T)$ is a u–v path P in the graph $(V(T), E(T) \cup E_L)$ such that P has no $t+1$ consecutive layer edges, and the intersection of every layer with P is consecutive in P (i.e., a stroll cannot reenter a previously visited layer). Two nodes $u, v \in V(T)$ are in a *t-wide ancestor–descendant* relation if there is a t-stroll between u and v. Then we define:

(2') There is a finite integer $t \geqslant 0$, such that every edge in $E(G) \setminus E_L$ is between a pair of t-wide ancestor–descendant of T.

A *generalized layered wheel* is one that satisfies Conditions (1), (2'), and (3). The (K_4, even hole)-free layered wheel of Sintiari and Trotignon satisfies Condition (2') with $t = 1$. Note that Condition (2') with $t = 0$ coincides with Condition (2).

We now define some further optional properties:

(4) For every $i \neq j \in \mathbb{N}$, there is at least one edge from L_i to L_j. G is then called *proper*.
(5) T has no leaf. (Every maximal branch from the root is infinite.) G is then called *neat*.
(6) There is an integer d such that every node of T has at most d children in T. G is then called *bounded*, more precisely *d-bounded*.
(7) Every node in every L_n has at most $f(n)$ children in T, for some particular function $f \colon \mathbb{N} \to \mathbb{N}$. G is then called *f-bounded*.
(8) There is an integer t such that for every node v of T, there is a subset X_v of ancestors of v of size at most t with the property that in $G - E_L$ every edge with exactly one endpoint within v and its descendants has its other endpoint in X_v. G is then called *upward-restricted*, more precisely *t-upward-restricted*.
(9) For every node v of T, the ancestors of v that are adjacent to v in G form a clique in G. G is then called *upward-simplicial*.
(10) For every triple of pairwise distinct nodes u, v, w, if u is an ancestor of v, v is an ancestor of w, and $uw \in E(G)$, then $uv \in E(G)$. G is then called *upward-nested*.

Note that, as a slight abuse of language, the above properties attribute to G what is actually a property of the pair (G, T). Figure 1 shows the first few layers of a layered wheel G with rooted tree T, which can be completed to a proper, neat, 3-bounded layered wheel, but which is not *upward-simplicial* nor *upward-nested*. Note that the class yielded

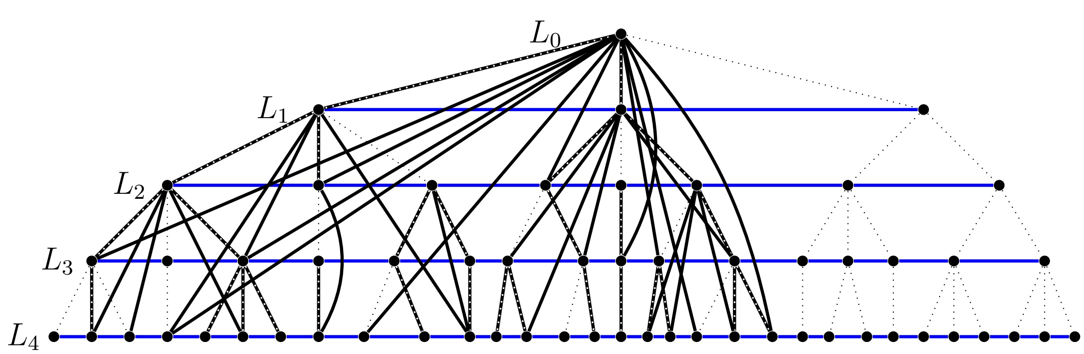

Figure 1: Possible first five layers of a proper, neat, 3-bounded layered wheel G with rooted tree T. The edges of T are dotted, the layer edges of G are in blue, and its other edges are in black (with white dots if they are also edges of T).

by a proper layered wheel has unbounded treewidth as it contains every clique as a minor (just contract each layer to a single vertex). The proof of Theorem 1.2 consists of building a particular proper, neat, bounded, upward-restricted layered wheel.

All layered wheels built so far are proper and neat. The so-called *(theta, triangle)-free layered wheel* of Sintiari and Trotignon [25] is further 17-bounded, upward-restricted, upward-simplicial, and upward-nested. The layered wheels of Chudnovsky and Trotignon [13] are f-bounded for well-chosen functions f, upward-simplicial, and upward-nested (but not always upward-restricted).

A family $\mathcal{F}$ of induced subgraphs of G satisfies the *bounded-branch* property if there is an integer h such that for every graph $G' \in \mathcal{F}$ and for every $v \in V(T)$ (not necessarily in $V(G')$), there is a *downward path* in T starting at v (i.e., a maximal, possibly infinite, path starting at v, and iteratively visiting a child of the current node while the current node is not a leaf) that contains at most h vertices of G'. Note that in a neat layered wheel, every downward path is infinite. Here is a crucial feature of upward-restricted layered wheels.

Theorem 1.5. *For every neat, upward-restricted layered wheel W, every class of finite graphs satisfying the bounded-branch property in W has bounded treewidth.*

Theorem 1.5 is essentially proved in [13], albeit not under this formalism. We shall then simply see how to engineer a neat, proper, upward-restricted layered wheel W_t such that the absence of any (a priori unknown) graph of treewidth at most t in finite induced subgraphs of W_t ensures the bounded-branch property.

Theorem 1.5 further leads to the following theorem.

Theorem 1.6. *The n-vertex induced subgraphs of any neat, upward-restricted layered wheel have treewidth $O(\log n)$.*

Theorem 1.6 unifies and extends what was individually observed on the finite induced subgraphs of the (theta, triangle)-free layered wheel [25], and on the K_t-free finite induced subgraphs of the layered wheels from [13]. Indeed, one can observe that the K_t-free finite induced subgraphs of an upward-simplicial, upward-nested layered wheel have the property of upward-restriction. On the other hand, it is not difficult to see that any proper, bounded layered wheel admits arbitrarily large n-vertex induced subgraphs of treewidth $\Omega(\log n)$. Thus a source of classes with logarithmic treewidth is given by considering the finite induced subgraphs of any proper, neat, bounded, upward-restricted layered wheel.

Someone interested in refuting Conjecture 4 by developing further layered wheels should be aware of the following obstacle.

Theorem 1.7. *The class of finite induced subgraphs of any neat, upward-restricted layered wheel has bounded twin-width.*

The *neat* condition mainly simplifies the proof of Theorem 1.7, and should not be needed. Thus, one would have to drop the upward-restricted condition, and rely on other specificities of one's layered wheel to bound the treewidth of subclasses $\mathcal{C}'$ of bounded twin-width.

1.2 Related work and further directions

We identify a few directions related to the current work.

Essentiality for other graph parameters. A reader familiar with the paper "*Induced subgraphs of induced subgraphs of large chromatic number*" of Girão et al. [17] likely noticed a parallel with Theorem 1.1. The authors show that for every graph H, there is a class $\mathcal{C}$ of unbounded chromatic number such that the H-free graphs of $\mathcal{C}$ have bounded chromatic number. (Furthermore, if H has at least one edge, the class $\mathcal{C}$ can be picked to have the same clique number as H.) In this sense, the current paper can be thought of as "*Induced subgraphs of induced subgraphs of large treewidth*." We note however that the construction in [17] contains, by design, arbitrarily large induced bicliques, so could not guide us in achieving Theorem 1.1.

We now know that no fixed induced subgraph can be removed from every class $\mathcal{C}$ of unbounded p while preserving that p is unbounded, for parameter p equal to chromatic number or treewidth. Other graph parameters can be considered such as clique-width, twin-width, etc. As our construction leads to a weakly sparse class, within which treewidth and clique-width are known to be functionally equivalent [18], our paper also offers a complete answer for clique-width. The case of twin-width remains interesting. It is known that the class of permutation graphs is a minimal hereditary class of unbounded twin-width [8]. This translates into the essentiality (for twin-width) of every permutation graph. To our knowledge, the essentiality (for twin-width) of any other graph is open. In general, we propose the following questions.

Meta-problem 1 (Characterize p-essential graphs)**.** *For a parameter of choice p, which graphs H are such that there is a hereditary class of unbounded p whose H-free graphs have bounded p?*

Observe that Ramsey's theorem can be rephrased as the fact that complete or edgeless graphs are the only p-essential graphs when p is the number of vertices.

Characterizing the essential families. For families rather than single graphs, our understanding of essentiality (for treewidth) is not quite complete. What about families $\mathcal{H}$ of unbounded treewidth? Such a family $\mathcal{H}$ is essential if and only if its hereditary closure is a minimal hereditary class of unbounded treewidth. Thus, any family $\mathcal{H}$ consisting of infinitely many complete graphs or of infinitely many complete bipartite graphs is essential. On the other hand, a family $\mathcal{H}$ consisting of infinitely many complete graphs (resp. complete bipartite graphs) plus at least one graph that is not a clique (resp. not a biclique) is *not* essential. Hence, we narrowed down the open cases to weakly sparse families. Is there a weakly sparse family whose hereditary closure is a minimal hereditary class of unbounded treewidth? This is precisely a question of Cocks [14], which we reformulate here.

Conjecture 5 (Cocks's Conjecture 1.5 in [14])**.** *A family of unbounded treewidth is essential if and only if it contains only complete graphs, or only complete bipartite graphs.*

Treewidth logarithmically bounded in the number of vertices. Sintiari and Trotignon [25] remarked that their layered wheel constructions have treewidth logarithmic in their number of vertices, and suggested relaxing the bounded treewidth condition and investigating logarithmic treewidth instead. Many (NP-hard) problems can be solved in polynomial time in n-vertex graphs of treewidth $O(\log n)$, such as every problem expressible in the so-called *Existential Counting Modal Logic* of Pilipczuk, a large fragment of Monadic Second-Order logic [20]. In the past five years, several classes have been shown to have logarithmic treewidth [9, 4, 6, 11, 3]. Our result shows that any graph is responsible for the transition from unbounded to bounded treewidth in some class. Is this true for the transition between superlogarithmic and logarithmic treewidth?

Problem 1. *For which families $\mathcal{H}$ is there a hereditary class $\mathcal{C}$ of superlogarithmic treewidth such that, for every $H \in \mathcal{H}$, the H-free graphs of $\mathcal{C}$ have at most logarithmic treewidth?*

Problem 1 is already open for singleton families $\mathcal{H} = \{H\}$. It should be noted that the Pohoata–Davies grid [21, 15] has no large clique, biclique, subdivided wall, or its line graph as an induced subgraph and has superlogarithmic treewidth: the $n \times n$ Pohoata–Davies grid has treewidth $\Theta(n)$.

2 Preliminaries

For any positive integer i, we denote by $[i]$ the set of integers $\{1, 2, \ldots, i\}$.

2.1 Induced subgraphs, neighborhoods, and some special graphs

We denote by $V(G)$ and $E(G)$ the set of vertices and edges of a graph G, respectively. A graph H is a *subgraph* of a graph G if H can be obtained from G by vertex and edge deletions. Graph H is an *induced subgraph* of G if H can be obtained from G by vertex deletions only. A graph G is *H-free* if G does not contain H as an induced subgraph. For $S \subseteq V(G)$, the *subgraph of G induced by S*, denoted $G[S]$, is obtained by removing from G all the vertices that are not in S (together with their incident edges). Then $G - S$ is a short-hand for $G[V(G) \setminus S]$. If $F \subseteq E(G)$, then $G - F$ is the graph $(V(G), E(G) \setminus F)$. We denote by $N_G(v)$ and $N_G[v]$, the open, respectively closed, neighborhood of v in G. For $S \subseteq V(G)$, we set $N_G(S) := (\bigcup_{v \in S} N_G(v)) \setminus S$ and $N_G[S] := N_G(S) \cup S$. A graph H is a *minor* (resp. *induced minor*) of a graph G if H can be obtained from a subgraph (resp. *induced subgraph*) of G by performing edge contractions.

A *subdivision* of a graph G is any graph obtained from G by replacing each edge of G by a path of at least one edge. The *s-subdivision* of G is the graph obtained from G by replacing each edge of G by a path of $s+1$ edges.

The t-vertex *complete graph*, denoted by K_t, is obtained by making adjacent every pair of two distinct vertices among t vertices, and the *complete bipartite graph* $K_{t,t}$ with bipartition (A, B) such that $|A| = |B| = t$ is obtained by making every vertex of A adjacent to every vertex of B. A *clique* of a graph G is a set of vertices inducing a complete graph. The *clique number* of a graph G is the largest t such that K_t is an (induced) subgraph of G. A graph class $\mathcal{C}$ is said *weakly sparse* if there is a finite integer t such that no graph of $\mathcal{C}$ contains $K_{t,t}$ as a subgraph. A class $\mathcal{C}$ is *hereditary* if for every $G \in \mathcal{C}$ and every induced subgraph H of G, $H \in \mathcal{C}$.

A *string graph* is the intersection graph of some collection of (non-self-intersecting) curves in the plane (usually called strings), or equivalently the intersection graph of a collection of connected sets of some planar graph. It is known that the 1-subdivision of any non-planar graph is not a string graph [24].

A (possibly infinite) graph is *chordal* if it does contain any cycle of length at least 4 as an induced subgraph. A vertex v in a graph G is *simplicial* if $N_G(v)$ (or $N_G[v]$) is a clique of G. The following is a classical fact.

Fact 1. *Every finite chordal graph admits a simplicial vertex.*

For two positive integers k, ℓ, the $k \times \ell$ *grid* is the graph on $k\ell$ vertices, say, $v_{i,j}$ with $i \in [k], j \in [\ell]$, such that $v_{i,j}$ and $v_{i',j'}$ are adjacent whenever either $i = i'$ and $|j - j'| = 1$ or $j = j'$ and $|i - i'| = 1$. For what comes next, it is helpful to identify vertex $v_{i,j}$ with the point (i, j) of $\mathbb{N}^2$. For $k \geqslant 2$ the $k \times k$ *wall* is the subgraph of the $2k \times k$ grid obtained by removing every "vertical edge" on an "even column" when the edge bottom endpoint is on an "odd row", and every "vertical edge" on an "odd column" when the edge bottom endpoint is on an "even row," and finally by deleting the two vertices of degree 1 that this process creates. See Figure 2 for an illustration of the 5×5 grid and 5×5 wall.

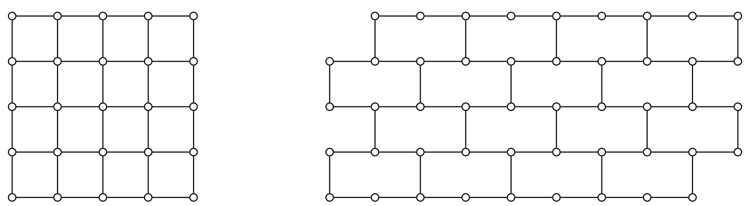

Figure 2: The 5×5 grid (left) and the 5×5 wall (right).

2.2 Treewidth, brambles, and separation number

A *tree-decomposition* of a graph G is a pair (T, β) where T is a tree and β is a map from $V(T)$ to $2^{V(G)}$ satisfying the following conditions:

- for every $uv \in E(G)$, there is an $x \in V(T)$ such that $\{u, v\} \subseteq \beta(x)$, and
- for every $v \in V(G)$, the set of nodes $x \in V(T)$ such that $v \in \beta(x)$ induces a non-empty subtree of T.

The *width* of (T, β) is defined as $\max_{x \in V(T)} |\beta(x)| - 1$, and the *treewidth* of G, denoted by $\mathrm{tw}(G)$, is the minimum width of (T, β) taken among every tree-decomposition (T, β) of G.

The notion of *bramble* was introduced by Seymour and Thomas [23] as a min-max dual to treewidth. A *bramble* of a graph G is a set $\mathcal{B} := \{B_1, \ldots, B_q\}$ of connected subsets of $V(G)$ such that for every $i, j \in [q]$ the pair B_i, B_j *touch*, i.e., $B_i \cap B_j \neq \emptyset$ or there is some $u \in B_i$ and $v \in B_j$ with $uv \in E(G)$. A *hitting set* of $\{B_1, \ldots, B_q\}$ is a set X such that for every $i \in [q]$, $X \cap B_i \neq \emptyset$, that is, X intersects every B_i. The *order* of bramble $\mathcal{B}$ is the minimum size of a hitting set of $\mathcal{B}$.

The *bramble number* of G, denoted by $\mathrm{bn}(G)$, is the maximum order of a bramble of G. The treewidth and bramble number are tied: for every graph G, $\mathrm{tw}(G) = \mathrm{bn}(G) - 1$ [23].

A subset S of vertices in a graph G is an *α-balanced separator* if every connected component of $G - S$ has at most $\alpha|V(G)|$ vertices. The *separation number* of a graph G, denoted by $\text{sn}(G)$, is the smallest integer s such that every (induced) subgraph of G has a $\frac{2}{3}$-balanced separator of size at most s. Treewidth and separation number are linearly tied.

Theorem 2.1 ([16])**.** *For every graph G, $sn(G) - 1 \leqslant tw(G) \leqslant 15sn(G)$.*

Small $\frac{2}{3}$-balanced separators are yielded by small(er) α-balanced separators with $\alpha < 1$.

Lemma 2.2. *Fix any graph G and $\alpha \in [\frac{2}{3}, 1)$. If every (induced) subgraph of G admits an α-balanced separator of size at most s, then G has separation number at most $\frac{\log(2/3)}{\log \alpha} \cdot s$.*

Proof. Take any subgraph H of G, and define S_i as an α-balanced separator of size at most s of the largest component of $H - \bigcup_{1 \leqslant j < i} S_j$, while this largest component has size more than $\frac{2}{3}|V(H)|$. Say that this process eventually defines the sets $S_1, \ldots, S_t \subseteq V(H)$. As the size of the largest component is multiplied by a factor of at most α at every iteration, and $\alpha^{\log(2/3)/\log \alpha} = \frac{2}{3}$, we have $t \leqslant \frac{\log(2/3)}{\log \alpha}$. Hence H has a $\frac{2}{3}$-balanced separator $\bigcup_{i \in [t]} S_i$ of size at most st, so G has separation number at most $\frac{\log(2/3)}{\log \alpha} \cdot s$. □

We will rely on the combination of Theorem 2.1 and Lemma 2.2.

Lemma 2.3. *Fix any graph G and $\alpha \in [\frac{2}{3}, 1)$. If every (induced) subgraph of G admits an α-balanced separator of size at most s, then G has treewidth at most $15 \cdot \frac{\log(2/3)}{\log \alpha} \cdot s$.*

2.3 Twin-width and oriented twin-width

A *partition sequence* $\mathcal{P}_n, \ldots, \mathcal{P}_1$ of an n-vertex graph G is such that $\mathcal{P}_i$ is a partition of $V(G)$ for every $i \in [n]$, $\mathcal{P}_n = \{\{v\} \mid v \in V(G)\}$, and for every $i \in [n-1]$, $\mathcal{P}_i$ is obtained from $\mathcal{P}_{i+1}$ by merging two parts $X, Y \in \mathcal{P}_{i+1}$ into one: $X \cup Y$. In particular, note that $\mathcal{P}_1 = \{V(G)\}$. For every $i \in [n]$, we denote by $\mathcal{R}(\mathcal{P}_i)$ the graph with one vertex per part of $\mathcal{P}_i$, and an edge between $P \neq P' \in \mathcal{P}_i$ whenever there are (possibly equal) $u, v \in P$ and $u', v' \in P'$ such that $uu' \in E(G)$ and $vv' \notin E(G)$. The *twin-width* of G, denoted by $\text{tww}(G)$, is the least integer d for which G admits a partition sequence $\mathcal{P}_n, \ldots, \mathcal{P}_1$ such that for every $i \in [n]$, the graph $\mathcal{R}(\mathcal{P}_i)$ has maximum degree at most d.

From the definition of twin-width, it can be observed that the twin-width of a graph is at least the twin-width of any of its induced subgraphs.

Observation 1 ([8])**.** *For any graph G and any induced subgraph H of G, $tww(H) \leqslant tww(G)$.*

We also define the digraph $\overrightarrow{\mathcal{R}}(\mathcal{P}_i)$, also on vertex set $\mathcal{P}_i$, with an arc from $P \in \mathcal{P}_i$ to $P' \in \mathcal{P}_i \setminus \{P\}$ whenever there are $u \neq v \in P$ and some $w \in P'$ such that $uw \in E(G)$ and $vw \notin E(G)$. Similarly, the *oriented twin-width* of G, denoted by $\text{otww}(G)$, is the least integer d for which G admits a partition sequence $\mathcal{P}_n, \ldots, \mathcal{P}_1$ such that for every $i \in [n]$, the digraph $\overrightarrow{\mathcal{R}}(\mathcal{P}_i)$ has maximum outdegree at most d.

It is immediate from the definition that the twin-width of any graph is upper bounded by its oriented twin-width. Rather surprisingly, twin-width and oriented twin-width are functionally equivalent.

Theorem 2.4 ([7])**.** *There is some c such that for every G, $tww(G) \leqslant otww(G) \leqslant 2^{2^{c \cdot tww(G)}}$.*

This simplifies the task of bounding the twin-width of a class. One can just upper bound its oriented twin-width. We see such an example with a proof of Theorem 1.7.

Theorem 1.7. *The class of finite induced subgraphs of any neat, upward-restricted layered wheel has bounded twin-width.*

Proof. Fix a neat, upward-restricted layered wheel W with rooted tree T. As W is upward-restricted there is an integer t such that for every node v of T, there is a set X_v of at most t ancestors of v with the property that every non-layer edge of W with exactly one endpoint among the descendants of v has its other endpoint in X_v.

For every natural number n, let $W_{\leqslant n}$ be the finite subgraph of W induced by its first n layers. By Observation 1, we just need to upper bound the twin-width of $W_{\leqslant n}$ (independently of n). By Theorem 2.4, it is enough to upper bound the oriented twin-width of $W_{\leqslant n}$.

We build a partition sequence $\mathcal{S}$ of $W_{\leqslant n}$ as follows. For every i from n down to 1, denote by $p_1, \ldots, p_h$ the parents of the nodes in the ith layer, from left to right. For every j from 1 to h, denote by $q_{j,1}, \ldots, q_{j,a_j}$ the children of p_j, from left to right. For every j from 1 to h, for every k from 2 to a_j, merge the part containing $\{q_{j,1}, \ldots, q_{j,k-1}\}$ with the part containing $q_{j,k}$. After this is done, for every j from 1 to h, merge the part containing $\{q_{j,1}, \ldots, q_{j,a_j}\}$ with the part $\{p_j\}$. Note that, as W is neat, this well-defines a partition sequence of $W_{\leqslant n}$.

Take any partition $\mathcal{P}$ of $\mathcal{S}$, and any part $P \in \mathcal{P}$. We claim that the outdegree of P in $\overrightarrow{\mathcal{R}}(\mathcal{P})$ is at most $t+3$. Note that P has either a unique topmost vertex in T, or all its topmost vertices in T share the same parent. In the former case, let v be this unique topmost vertex. Part P has at most $t+3$ outneighbors in $\overrightarrow{\mathcal{R}}(\mathcal{P})$: at most one part to its left, at most two parts to its right, and at most one part per vertex in X_v. In the latter case, let v be the shared parent. Part P has again at most $t+3$ outneighbors in $\overrightarrow{\mathcal{R}}(\mathcal{P})$: at most one part to its left, at most one part to its right, and at most one part per vertex in $X_v \cup \{v\}$. ☐

3 Chordal trigraphs and their tree representations

It will be convenient to work with *trigraphs*.[4] Trigraphs are graphs with two disjoint edge sets: black edges and red edges. Red edges should be viewed here as virtual edges, and will not be part of the eventual construction. They are however useful for the inductive definitions and in the proofs.

A *trigraph* is a triple (V, E_B, E_R), where E_B and E_R are disjoint subsets of $\binom{V}{2}$ called the set of *black* (or *real*) and *red* (or *virtual*) edges, respectively. For any $u, v \in V(G)$, the *adjacency type* of uv is defined as 'black' if $uv \in E_B(G)$, or 'red' if $uv \in E_R(G)$, or 'non-edge' otherwise. The graph $(V, E_B \cup E_R)$, denoted by $\mathcal{T}(G)$, is called the *total graph* of G, and (V, E_B) the *real* graph of G. A trigraph is said to be *chordal* if its total graph is chordal.

Two trigraphs G and H are *isomorphic* if there is a bijection f from $V(G)$ to $V(H)$ such that uv (in G) and $f(u)f(v)$ (in H) have the same adjacency type, for every $u, v \in V(G)$. When $X \subseteq V(G)$, we denote by $G[X]$ *the subtrigraph of G induced by X*, that is, the trigraph

[4] Here the trigraphs and red edges are unrelated to twin-width.

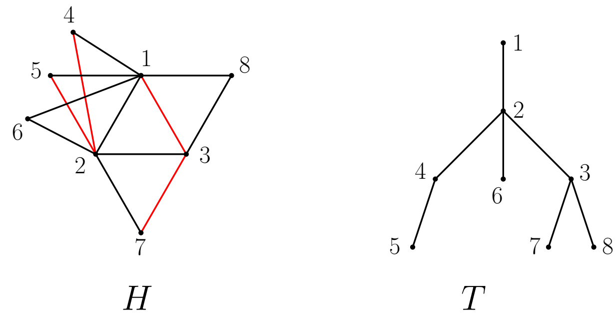

Figure 3: A chordal trigraph H and a tree representation T of H.

$(X, E_B(G) \cap \binom{X}{2}, E_R(G) \cap \binom{X}{2})$. We say that that a trigraph H is an *induced subtrigraph* of a trigraph G if for some set $X \subseteq V(G)$, H is isomorphic to $G[X]$. A trigraph is *H-free* if H is not an induced subtrigraph of G.

When T is a rooted tree, we use the words *parent*, *child*, *ancestor*, and *descendant* with their usual meaning. A node is its own ancestor and descendant. We use the notation $\mathsf{parent}(v)$ for the parent of a non-root node v, $\mathsf{ancestors}(v)$ for the set of ancestors of v, and $\mathsf{descendants}(v)$ for the set of descendants of v. Nodes u, v are *siblings* in T if they have the same parent. The *depth* of a node is its distance to the root.

A *tree representation* of a chordal trigraph H' with total graph $H := \mathcal{T}(H')$ is a rooted tree T on the same vertex set such that:

(1) for every $v \in V(H)$, $N_H(v) \cap \mathsf{ancestors}(v)$ is a clique in H,
(2) for every non-root vertex $v \in V(H)$,

$$N_H(v) \subseteq \mathsf{descendants}(v) \cup (\mathsf{ancestors}(v) \cap N_H[\mathsf{parent}(v)]), \text{ and}$$

(3) if u and v are siblings in T, then u and v have a common ancestor x such that ux and vx have distinct adjacency types in H'.

Condition (3), the only condition depending on H' rather than merely on its total graph, is called *sibling condition*. It imposes that siblings cannot have the exact same adjacency types toward their (common) strict ancestors. An example is represented in Figure 3. Observe that vertices 4 and 5 are adjacent in T but not in H; so in general, T is *not* a spanning tree of H. More classical representations of chordal graphs would make vertices 4 and 5 siblings. Here they cannot because of our sibling condition. Nonetheless, the following holds.

Lemma 3.1. *Every finite chordal trigraph admits a tree representation.*

Proof. Fix any finite chordal trigraph H' and set $H := \mathcal{T}(H')$. We prove the property by induction on $|V(H)|$. If H (or H') has a single vertex, then a single-node tree represents H'.

Suppose now that $|V(H)| \geqslant 2$. By Fact 1, there exists a simplicial vertex v in H. Let T' be a tree representation of $H' - \{v\} := H'[V(H') \setminus \{v\}]$ obtained from the induction hypothesis. If v is isolated in H, then define T by adding v as a child of any (fixed) leaf in T'.

Otherwise, let u be *the* deepest vertex in T' of $N_H(v)$. Since v is a simplicial vertex, $N_H(v)$ is a clique, so u is well-defined. Indeed by Condition (2), every pair of adjacent vertices in $\mathcal{T}(H' - \{v\})$ are in an ancestor–descendant relation in T', and every clique of $\mathcal{T}(H' - \{v\})$ lies in a single branch of T'.

If u has no child w in T' such that for every strict ancestor x of w, wx and vx have the same adjacency type in H', then define T by adding v as a child of u in T'. Otherwise, such a vertex w exists, and it is unique by the sibling condition. Let w' be a deepest descendant of w (possibly w itself) such that for every strict ancestor x of w' in T', we have that $w'x$ and vx have the same adjacency type in H'. Then we build T from T' by adding v as a child of w'.

In all cases, one can easily check that T is indeed a tree representation of H'. □

A *chordal completion* of a graph G is any chordal graph H such that $V(H) = V(G)$ and $E(G) \subseteq E(H)$. It is known that the treewidth of a graph G is the minimum of the clique number of H minus one, overall chordal completions H of G. A chordal completion of a graph G will also be thought as a chordal trigraph H such that $V(H) = V(G)$ and $E_B(H) = E(G)$.

4 The new upward-restricted layered wheel

We now construct the promised layered wheel W_t, in order to prove Theorem 1.2. For that, we build a countably infinite chordal trigraph G_t parameterized by a positive integer t.

Construction of G_t, T_t, and W_t. We now inductively build the trigraph G_t and a tree T_t on the same vertex set as G_t. Note that T_t is a graph (not a trigraph), and is *not* a subgraph of $\mathcal{T}(G_t)$.

The vertex set of G_t is partitioned into infinitely many layers $L_0, L_1, L_2, \ldots$, each inducing a path of real edges of finite size. The first layer consists of a single vertex, which is the root of T_t. For every integer $i \geqslant 0$, let $L_{\leqslant i} := \bigcup_{0 \leqslant j \leqslant i} L_j$ and assume that $G_t[L_{\leqslant i}]$ and $T_t[L_{\leqslant i}]$ are defined. The induced path L_{i+1} is built as follows.

For every $u \in L_i$ taken from left to right, where the first (resp. last) vertex introduced in L_i is leftmost (resp. rightmost), initialize $\mathsf{children}(u)$ to the empty set. For every ordered pair (B, R) of disjoint subsets of

$$N^{\uparrow}[u] := (N_{\mathcal{T}(G_t)}(u) \cap L_{\leqslant i-1}) \cup \{u\}$$

such that $|B \cup R| \leqslant t$, append a vertex $v_{B,R}$ to L_{i+1} at its rightmost end, putting the new edge in $E_B(G_t)$. Add $v_{B,R}$ to the set $\mathsf{children}(u)$. In G_t, add a black edge between $v_{B,R}$ and every vertex of B and a red edge between $v_{B,R}$ and every vertex of R. Every vertex of $\mathsf{children}(u)$ is a child of u in T_t. This finishes the inductive definitions of the trigraph G_t and the tree T_t.

The layered wheel W_t is the real graph of G_t (with rooted tree T_t), and $\mathcal{C}_t$ is the class of finite induced subgraphs of W_t. Note that $\mathcal{T}(G_t)$ is also a layered wheel with rooted tree T_t, and a spanning supergraph of W_t. In Figures 4 and 5, we illustrate the first 3 layers of T_t, G_t and W_t when $t = 1$, respectively $t \geqslant 2$.

We make some observations.

Observation 2. *For every $u \in V(G_t)$, $|N^{\uparrow}[u]| \leqslant t + 1$.*

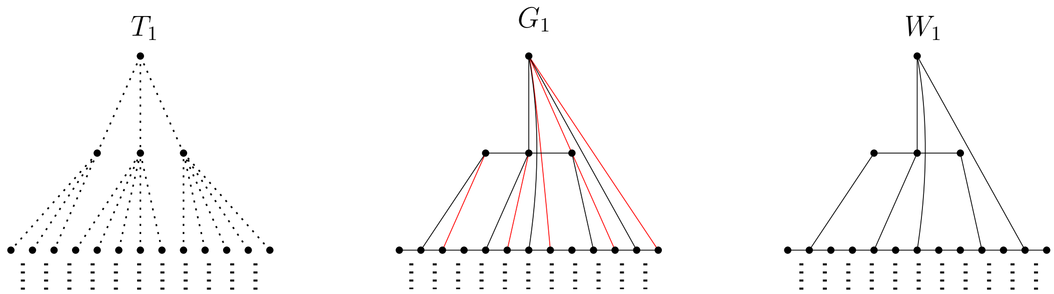

Figure 4: First three layers of T_1, G_1, W_1.

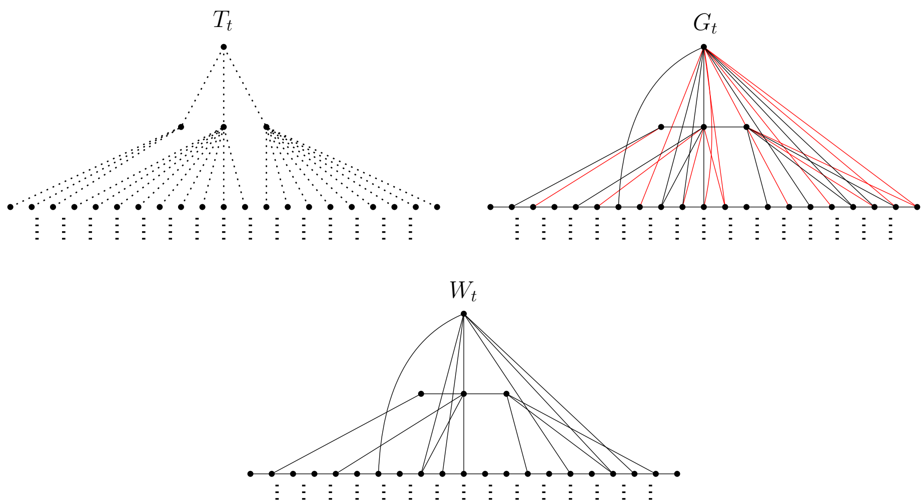

Figure 5: First three layers of T_t, G_t, W_t when $t \geqslant 2$.

Observation 3. *The layered wheels $\mathcal{T}(G_t)$, and hence W_t, with rooted tree T_t are upward-restricted.*

Proof. For every $u \in V(G_t)$, the required set X_u is simply $N^{\uparrow}[u]$ (with $|X_u| \leqslant t+1$). $\square$

Observation 4. *Every graph of $\mathcal{C}_t$ has clique number at most $t+1$.*

Proof. Let K be a (finite) clique in W_t, and let $u \in K$ be a deepest vertex of K, i.e., no vertex of K is in a layer of strictly larger index. Thus $K \setminus N^{\uparrow}[u]$ is empty or consists of exactly one sibling u' of u. In the former case, we are done, by Observation 2. In the latter case, $|K \setminus \{u, u'\}| \leqslant t-1$, as otherwise u and u' would have the exact same neighbors and non-neighbors among their (common) ancestors, which is ruled out by our construction. $\square$

The layered wheels $\mathcal{T}(G_t)$ and W_t, with rooted tree T_t, are 3^{t+1}-bounded.

Observation 5. *Every node of T_t has fewer than 3^{t+1} children.*

Proof. Indeed the number of ordered pairs (B, R) such that B and R are disjoint subsets of a set of size at most $t+1$ with $|B \cup R| \leqslant t$ is at most $\max(3^{t+1} - 2^{t+1}, 3^t) < 3^{t+1}$. $\square$

A *downward path* from $u \in V(G_t) = V(T_t)$ is an infinite path $u_1u_2\dots$ in T_t such that $u_1 = u$ and for every $i > 1$, u_{i-1} is the parent of u_i in T_t (note that a downward path is not necessarily a path in G_t). The *upward path* from $u \in V(G_t)$ is the unique path in T_t from u to the root of T_t. Observe that the union of the upward path from u and any downward path from u contains exactly one vertex in each layer: it is a branch of T_t.

We now show the main technical lemma of this section: The finite H-free induced subgraphs of G_t satisfy the bounded-branch property for every chordal trigraph H whose total graph has clique number at most $t+1$.

Lemma 4.1 (Bounded-branch property)**.** *For every chordal trigraph H with $\mathcal{T}(H)$ of clique number at most $t+1$, every set $X \subseteq V(G_t)$ such that $G_t[X]$ is H-free, and every $u \in V(G_t)$, there exists a downward path P_u from u satisfying*

$$|V(P_u) \cap X| \leqslant |V(H)| - 1.$$

Proof. Let T_H be a tree representation of H, as defined in Section 3. We number $v_1, \dots, v_h$ the vertices of trigraph H such that for all $i \in [h]$, v_i is a leaf of $T_H[\{v_1, \dots, v_i\}]$. Fix any set $X \subseteq V(G_t)$ such that $G_t[X]$ is H-free, and any $u \in V(G_t)$.

Let us prove for every integer $k \in [h]$, the following disjunction, which we call $\mathcal{P}_k$.

(1) There exists a subtree S_k of T_t, rooted at u, such that all the following facts hold:
- $G_t[X \cap V(S_k)]$ is isomorphic to $H[\{v_1, \dots, v_k\}]$,
- denoting by v'_i the image of v_i in the isomorphism from $H[\{v_1, \dots, v_k\}]$ to $G_t[X \cap V(S_k)]$, for any pair of vertices v_i, v_j with $i, j \in [k]$, v_i is an ancestor of v_j in T_H if and only if v'_i is an ancestor of v'_j in S_k, and
- for any pair of vertices v'_i and v'_j such that v_i is the parent of v_j in T_H, if $x \neq v'_i$ is on the path from v'_i to v'_j in S_k, then for every $w \in \mathsf{ancestors}_{T_t}(x) \cap \{v'_1, \dots, v'_k\}$, xw and v'_jw have the same adjacency type in H,

(2) or there exists a downward path P_u from u such that $|V(P_u) \cap X| \leqslant k - 1$.

The lemma indeed follows because only the second outcome of $\mathcal{P}_h$ can hold since $G_t[X]$ is H-free. We now prove $\mathcal{P}_k$ by induction on k.

Consider the base case $k = 1$. If there is some descendant $v \in X$ of u in T_t, then add a shortest path from u to such a vertex v. This path is the tree S_1 and set $v'_1 := v$, which satisfies Disjunct (1). Otherwise, u has no descendant in X, and any downward path P_u from u satisfies Disjunct (2).

Suppose now that $\mathcal{P}_k$ holds for some $k \in [h-1]$. Let us prove $\mathcal{P}_{k+1}$. If Disjunct (2) of $\mathcal{P}_k$ holds, then the same downward path P_u from u satisfies Disjunct (2) of $\mathcal{P}_{k+1}$. We thus assume that Disjunct (1) of $\mathcal{P}_k$ holds. Since $G_t[X \cap V(S_k)]$ is isomorphic to $H[\{v_1, \dots, v_k\}]$, some vertex $v'_i \in X \cap V(S_k)$ is such that v_i is the parent of v_{k+1} in T_H with $i \in [k]$.

From the definition of tree representations, for every $v \in V(H)$, $N_{\mathcal{T}(H)}(v) \cap \mathsf{ancestors}_{T_H}(v)$ is a clique in $\mathcal{T}(H)$, and in particular has size at most t (indeed v is fully adjacent to this clique in $\mathcal{T}(H)$). By construction of G_t, there is a downward path Q in T_t from v'_i such that for every $x \in V(Q) \setminus \{v'_i\}$ and for every $j \in [k]$ with $v_j \in \mathsf{ancestors}_{T_H}(v_{k+1})$, xv'_j (in G_t) and $v_{k+1}v_j$ (in H) have the same adjacency type. In particular, for any vertex w of $V(Q) \setminus \{v'_i\}$, the trigraph $G_t[\{v'_1, \dots, v'_k, w\}]$ is isomorphic to $H[\{v_1, \dots, v_{k+1}\}]$. Moreover, $V(Q) \setminus \{v'_i\}$ does not contain any v'_j with $j \in [k]$, because of the sibling condition in T_H, together with the third condition of Disjunct (2).

If $V(Q) \setminus \{v'_i\}$ contains no vertex of X, then we build the downward path P_u by taking the union of Q and the (unique) path from u to v'_i in S_k. Vertices from $V(P_u) \cap X$ are all in S_k, so $V(P_u) \cap X$ contains at most k vertices.

Otherwise, $V(Q) \setminus \{v'_i\}$ contains some vertex in X. Let Q' be the shortest subpath of Q from v'_i to some vertex $w \in X \setminus \{v'_i\}$. We set $v'_{k+1} := w$, and we build the tree S_{k+1} by adding Q' to S_k. One can finally observe that S_{k+1} satisfies the three conditions of Disjunct (2). □

We recall that W_t is the real graph of G_t, and $\mathcal{C}_t$, the class of finite induced subgraphs of W_t. It is easy to see that $\mathcal{C}_t$ has unbounded treewidth.

Observation 6. *$W_t[L_{\leqslant i}]$ has treewidth at least i. In particular, $\mathcal{C}_t$ has unbounded treewidth.*

Proof. The $i+1$ pairwise disjoint connected sets $L_0, L_1, \ldots, L_i$ form a bramble in $W_t[L_{\leqslant i}]$. Hence the treewidth of $W_t[L_{\leqslant i}]$ is at least $(i+1) - 1 = i$. □

We will now prove that for any (finite) H of treewidth at most t, the finite H-free induced subgraphs of W_t have treewidth upper bounded by a function of t and $|V(H)|$. For that, we use the functional equivalence of treewidth and separation number, and we will rely on the following separator.

Observation 7. *For every $u \in V(G_t)$, let u^- be the vertex to the immediate left (in their layer) of the leftmost child of u, and u^+, the vertex to the immediate right of the rightmost child of u. For any downward paths P_{u^-}, P_{u^+} from u^-, u^+, respectively, the set*

$$S := N^{\uparrow}[u] \cup N^{\uparrow}[u^-] \cup N^{\uparrow}[u^+] \cup P_{u^-} \cup P_{u^+}$$

separates in $\mathcal{T}(G_t)$, and hence in W_t, the descendants of u (as well as some descendants of u^- and u^+) from the rest of the graph. (If one u' of $\{u^-, u^+\}$ is not defined, simply remove $N^{\uparrow}[u'] \cup P_{u'}$ from the definition of S.)

Proof. By Observation 3, $\mathcal{T}(G_t)$, and hence W_t, are upward-restricted. Therefore, for any $u' \in \{u, u^-, u^+\}$ and any descendant w of u',

$$N^{\uparrow}[w] \cap \mathsf{ancestors}_{T_t}(u') \subseteq N^{\uparrow}[u'] \subseteq S. \tag{1}$$

Let D be the connected component of $\mathcal{T}(G_t) - S$ containing the descendants of u. Let xy be an edge with exactly one endpoint in D. If x and y are both descendants of u^- or of u^+, then one of x and y is in $P_{u^-} \cup P_{u^+} \subseteq S$. Otherwise exactly one of x, y is a descendant of u, u^-, or u^+, and the other is in S by Inclusion (1). □

The next lemma locates an appropriate node u on which to apply Observation 7.

Lemma 4.2. *For every $n \geqslant 8$ and every n-vertex induced subgraph G of $\mathcal{T}(G_t)$ (hence of W_t), there is some u of T_t such that there are at least $\frac{n}{8 \cdot 3^{t+1}}$ and at most $\frac{n}{4}$ descendants of u that are in $V(G)$, and for every node $v \in V(T_t)$ on the layer of u, v also has at most $\frac{n}{4}$ descendants in $V(G)$.*

Proof. We start from the root of T_t, and move down along edges of T_t according to the following rule. If the current node has at most $n/4$ vertices of G among its descendants in T_t, stop. Otherwise, move to a child of the current node with a largest number of descendants within $V(G)$, that is, at least $\left(\frac{n}{4}-1\right)/3^{t+1} \geqslant \frac{n}{8 \cdot 3^{t+1}}$, by Observation 5. Thus, the node we stop at has the required property. □

Graph W_t is a spanning subgraph of $\mathcal{T}(G_t)$. In particular, for every $X \subseteq V(G_t) = V(W_t)$, we have $\mathrm{tw}(W_t[X]) \leqslant \mathrm{tw}(\mathcal{T}(G_t)[X])$. We can now conclude.

Lemma 4.3. *For every graph H of treewidth at most t and every finite H-free induced subgraph G of W_t, the treewidth of G is at most $15 \cdot \frac{\log(2/3)}{\log \alpha} \cdot (3t + 2|V(H)| + 1)$ with $\alpha := 1 - \frac{1}{8 \cdot 3^{t+1}}$.*

Proof. Let $X \subseteq V(W_t)$ be the finite set such that $G = W_t[X]$. Let H' be a chordal trigraph whose real graph is H, and whose total graph is a chordal graph of clique number at most $t+1$. As H is not an induced subgraph of $W_t[X]$, H' is not an induced subtrigraph of $G_t[X]$.

If $n := |V(G)| < 8$, we are done. Otherwise, let u be the node of T_t obtained by applying Lemma 4.2 to G. Let u^- be the vertex to the immediate left (in their layer) of the leftmost child of u, and u^+, the vertex to the immediate right of the rightmost child of u. Let P_{u^-}, P_{u^+} be the downward paths given by Lemma 4.1 for H', X, u^- and H', X, u^+, respectively. Let

$$S := (N^{\uparrow}[u] \cup N^{\uparrow}[u^-] \cup N^{\uparrow}[u^+] \cup P_{u^-} \cup P_{u^+}) \cap X.$$

In the case when u^- (resp. u^+) does not exist, simply remove $P_{u^-} \cup N^{\uparrow}[u^-]$ (resp. $P_{u^+} \cup N^{\uparrow}[u^+]$) from S. We have

$$|S| \leqslant 3(t+1) + 2(|V(H')| - 1) = 3t + 2|V(H)| + 1$$

by Observation 2 and Lemma 4.1.

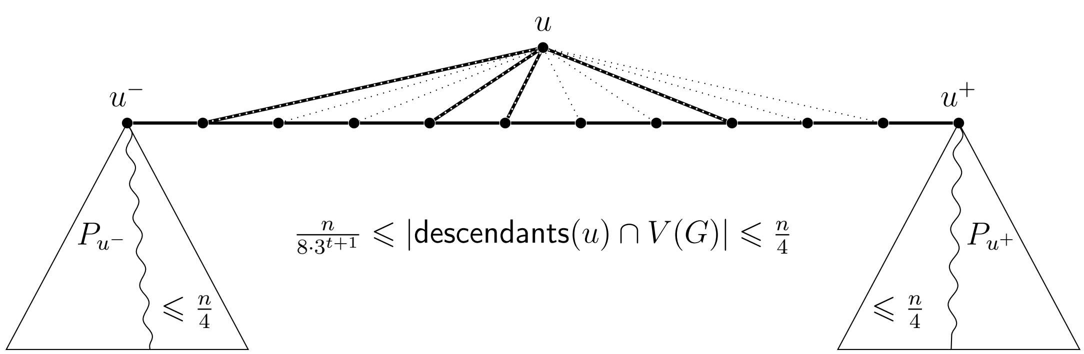

Figure 6: Lower bound of $\frac{n}{8 \cdot 3^{t+1}}$ and upper bound of $\frac{3n}{4}$ in the number of vertices of G below the separator S.

By Observation 7, the descendants of u, and part of the descendants of u^- and of u^+, are disconnected from the other vertices of $G - S$. By Lemma 4.2, S is an α-balanced separator with $\alpha := 1 - \frac{1}{8 \cdot 3^{t+1}}$. Indeed there are at least $\frac{n}{8 \cdot 3^{t+1}}$ vertices of $G - S$ *under* S (i.e., within the descendants of u, u^-, or u^+), and at most $\frac{n}{4} + \frac{n}{4} + \frac{n}{4} = \frac{3n}{4}$. The latter holds because u^-

and u^+ are below the layer of u so, in particular, they have at most $n/4$ descendants in X; see Figure 6.

Thus G (and its subgraphs) has an α-balanced separator of size at most $3t+2|V(H)|+1$, and we conclude by Lemma 2.3. □

Lemma 4.3 and Observation 6 establish Theorem 1.2, and hence Theorem 1.1.

5 Key lemmas on abstract layered wheels

We finish this paper by proving Theorems 1.5 and 1.6. This essentially requires adapting Lemma 4.2, which uses that the layered wheels $\mathcal{T}(G_t)$ and W_t are (3^{t+1}-)bounded, an assumption which we want to drop.

Lemma 5.1. *For every layered wheel W with rooted tree T, for every integer $n \geqslant 8$ and n-vertex induced subgraph G of W, either*

- *there is some $u \in V(T)$ such that there are at least $\frac{n}{16}$ and at most $\frac{n}{4}$ descendants of u that are in $V(G)$, and for every node $v \in V(T)$ on the layer of u, v also has at most $\frac{n}{4}$ descendants in $V(G)$, or*
- *there is some $u \in V(T)$ and a child u^+ of u such that the total number of descendants in $V(G)$ of all children of u both to the right of the leftmost child of u and to the left of u^+ is at least $\frac{n}{16}$ and at most $\frac{n}{8}$.*

Proof. Start at the root of T, and while the number of descendants in $V(G)$ of the current node w is at least $\frac{n}{4}$ (first stopping condition) move to a child of w with the largest number of descendants in $V(G)$ if this number is at least $\frac{n}{16}$, and stop if none exists (second stopping condition).

If the while loop stops because of the first condition, the first item of the disjunction holds. If the while loop stops because of the second condition, let $u \in V(T)$ be the vertex we arrived at. Number $u_1, u_2, \ldots, u_h$ the children of u from left to right. Let $p \in [h]$ be the smallest index such that the total number of descendants in $V(G)$ of vertices in $\{u_2, \ldots, u_p\}$ is at least $\frac{n}{16}$ and at most $\frac{n}{8}$. By construction, for every child u_i of u the number of descendants in $V(G)$ of u_i is less than $\frac{n}{16}$. Thus, as the number of descendants in $V(G)$ of u is at least $\frac{n}{4}$, p is well-defined and at most equal to $h-1$. Vertex u^+ is then defined as u_{p+1}, and the second item holds. □

We are now equipped for proving Theorem 1.5, which we recall.

Theorem 1.5. *For every neat, upward-restricted layered wheel W, every class of finite graphs satisfying the bounded-branch property in W has bounded treewidth.*

Proof. Fix any neat, t-upward-restricted layered wheel W with rooted tree T. Let $\mathcal{C}$ be a class of finite induced subgraphs of W satisfying the bounded-branch property for some integer h. Fix any $G \in \mathcal{C}$. Let $u \in V(T)$ be the vertex given by applying Lemma 5.1 to T, G.

If u satisfies the first item of Lemma 5.1, we note that Observation 7 only uses the fact that the layered wheel is neat (for the downward paths to intersect every layer below their starting point) and upward-restricted. We then conclude, following the proof of Lemma 4.3, that G has a $\frac{15}{16}$-balanced separator of size at most $3(t+1)+2h$.

Let us thus assume that u and its child u^+ satisfy the second item of Lemma 5.1. Let u^- be the leftmost child of u. Let P_{u^-} (resp. P_{u^+}) be any downward path in T from u^- (resp. u^+) such that $|V(P_{u^-}) \cap V(G)| \leqslant h$ (resp. $|V(P_{u^+}) \cap V(G)| \leqslant h$). As the layered wheel is neat, these downward paths are infinite. Similarly to Observation 7,

$$S := N^{\uparrow}[u] \cup N^{\uparrow}[u^-] \cup N^{\uparrow}[u^+] \cup P_{u^-} \cup P_{u^+} = N^{\uparrow}[u] \cup P_{u^-} \cup P_{u^+}$$

separates in G the descendants of children between u^- and u^+ in the left-to-right order to the rest of the graph. Moreover, this number of descendants is at least $\frac{n}{16}$ and at most $\frac{n}{8} + \frac{2n}{16} = \frac{n}{4}$. This means that G has a $\frac{15}{16}$-balanced separator of size at most $t + 1 + 2h$.

We conclude by Lemma 2.3 that $\mathrm{tw}(G) \leqslant 15 \cdot \frac{\log(2/3)}{\log(15/16)} \cdot (3(t+1) + 2h)$. □

We can further adapt the proof of Theorem 1.5 to show that finite induced subgraphs of neat, upward-restricted layered wheels have treewidth at most logarithmic in their number of vertices.

Theorem 1.6. *For every neat, t-upward-restricted layered wheel W on rooted tree T, every n-vertex induced subgraph of W has treewidth at most*

$$15 \cdot \frac{\log(2/3)}{\log(15/16)} \cdot (3(t+1) + 2h \log n) = O(t + \log n),$$

for some finite integer h depending only on T.

Proof. In the proof of Theorem 1.5, every time a downward path from a specific vertex v is added to a balanced separator, choose a downward path from v that contains as few vertices of G as possible. By definition of a layered wheel, there is an integer h such that T has no path of length h only consisting of degree-2 vertices. Thus, from any node v with fewer than n descendants in $V(G)$, there is a downward path from v containing at most $h \log n$ vertices of G, by greedily choosing a branch with fewer descendants in $V(G)$. □

On the other hand, the class of the finite induced subgraphs of any proper, bounded layered wheel has at least logarithmic treewidth.

Observation 8. *For every proper, t-bounded layered wheel W with rooted tree T, and every positive integer n', W admits an n-vertex induced subgraph G such that $n \geqslant n'$ and*

$$tw(G) \geqslant \frac{\log n}{\log t} - 1.$$

Proof. We can assume that $t \geqslant 2$, as every proper 1-bounded layered wheel is a clique. Let i be the smallest integer such that the subgraph G induced by the first $i + 1$ layers of W has at least n' vertices. By Observation 6, which only uses the fact that the layered wheel is proper, $\mathrm{tw}(G) \geqslant i$. As the layered wheel is t-bounded, $n \leqslant t^{i+1}$, so $i \geqslant \log n / \log t - 1$. □

By Theorem 1.6 and Observation 8 the class of finite induced subgraphs of any proper, neat, bounded, upward-restricted layered wheel has logarithmic treewidth.

References

[1] Pierre Aboulker, Isolde Adler, Eun Jung Kim, Ni Luh Dewi Sintiari, and Nicolas Trotignon. On the tree-width of even-hole-free graphs. *European Journal of Combinatorics*, 98:103394, 2021.

[2] Tara Abrishami, Bogdan Alecu, Maria Chudnovsky, Cemil Dibek, Peter Gartland, Sepehr Hajebi, Daniel Lokshtanov, Paweł Rzążewski, Sophie Spirkl, and Kristina Vušković. Induced subgraphs and tree decompositions I–XVIII, 2020-.

[3] Shinwoo An, Eunjin Oh, and Jie Xue. Sparse outerstring graphs have logarithmic treewidth. In Timothy M. Chan, Johannes Fischer, John Iacono, and Grzegorz Herman, editors, *32nd Annual European Symposium on Algorithms, ESA 2024, September 2-4, 2024, Royal Holloway, London, United Kingdom*, volume 308 of *LIPIcs*, pages 10:1–10:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

[4] Marthe Bonamy, Édouard 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. *Journal of Combinatorial Theory, Series B*, 167:215–249, 2024.

[5] Édouard Bonnet. Sparse induced subgraphs of large treewidth. *CoRR*, abs/2405.13797, 2024.

[6] Édouard Bonnet and Julien Duron. Stretch-width. 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 8:1–8:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.

[7] Édouard Bonnet, Eun Jung Kim, Amadeus Reinald, and Stéphan Thomassé. Twin-width VI: the lens of contraction sequences. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 1036–1056. SIAM, 2022.

[8] Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. *J. ACM*, 69(1):3:1–3:46, 2022.

[9] Maria Chudnovsky, Tara Abrishami, Sepehr Hajebi, and Sophie Spirkl. Induced subgraphs and tree decompositions III. Three-path-configurations and logarithmic treewidth. *Advances in Combinatorics*, page 1–29, September 2022.

[10] Maria Chudnovsky, David Fischer, Sepehr Hajebi, Sophie Spirkl, and Bartosz Walczak. Treewidth and outerstring graphs, in preparation, 2025.

[11] Maria Chudnovsky, Peter Gartland, Sepehr Hajebi, Daniel Lokshtanov, and Sophie Spirkl. Induced subgraphs and tree decompositions XV. Even-hole-free graphs with bounded clique number have logarithmic treewidth. *CoRR*, abs/2402.14211, 2024.

[12] Maria Chudnovsky, Sepehr Hajebi, and Sophie Spirkl. Induced subgraphs and tree decompositions XVI. Complete bipartite induced minors. *CoRR*, abs/2410.16495, 2024.

[13] Maria Chudnovsky and Nicolas Trotignon. On treewidth and maximum cliques. *CoRR*, abs/2405.07471, 2024.

[14] Daniel Cocks. t-sails and sparse hereditary classes of unbounded tree-width. *European Journal of Combinatorics*, 122:104005, 2024.

[15] James Davies. Oberwolfach report 1/2022. 2022.

[16] Zdenek Dvorák and Sergey Norin. Treewidth of graphs with balanced separations. *J. Comb. Theory B*, 137:137–144, 2019.

[17] António Girão, Freddie Illingworth, Emil Powierski, Michael Savery, Alex Scott, Youri Tamitegama, and Jane Tan. Induced subgraphs of induced subgraphs of large chromatic number. *Combinatorica*, 44:37–62, 2024.

[18] Frank Gurski and Egon Wanke. The tree-width of clique-width bounded graphs without $K_{n,\ n}$. In Ulrik Brandes and Dorothea Wagner, editors, *Graph-Theoretic Concepts in Computer Science, 26th International Workshop, WG 2000, Konstanz, Germany, June 15-17, 2000, Proceedings*, volume 1928 of *Lecture Notes in Computer Science*, pages 196–205. Springer, 2000.

[19] Sepehr Hajebi. Chordal graphs, even-hole-free graphs and sparse obstructions to bounded treewidth. *CoRR*, abs/2401.01299, 2024.

[20] Michał Pilipczuk. Problems parameterized by treewidth tractable in single exponential time: A logical approach. In Filip Murlak and Piotr Sankowski, editors, *Mathematical Foundations of Computer Science 2011 - 36th International Symposium, MFCS 2011, Warsaw, Poland, August 22-26, 2011. Proceedings*, volume 6907 of *Lecture Notes in Computer Science*, pages 520–531. Springer, 2011.

[21] Andrei Cosmin Pohoata. *Unavoidable induced subgraphs of large graphs*. Senior theses, Princeton University, 2014.

[22] Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986.

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

[24] Frank W. Sinden. Topology of thin film RC circuits. *Bell System Technical Journal*, 45(9):1639–1662, 1966.

[25] Ni Luh Dewi Sintiari and Nicolas Trotignon. (Theta, triangle)-free and (even hole, K_4)-free graphs—Part 1: Layered wheels. *Journal of Graph Theory*, 97(4):475–509, 2021.