Graph Minor Hierarchies

Reinhard Diestel & Daniela Kühn

12.10.2002

We generalize tree-decompositions to decompositions modelled on graphs other than trees, and study how such more general decompositions might be used to establish structural complexity hierarchies of graph properties.

Introduction

In their work on graph minors, Robertson and Seymour begin by describing graphs whose structure is particularly simple, graphs that look roughly like thickened paths [8, 3]. They say that such graphs have small 'path-width'. Of course, not every graph looks roughly like a thickened path, but it is possible to describe those that do not: every graph of large path-width contains a particular such graph as a minor, a large (binary) tree.

As a natural next step, Robertson and Seymour consider the graphs that look roughly like thickened trees, the graphs of small 'tree-width' [9; 5, 2]. As before, not all graphs have small tree-width, but those that do not must all contain a particular kind of graph of large tree-width as a minor: a large grid.

And why stop here? In the same way as above one might extend the class of graphs described so far by including those that have small 'grid-width', then continue with the minor-minimal graphs of 'unbounded grid-width' if those can be determined, and so on.

Since the proof of their graph minor theorem takes a different turn for the graphs of large tree-width—proceeding inductively up the genus of the graphs considered, see [2; Ch. 12.5]—Robertson and Seymour have no need to pursue this emerging hierarchy further. However, it seems worthwhile attempting to do so. For every element of a class of graphs of 'bounded $\mathcal{H}$-width', where $\mathcal{H}$ is some previously described class of simpler graphs, will inherit some of the properties of the graphs in $\mathcal{H}$. (For example, graphs of bounded tree-width inherit many of the algorithmic advantages of trees, their well-quasi-ordering, and so on.) Thus, given a graph property shared by the graphs at the lower levels of this hierarchy, it may well be worth asking how far up the hierarchy it holds—once the hierarchy has been established.

This paper studies how to set up such hierarchies of graph properties. It turns out that finding suitable definitions of $\mathcal{H}$-decompositions and $\mathcal{H}$-width for more general classes $\mathcal{H}$ than paths and trees is a more delicate problem than might be anticipated. We suggest a number of ways in which this could be done, but there seems to be no general rule of how best to define the concepts involved, once and for all. Rather, they will depend on the class of properties to be studied, and even then may have to be fine-tuned as part of that study.

All the graphs considered in this paper are simple and finite. The notation adopted is that of [2], where also any standard theorems referred to below can be found.

1. Graph properties

If we wish to investigate how some graphs can be modelled on others (eg., the former having 'bounded width' in terms of the latter), our primary objects of study will not be individual graphs but classes of graphs. In this section, we introduce such classes (to be called 'properties') and see how they can be compared in terms of the graph minor relation and related orderings.

A *graph property* in this paper is taken to mean an infinite class of isomorphism types of finite graphs. Thus, every graph property we consider contains arbitrarily large (unlabelled) graphs.

Given two graphs G_1 and G_2, we write $G_1 \sqsubseteq G_2$ and call G_1 a *preminor* of G_2 if the vertices $v \in G_1$ can be mapped to disjoint sets $X_v \subseteq V(G_2)$ so that G_2 contains an X_v–X_w edge whenever vw is an edge of G_1. Note that the sets X_v, the *branch sets* of this mapping, need not be connected; if they are, then G_1 is a *minor* of G_2 and we write $G_1 \preccurlyeq G_2$. If all the branch sets can be chosen with no more than k vertices, we also write $G_1 \sqsubseteq_k G_2$ and $G_1 \preccurlyeq_k G_2$, respectively.

Given two graph properties $\mathcal{P}_1$ and $\mathcal{P}_2$, we write $\mathcal{P}_1 \sqsubseteq \mathcal{P}_2$ if for every $G_1 \in \mathcal{P}_1$ there is a $G_2 \in \mathcal{P}_2$ such that $G_1 \sqsubseteq G_2$, and similarly for $\preccurlyeq$, $\sqsubseteq_k$ and $\preccurlyeq_k$. More generally if there exists a $k \in \mathbb{N}$ such that $\mathcal{P}_1 \sqsubseteq_k \mathcal{P}_2$ (resp. $\mathcal{P}_1 \preccurlyeq_k \mathcal{P}_2$), we also write $\mathcal{P}_1 \sqsubseteq_* \mathcal{P}_2$ (resp. $\mathcal{P}_1 \preccurlyeq_* \mathcal{P}_2$) and say that the graphs in $\mathcal{P}_1$ are *bounded (pre-)minors* of those in $\mathcal{P}_2$. Note that, unlike $\sqsubseteq_k$ and $\preccurlyeq_k$, the relations $\sqsubseteq_*$ and $\preccurlyeq_*$ are transitive. Finally, if every graph in $\mathcal{P}_1$ is a subgraph of some graph in $\mathcal{P}_2$ we write $\mathcal{P}_1 \Subset \mathcal{P}_2$; note that $\Subset$ is equivalent to $\sqsubseteq_1$ and to $\preccurlyeq_1$.

The relations between graph properties that we shall mainly be interested in are the three transitive relations $\preccurlyeq$, $\preccurlyeq_*$, $\sqsubseteq_*$ and variants of these such as compositions with topological minors. Since $\preccurlyeq_*$ is a refinement (a subset) of both $\preccurlyeq$ and $\sqsubseteq_*$, and these are refinements of $\sqsubseteq$ (Fig. 1), a statement about $\preccurlyeq_*$ or $\sqsubseteq$ will often imply analogous statements about the other three relations, which we shall not always mention explicitly.

$$
\begin{tikzcd}
 & \sqsubseteq & \\
\sqsubseteq_* \arrow[ur] & & \preccurlyeq \arrow[ul] \\
 & \preccurlyeq_* \arrow[ul] \arrow[ur] &
\end{tikzcd}
$$

FIGURE 1. Our four minor relations between graph properties

Let $\leqslant$ denote any reflexive and transitive relation between graph properties (such as $\sqsubseteq$, $\sqsubseteq_*$, $\preccurlyeq$ or $\preccurlyeq_*$). We shall call $\mathcal{P}_1$ and $\mathcal{P}_2$ *equivalent* (wrt. $\leqslant$) and write $\mathcal{P}_1 \sim \mathcal{P}_2$ if both $\mathcal{P}_1 \leqslant \mathcal{P}_2$ and $\mathcal{P}_1 \geqslant \mathcal{P}_2$. Then $\leqslant$ induces a partial ordering on the equivalence classes of graph properties thus defined. We write $\mathcal{P}_1 < \mathcal{P}_2$ (ie. $\sqsubset$, $\sqsubset_*$, $\prec$, $\prec_*$ etc.) to express that $\mathcal{P}_1 \leqslant \mathcal{P}_2$ but $\mathcal{P}_1 \not\geqslant \mathcal{P}_2$, ie. that while $\mathcal{P}_1 \leqslant \mathcal{P}_2$ the two properties are not equivalent. Informally, we shall think of $\mathcal{P}_1$ as lying *above* $\mathcal{P}_2$ if $\mathcal{P}_1 \geqslant \mathcal{P}_2$.

A graph property is *sparse* if the average degrees of its graphs are bounded above by some constant; it is *dense* if its graphs have average degrees bounded below by some linear function of their order. Of particular interest among the sparse properties are those properties $\mathcal{P}$ such that every $\mathcal{P}' \Subset \mathcal{P}$ is also sparse (or equivalently, such that $\mathcal{P}' := \{ G' \subseteq G \mid G \in \mathcal{P} \}$ is sparse). These properties can be described as follows.

Proposition 1.1. *The following assertions are equivalent for graph properties $\mathcal{P}$:*

(i) *$\mathcal{P}' := \{ G' \subseteq G \mid G \in \mathcal{P} \}$ is sparse;*

(ii) *$\mathcal{P}$ has bounded arboricity;*

(iii) *$\mathcal{P}$ has bounded colouring number;*

(iv) *the graphs in $\mathcal{P}$ admit orientations with bounded out-degrees.*

Proof. (i)$\leftrightarrow$(ii): This equivalence follows from Nash-Williams's characterization of the graphs of arboricity at most k for given k; see [2; Thm 3.5.4].

(i)$\rightarrow$(iii) is an immediate consequence of the easy fact that the colouring number of a graph is exactly one greater than the greatest minimum degree of its subgraphs [2; Prop 5.2.2].

(iii)$\rightarrow$(iv): A graph of colouring number k has a vertex enumeration in which each vertex is preceded by fewer than k of its neighbours. The orientation from later to earlier vertices thus has out-degrees $< k$.

(iv)$\rightarrow$(i): If a graph has an orientation with out-degrees at most k, it has at most k times as many edges as vertices and hence an average degree of at most $2k$. $\square$

Although most properties we shall consider will be sparse, a few observations about dense properties may serve to put our later investigations into perspective. As we shall see, the complete graphs and the complete bipartite graphs will play a special role as prototype dense properties. Let us denote them as

$$\mathcal{K} := \{ K^n \mid n \in \mathbb{N} \} \quad \text{and} \quad \mathcal{K}^2 := \{ K_{n,n} \mid n \in \mathbb{N} \}.$$

Not surprisingly in the context of minors, there is only one dense property up to equivalence (with respect to any of our four relations), which lies above all other properties:

Proposition 1.2. *If $\mathcal{Q}$ is any dense property, then $\mathcal{P} \preccurlyeq_* \mathcal{Q}$ for every graph property $\mathcal{P}$. In particular, every two dense properties are equivalent.*

Proof. As clearly $\mathcal{P} \preccurlyeq_* \mathcal{K}$, it suffices to show that $\mathcal{K} \preccurlyeq_* \mathcal{Q}$. We show that, in fact, $\mathcal{K} \preccurlyeq_2 \mathcal{Q}$. Since $\mathcal{Q}$ is dense, there exists an $\epsilon > 0$ such that every graph $G \in \mathcal{Q}$ has at least $\epsilon|G|^2$ edges. The Erdős-Stone theorem therefore implies that, given $n \in \mathbb{N}$, every large enough graph in $\mathcal{Q}$ has a $K_{n,n}$ subgraph. Contracting every edge in some maximal matching in this $K_{n,n}$ we obtain a K^n minor as desired. □

For the two bounded-minor relations $\sqsubseteq_*$ and $\preccurlyeq_*$, Proposition 1.2 has a pretty counterpart. Ideally, one would perhaps like to prove as its obvious converse that *only* dense properties can be equivalent to $\mathcal{K}$, ie. that the equivalence class of $\mathcal{K}$ is exactly the class of dense properties. But that is trivially false: by adding enough isolated vertices to its graphs we can make any dense property sparse while keeping it equivalent to the original property. However, every property equivalent to a dense property (and hence to $\mathcal{K}$) contains a dense core, and thus arises from a dense property in this way:

Proposition 1.3. *Under each of the relations $\sqsubseteq_*$ and $\preccurlyeq_*$, a property $\mathcal{P}$ is equivalent to some dense property (and hence to $\mathcal{K}$) if and only if $\mathcal{K}^2 \Subset \mathcal{P}$.*

Proof. The 'if' implication is clear from Proposition 1.2. For 'only if', it suffices to show that every property $\mathcal{P} \sqsupseteq_* \mathcal{K}$ satisfies $\mathcal{K}^2 \Subset \mathcal{P}$. Let $k \in \mathbb{N}$ be such that $\mathcal{P} \sqsupseteq_k \mathcal{K}$, and let $s \in \mathbb{N}$ be given. We know that for every $r \in \mathbb{N}$ there is a graph $G \in \mathcal{P}$ that contains r disjoint sets of at most k vertices each and at least one edge between every two of these sets; let us show that the subgraph $H \subseteq G$ induced by these sets contains $K_{s,s}$ if r is large enough.

By the Erdős-Stone theorem it suffices to show that H has at least ϵn^2 edges for $n := |H| \to \infty$ and some $\epsilon > 0$ depending only on k. As $r \leqslant n \leqslant kr$ and H has at least $\binom{r}{2} \geqslant r^2/4 \geqslant n^2/4k^2$ edges, this is indeed the case. □

Using the lower bound of $c_r n^{2-\frac{2}{r+1}}$ for the extremal function $\mathrm{ex}\,(n, K_{r,r})$ (see [2, p. 152]), one can use Proposition 1.3 to show that properties inequivalent to $\mathcal{K}$ can have almost linear average degrees:

Corollary 1.4. *For every constant c and every $\epsilon > 0$ there exists a (non-dense) graph property $\mathcal{P} \not\sqsupseteq_* \mathcal{K}$ whose elements G all have average degrees at least $c|G|^{1-\epsilon}$.* □

By a well-known theorem of Mader [2], every property $\mathcal{P}$ of unbounded average degree satisfies $\mathcal{K} \preccurlyeq \mathcal{P}$. The above lower bound for $\mathrm{ex}\,(n, K_{r,r})$ therefore shows that $\mathcal{K} \preccurlyeq \mathcal{P}$ cannot imply $\mathcal{K}^2 \Subset \mathcal{P}$, so Proposition 1.3 does not extend to $\preccurlyeq$ or to $\sqsubseteq$.

On the other hand, the unbounded minor relations $\preccurlyeq$ and $\sqsubseteq$ differ from their bounded counterparts also in a very pleasant way:

Lemma 1.5. *All strictly descending chains $\mathcal{P}_1 \succ \mathcal{P}_2 \succ \ldots$ and $\mathcal{P}_1 \sqsupset \mathcal{P}_2 \sqsupset \ldots$ are finite.*

Proof. We prove the lemma for the minor relation; the proof for preminors is exactly the same. Let $\mathcal{P}_1 \succ \mathcal{P}_2 \succ \ldots$ be a strictly descending chain of graph properties. Since $\mathcal{P}_1 \not\preccurlyeq \mathcal{P}_2$, there is a graph $G_1 \in \mathcal{P}_1$ that is not a minor of any graph in $\mathcal{P}_2$. Then G_1 is not even a minor of a graph in $\mathcal{P}_i$ for any $i \geqslant 2$, because $\mathcal{P}_i \preccurlyeq \mathcal{P}_2$. Similarly, $\mathcal{P}_2$ contains a graph G_2 that is not a minor of any graph in $\mathcal{P}_i$ with $i \geqslant 3$, and so on. We thus obtain a sequence $G_1, G_2, \ldots$ with $G_i \in \mathcal{P}_i$ for each i and such that $G_i \not\preccurlyeq G_j$ whenever $i < j$. By the Robertson-Seymour graph minor theorem every such sequence is finite, and hence so is our chain $\mathcal{P}_1 \succ \mathcal{P}_2 \succ \ldots$. □

(The fact that the finite graphs are well-quasi-ordered under the preminor relation follows from the graph minor theorem, but is also not difficult to show directly.)

Theorem 3.5 below will imply that for both $\preccurlyeq_*$ and $\sqsubseteq_*$ there are indeed infinite strictly descending chains. In the case of $\preccurlyeq_*$, one can obtain such chains simply by subdividing, and easily write down explicit examples: the classes $\mathcal{P}_n := \{ G_n^i \mid i \in \mathbb{N} \}$, for instance, where G_n^i is obtained from an i-star by taking for each leaf a cycle of length $2^{2^n i}$ and identifying the leaf with a vertex on that cycle, form an infinite strictly decreasing chain as $n = 1, 2, \ldots$.

Regarding Lemma 1.5, one may ask whether the graph properties are in fact well-quasi-ordered by the relation $\preccurlyeq$, ie. whether there are also no infinite $\preccurlyeq$-antichains of properties. This is an open problem. If the finite graphs are better-quasi-ordered as minors (which all the experts seem to believe) then the graph properties would likewise be better-quasi-ordered (and hence well-quasi-ordered). Incidentally, Lemma 1.5 implies the graph minor theorem just as easily as the other way round: if $G_0, G_1, \ldots$ is an infinite sequence of graphs such that $G_i \not\preccurlyeq G_j$ whenever $i < j$, then its tails $\{ G_i, G_{i+1}, \ldots \}$ form an infinite descending $\succ$-chain of graph properties.

2. Divisibility of properties

This section briefly addresses a fundamental problem concerning graph properties that will become relevant later to the hierarchies we seek to establish, but is at this point included more for its own interest. The section may be skipped without loss at first reading. Throughout this section, $\leqslant$ stands for any of our three relations $\preccurlyeq$, $\preccurlyeq_*$, $\sqsubseteq_*$, and similar observations hold for related relations between graph properties.

Given a graph property $\mathcal{P}$, we can obtain numerous equivalent properties just by 'adding junk': for every property $\mathcal{P}' < \mathcal{P}$, the property $\mathcal{P} \cup \mathcal{P}'$ is equivalent to $\mathcal{P}$. This process is not easily reversible: if we are given $\mathcal{P} \cup \mathcal{P}'$ as a single property, we may not readily be able to identify and discard its 'inessential' part $\mathcal{P}'$. So it seems that properties not containing such 'junk' are particularly interesting representatives of their equivalence types.

To make this precise, let us call a property $\mathcal{P}$ *lean* if every property $\mathcal{P}' \subseteq \mathcal{P}$ is equivalent to $\mathcal{P}$. (Recall that $\mathcal{P}'$, by our definition of a graph property, contains arbitrarily large graphs.) The stars, for example, form a lean property, and so do the paths. Note that if $\mathcal{P}$ is equivalent to some lean property $\mathcal{Q}$ it also has a lean subset. For example, when $\leqslant$ means $\preccurlyeq_*$, choose k so that $\mathcal{Q} \preccurlyeq_k \mathcal{P}$, and for each $H \in \mathcal{Q}$ choose $G = G(H) \in \mathcal{P}$ with $H \preccurlyeq_k G$. Then $\mathcal{P} \sim \mathcal{Q} \preccurlyeq_k \mathcal{P}_\mathcal{Q} := \{ G(H) \mid H \in \mathcal{Q} \} \subseteq \mathcal{P}$ (so all these properties are equivalent), and $\mathcal{P}_\mathcal{Q}$ is lean: for any property $\mathcal{P}' \subseteq \mathcal{P}_\mathcal{Q}$ the set $\mathcal{Q}' := \{ H \mid G(H) \in \mathcal{P}' \} \subseteq \mathcal{Q}$ is infinite (like $\mathcal{P}'$) and hence equivalent to $\mathcal{Q}$, so $\mathcal{P} \leqslant \mathcal{Q} \leqslant \mathcal{Q}' \leqslant \mathcal{P}'$ as desired.

So the question arises whether every graph property is equivalent to some lean property, in which case we could take those as their standard representatives. However, this is not the case; for example, for all our three relations the property consisting of the stars *and* the paths is not equivalent to any lean property. More generally, let us call a subset $\mathcal{P}'$ of a property $\mathcal{P}$ *small* if either $\mathcal{P}'$ is finite or $\mathcal{P}' < \mathcal{P}$, and let us call $\mathcal{P}$ *divisible* if it is the union of two small subsets. Thus, the property of stars and paths considered above is divisible, and so is the property consisting of all the stars and the path of length five. Lean properties, on the other hand, are indivisible.

Lemma 2.1. *If $\mathcal{P}$ is indivisible and $\mathcal{P} \leqslant \mathcal{Q}_1 \cup \mathcal{Q}_2$ (where one of the $\mathcal{Q}_i$ may be finite), then $\mathcal{P} \leqslant \mathcal{Q}_1$ or $\mathcal{P} \leqslant \mathcal{Q}_2$.*

Proof. We prove the assertion for $\leqslant$ meaning $\preccurlyeq_*$; the other cases are similar. Choose k so that $\mathcal{P} \preccurlyeq_k \mathcal{Q}_1 \cup \mathcal{Q}_2$. Then we may write $\mathcal{P}$ as $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2$ where $\mathcal{P}_i := \{ G \in \mathcal{P} \mid \exists\, H \in \mathcal{Q}_i \colon G \preccurlyeq_k H \}$. As $\mathcal{P}$ is indivisible, the $\mathcal{P}_i$ cannot both be small. So one of them satisfies $\mathcal{P} \preccurlyeq_* \mathcal{P}_i \preccurlyeq_* \mathcal{Q}_i$, as desired. □

Lemma 2.1 implies that the small subsets of any indivisible property form a set-theoretic ideal: finite unions and subsets of small sets are again small. Furthermore, it suggests that indivisible properties behave like primes in factoring. This is not just a coincidence: when $\leqslant$ stands for unbounded minors or preminors, one can even use Lemma 1.5 to prove that every graph property partitions, uniquely up to equivalence, into a finite set of indivisible subproperties. See [4] for details.

Moreover, unlike leanness, divisibility and indivisibility are invariant under equivalence:

Corollary 2.2. *If $\mathcal{P}$ is indivisible and $\mathcal{P} \sim \mathcal{Q}$ then $\mathcal{Q}$ is indivisible.* □

By Corollary 2.2, only an indivisible property can be equivalent to a lean property. Our general problem thus is as follows:

Problem 2.3. *Is every indivisible property equivalent to some lean property?*

For unbounded minors and preminors, the indivisible properties have a simple structural characterization [4] which implies a positive answer to Problem 2.3; in general we do not know the answer.

Perhaps surprisingly, a property that is indivisible but not lean may well be more natural and simpler to describe than an equivalent lean property. For example, the property of being a tree is indivisible but not lean. It is equivalent to the 'diagonal' lean property $\mathcal{T} = \{ T_n \mid n \in \mathbb{N} \}$, where T_n is the tree of height n in which every vertex not at the nth level has degree n, but $\mathcal{T}$ is less straightforward to describe. (On the other hand, pointing out the equivalence to $\mathcal{T}$ is perhaps the easiest way to prove that the property of being a tree is indivisible.)

3. General graph decompositions

Just as the concept of tree-decomposition provides a way of 'roughly' modelling a given graph G on a tree T, we shall now define a more general concept of modelling a given graph G on another graph H.

Let G, H be graphs. Consider a family $\mathcal{D} = (G_h)_{h \in H}$ of induced subgraphs of G indexed by the vertices of H. We call $\mathcal{D}$ an *H-decomposition* of G (into the *parts* G_h) if

(D1) every vertex of G lies in some G_h; and

(D2) given an edge $e = gg' \in G$, either e lies in some G_h or there exists an edge $hh' \in H$ such that $g \in G_h$ and $g' \in G_{h'}$.

We shall call this decomposition $\mathcal{D}$ *connected* if

(D3) whenever a vertex $g \in G$ lies in $G_{h_1} \cap G_{h_2}$ for some $h_1, h_2 \in H$, there is a path $P = h_1 \dots h_2$ in H such that $g \in G_h$ for every $h \in P$.

For vertices $g \in G$ we write

$$H_g := H[\{ h \mid g \in G_h \}] \tag{1}$$

and call these graphs H_g the *co-parts* of $\mathcal{D}$.

We call $\mathrm{wd}(\mathcal{D}) := \max_{h \in H} |G_h|$ the *width* of the decomposition $\mathcal{D}$, and $\mathrm{sp}(\mathcal{D}) := \max_{v \in G} |H_v|$ its *spread*. The maximum of these two numbers is the *size* of the decomposition $\mathcal{D}$. When $\mathcal{H}$ is a graph property then the *$\mathcal{H}$-width* (resp. *$\mathcal{H}$-size*) of G is the least width (resp. size) of an H-decomposition of G with $H \in \mathcal{H}$. The *connected $\mathcal{H}$-width* (resp. *-size*) of G are defined analogously with respect to connected decompositions.

Let us spend a moment to see how a connected T-decomposition of G, when T is a tree, corresponds to a traditional tree-decomposition. The conditions (D1) and (D3) correspond exactly to the conditions (T1) and (T3) in the standard definition of a tree-decomposition [2]. (D2) however is slightly weaker than (T2): it says that every edge of G is either accommodated in one of the parts *or reflected by an edge of H* (which is not an option in (T2)).

Allowing in (D2) that edges of G may have their ends in different parts ensures that every graph H has the trivial H-decomposition into singletons. This relaxation of (T2) is not only natural, it is essential when H can be an arbitrary graph: without it, a graph with an H-decomposition of width at most k could never have more than $\binom{k}{2}|H|$ edges, and so even the 'H-width' of H itself (the least width of an H-decomposition of H) would be unbounded as H gets large and dense. In order to avoid such anomalities we need to allow edges as in (D2) here, even though they can be avoided when H is a tree.*

It is possible to rewrite the conditions (D1)–(D3) more elegantly (though perhaps less accessibly) in terms of the co-parts H_g. Recall that two subgraphs of a graph H are said to *touch* (in H) if they have a vertex in common or H contains an edge between them. Each of the conditions (D1)–(D3) then is clearly equivalent to the corresponding following condition:

(C1) every H_g is non-empty;

(C2) for every edge $gg' \in G$, the graphs H_g and $H_{g'}$ touch in H;

(C3) every H_g is connected.

Since

$$h \in H_g \quad \Leftrightarrow \quad g \in G_h\,, \tag{2}$$

the parts G_h of a decomposition can be reobtained from their co-parts H_g, ie. the decomposition $(G_h)_{h \in H}$ is uniquely identified by the family $(H_g)_{g \in G}$. We shall often use this fact in that we present a decomposition either as a family of parts or as a family of co-parts, whichever is more convenient.

Note that, by the symmetry in (2), the G_h are in fact obtained from the H_g exactly as those were defined from the G_h in (1): for all $h \in H$, we have $G_h = G[\{ g \mid h \in H_g \}]$. We do not explore this duality further here, but remark that it includes planar duality as a special case: if G and H are dual plane graphs then the H-decomposition of G into its face boundaries (each associated with the vertex of H corresponding to that face) has as co-parts the face boundaries of H associated with the vertices of G, and vice versa. Perhaps some features of planar duality (such as flow-colouring duality) might be extendable along these lines.

Most of this section will be needed to establish some easy technical lemmas about these decompositions. Their straightforward proofs are included for completeness, but the reader is encouraged to quickly verify the assertions

---

* In terms of tree-width it matters little whether edges as in (D2) are allowed or not: any connected T-decomposition $\mathcal{D}$ with T a tree can be turned into a traditional tree-decomposition with parts $\{ G_t \cup G_{t'} \mid tt' \in E(T) \}$, whose width thus exceeds that of $\mathcal{D}$ by no more than a factor of 2.

directly and thus become better acquainted with the two sets of conditions above.

Most importantly, decompositions and the minor relation are naturally compatible as follows. When G and H are two graphs then, intuitively, G having an H-decomposition of small width and spread is a way of expressing that G is 'not much bigger' than H, while $G \preccurlyeq H$ is a way of saying that G is smaller than H. And these two concepts of the size (or complexity) of a graph are indeed compatible in the way such intuition suggests: if $\mathcal{P} \preccurlyeq \mathcal{H}$ then the graphs in $\mathcal{P}$ have bounded connected $\mathcal{H}$-width, and if $\mathcal{P} \preccurlyeq_* \mathcal{H}$ they even have bounded connected $\mathcal{H}$-size. Quantitatively:

Lemma 3.1.

(i) *$G \sqsubseteq_k H$ if and only if G has an H-decomposition of width at most 1 and spread at most k.*

(ii) *$G \preccurlyeq_k H$ if and only if G has a connected H-decomposition of width at most 1 and spread at most k.*

Proof. If $G \sqsubseteq_k H$ then every vertex $g \in G$ has a corresponding branch set $X_g \subseteq V(H)$. The graphs $H_g := H[X_g]$ satisfy (C1) and (C2), and they satisfy (C3) if $G \preccurlyeq_k H$. Thus $(H_g)_{g \in G}$ is the family of co-parts of an H-decomposition of G, which has width 1 since the H_g are disjoint and spread $\leqslant k$ since $|X_g| \leqslant k$ for all g.

Conversely if G has an H-decomposition $\mathcal{D} = (G_h)_{h \in H}$ of width at most 1 then $G \sqsubseteq H$ with branch sets $V(H_g)$ $(g \in G)$: these sets are non-empty by (C1), disjoint by $\mathrm{wd}(\mathcal{D}) \leqslant 1$, connected by (C3) if $\mathcal{D}$ is connected. Moreover, if $gg' \in E(G)$ then H has an H_g–$H_{g'}$ edge by (C2). Finally, if $\mathrm{sp}(\mathcal{D}) \leqslant k$ then these branch sets have size at most k. □

General decompositions may thus be viewed as a way of relaxing the minor relation, which may prove useful also in contexts otherwise unrelated to our purpose here.

Our next lemma shows that bounded-size decompositions define a transitive relation between graph properties: if $\mathcal{P}_1$ has bounded $\mathcal{P}_2$-size and $\mathcal{P}_2$ has bounded $\mathcal{P}_3$-size then $\mathcal{P}_1$ has bounded $\mathcal{P}_3$-size (and likewise for connected size, width etc.).

Lemma 3.2. *Let G, H, H' be disjoint graphs. If G has an H-decomposition of width at most k and spread at most ℓ, and H has an H'-decomposition of width at most k' and spread at most ℓ', then G has an H'-decomposition of width at most kk' and spread at most $\ell\ell'$. This latter decomposition is connected if the other two decompositions are connected.*

Proof. Let $\mathcal{D} = (G_h)_{h \in H}$ be an H-decomposition of G and $\mathcal{D}' = (H_{h'})_{h' \in H'}$ an H'-decomposition of H, both of the required width and spread. For all

$h' \in H'$ put

$$G_{h'} := G\Big[\bigcup\{\, G_h \mid h \in H_{h'} \,\}\Big].$$

We show that $\mathcal{D}'' := (G_{h'})_{h' \in H'}$ is an H'-decomposition of G which is connected if $\mathcal{D}$ and $\mathcal{D}'$ are connected.

The key observation is that, for every vertex $g \in G$,

$$H'_g = H'\Big[\bigcup\{\, H'_h \mid h \in H_g \,\}\Big].$$

This shows at once that $\mathcal{D}''$ will have the desired spread (and, clearly, $\mathcal{D}''$ also has the desired width). Moreover, (C1) and (C2) for $\mathcal{D}''$ follow at once from the corresponding conditions for $\mathcal{D}$ and $\mathcal{D}'$. Regarding (C3) for $\mathcal{D}''$, note that if both $\mathcal{D}$ and $\mathcal{D}'$ are connected then H_g is connected by (C3) for $\mathcal{D}$, while every H'_h is connected by (C3) for $\mathcal{D}'$. Now as H'_h and $H'_{h'}$ touch whenever hh' is an edge of H_g (by (C2) for $\mathcal{D}'$), this implies that H'_g is connected. □

Corollary 3.3.

(i) *If $G \sqsubseteq_\ell H$ and H has an H'-decomposition of width at most k and spread at most ℓ', then G has an H'-decomposition of width at most k and spread at most $\ell\ell'$.*

(ii) *If $G \preccurlyeq_\ell H$ and H has a connected H'-decomposition of width at most k and spread at most ℓ', then G has a connected H'-decomposition of width at most k and spread at most $\ell\ell'$.*

Proof. Immediate by Lemmas 3.1 and 3.2. □

The message of Corollary 3.3 is that H-decompositions are passed on to minors in a canonical way, without increasing the width, and with an increase in spread by at most the factor ℓ that bounded the branch set size of the minor. Indeed, it is not at all difficult to write down the H'-decomposition $(G_{h'})_{h' \in H'}$ of G explicitly: $G_{h'}$ contains exactly those vertices of G whose corresponding branch sets in H meet $H_{h'}$, where $(H_{h'})_{h' \in H'}$ is the given H'-decomposition of H. In the special case when $G \subseteq H$ this yields $G_{h'} = G \cap H_{h'}$, as expected.

Our next lemma shows that, for instance, the average degree of a property of bounded $\mathcal{H}$-size cannot exceed the average degree of $\mathcal{H}$ (and its subgraphs) by more than a constant.

Lemma 3.4. *Let $\mathcal{H}$ be a graph property closed under taking subgraphs, and let $f\colon \mathbb{N} \to \mathbb{R}$ be a non-decreasing function such that $\|H\| \leqslant f(|H|)$ for all $H \in \mathcal{H}$ and $f(kn) \leqslant k^2 f(n)$ for all $k, n \in \mathbb{N}$. Then for every property $\mathcal{P}$ of bounded $\mathcal{H}$-size there is a constant c such that $\|G\| \leqslant c \cdot f(|G|)$ for all $G \in \mathcal{P}$.*

Proof. Let $k \in \mathbb{N}$ be such that every $G \in \mathcal{P}$ has an H-decomposition of size at most k with $H \in \mathcal{H}$. Deleting empty parts (and using that $\mathcal{H}$ is closed under taking subgraphs) we may assume that $|H| \leqslant k\,|G|$. To estimate the number of edges of G, let us divide them into those edges that lie within some part G_h of this decomposition and those that do not. Every vertex $g \in G$ is incident with at most $k(k-1)$ edges of the first type, because it has at most $k-1$ neighbours in each of the at most k parts it lies in. So G has at most $\binom{k}{2}|G|$ edges of type 1. As every edge hh' of H reflects at most k^2 edges of G (those between G_h and $G_{h'}$), we have at most $k^2||H||$ edges of type 2. Since $|H| \leqslant k\,|G|$, this gives

$$||G|| \leqslant \binom{k}{2}|G| + k^2 f(|H|) \leqslant \binom{k}{2}|G| + k^4 f(|G|)$$

as desired. □

Lemma 3.4 can be used to show that there are infinite sequences $\mathcal{P}_0, \mathcal{P}_1, \ldots$ of graph properties that decrease strictly with respect to bounded preminors. Indeed all we have to do is choose the graphs in $\mathcal{P}_{n+1}$ as (spanning) subgraphs of the graphs in $\mathcal{P}_n$ but with substantially fewer edges, and so that their own subgraphs conform to the same upper bound on their number of edges (in terms of their order) as the graphs in $\mathcal{P}_{n+1}$ themselves. Then Lemma 3.4 (with $\mathcal{P} := \mathcal{P}_n$ and $\mathcal{H}$ the closure of $\mathcal{P}_{n+1}$ under subgraphs) implies $\mathcal{P}_n \sqsupset_* \mathcal{P}_{n+1}$ as desired. More precisely, we have the following result:

Theorem 3.5. *There are infinite sequences $\mathcal{P}_0, \mathcal{P}_1, \ldots$ of graph properties such that, for all $i < j$, we have $\mathcal{P}_i \Supset \mathcal{P}_j$ and $\mathcal{P}_i$ has unbounded $\mathcal{P}_j$-size (and hence $\mathcal{P}_i \sqsupset_* \mathcal{P}_j$ and $\mathcal{P}_i \succ_* \mathcal{P}_j$ by Lemma 3.1).*

The proof of Theorem 3.5 will not be needed for the rest of this paper, but we include it for completeness. It uses the following special case of a result of Bondy and Simonovits [1], which, with an unspecified constant c_ℓ instead of 50ℓ, is due to Erdős [7]:

Lemma 3.6. *For each even integer $\ell \geqslant 4$ every C_ℓ-free graph G has most $50\ell|G|^{1+2/\ell}$ edges.*

Proof of Theorem 3.5. We will construct an infinite sequence $\mathcal{P}_0 \Supset \mathcal{P}_1 \Supset \ldots$ of graph properties such that

(i) for each $i \geqslant 0$ there are positive constants $c_i, \varepsilon_i \leqslant 1$ such that every graph $G \in \mathcal{P}_i$ satisfies $||G|| \geqslant c_i|G|^{1+\varepsilon_i}$ and $\Delta(G) \leqslant |G|^{\varepsilon_i}$;

(ii) for each $i \geqslant 1$ there is an even integer ℓ_i such that all the graphs in $\mathcal{P}_i$ are C_{ℓ_i}-free, $\ell_i > 2/\varepsilon_{i-1}$ and, if $i \geqslant 2$, $\ell_i > \ell_{i-1}$.

Let us first show that for such a sequence each $\mathcal{P}_i$ must have unbounded $\mathcal{P}_j$-size for all $j > i$. So suppose on the contrary that, for some $i < j$, $\mathcal{P}_i$ has bounded $\mathcal{P}_j$-size. Lemma 3.6 implies that $||H|| \leqslant 50\ell_j|H|^{1+2/\ell_j}$ for every subgraph H of a graph in $\mathcal{P}_j$. So by Lemma 3.4 there is a constant c such that $||G|| \leqslant 50c\ell_j|G|^{1+2/\ell_j}$ for all $G \in \mathcal{P}_i$. But if $|G|$ is sufficiently large, then $50c\ell_j|G|^{1+2/\ell_j} < c_i|G|^{1+\varepsilon_i}$ (since $\ell_j > 2/\varepsilon_i$), contradicting our assumption on the graphs in $\mathcal{P}_i$.

To construct such a sequence $\mathcal{P}_0 \supseteq \mathcal{P}_1 \supseteq \dots$ put $\mathcal{P}_0 := \{ K_n \,|\, n \geqslant 2 \}$, $c_0 := 1/4$ and $\varepsilon_0 := 1$ suppose that for some $i \geqslant 0$ we have already defined $\mathcal{P}_0, \dots, \mathcal{P}_i$. Pick an even integer $\ell_{i+1} > \max\{ 2/\varepsilon_i, \ell_i \}$ and set

$$\varepsilon_{i+1} := \frac{\varepsilon_i}{\ell_{i+1}-1}, \qquad c_{i+1} := \frac{c_i^2}{16}.$$

Our aim is to find for every graph $G \in \mathcal{P}_i$ a $C_{\ell_{i+1}}$-free subgraph H with $V(H) = V(G)$, $||H|| \geqslant c_{i+1}|H|^{1+\varepsilon_{i+1}}$ and $\Delta(H) \leqslant |H|^{\varepsilon_{i+1}}$. Then we can take $\mathcal{P}_{i+1}$ to be the set of all these graphs H.

So let $n := |G|$ and consider a random subgraph G_p of G with vertex set $V(G)$ which is obtained by including every edge of G with probability

$$p := c_i n^{\frac{\varepsilon_i}{\ell_{i+1}-1}}/4n^{\varepsilon_i},$$

independently of all other edges of G. Given an edge $e \in G$, call e *bad* if $e \in G_p$ and if G_p contains at least $4p\Delta(G)$ other edges which are adjacent to e. The expected number of edges of G_p which are adjacent to e in G is at most $2p\Delta(G)$. So Markov's inequality implies that

$$\mathbb{P}(\text{at least } 4p\Delta(G) \text{ edges of } G_p \text{ are adjacent to } e \text{ in } G) \leqslant 1/2.$$

Thus with probability at most $p/2$ a given edge $e \in G$ is bad. Therefore,

$$\mathbb{E}(\text{number of bad edges}) \leqslant p||G||/2. \tag{3}$$

Let us now estimate the expected number of $C_{\ell_{i+1}}$'s in G_p. As every $C_{\ell_{i+1}}$ in G is determined by first choosing a vertex of G and then a path of length $\ell_{i+1}-1$ starting at this vertex, G contains at most $n(\Delta(G))^{\ell_{i+1}-1}$ distinct $C_{\ell_{i+1}}$. Since each of them lies with probability $p^{\ell_{i+1}}$ in G_p, we have

$$\begin{aligned} \mathbb{E}(\text{number of } C_{\ell_{i+1}}\text{'s in } G_p) &\leqslant n(\Delta(G))^{\ell_{i+1}-1}p^{\ell_{i+1}} \\ &\leqslant pn^{1+\varepsilon_i} \cdot n^{-\varepsilon_i} n^{\varepsilon_i(\ell_{i+1}-1)} p^{\ell_{i+1}-1} \\ &\underset{(c_i \leqslant 1)}{\leqslant} p||G||/4. \end{aligned} \tag{4}$$

Furthermore,

$$\mathbb{E}(||G_p||) = p||G||. \tag{5}$$

Let X_p denote the difference of the number of all those edges of G_p which are not bad with the number of $C_{\ell_{i+1}}$'s in G_p. Then (3), (4) and (5) together imply that $\mathbb{E}(X_p) \geqslant p||G||/4$. So there exists an outcome G_p with $X_p \geqslant p||G||/4$. Let H be the subgraph of G_p which is obtained by deleting all bad edges as well as one edge on each $C_{\ell_{i+1}}$ in G_p. Then H is $C_{\ell_{i+1}}$-free and

$$||H|| \geqslant \frac{p||G||}{4} \geqslant \frac{pc_i n^{1+\varepsilon_i}}{4} = c_{i+1} n^{1+\varepsilon_{i+1}} = c_{i+1}|H|^{1+\varepsilon_{i+1}}.$$

Moreover, since H does not contain bad edges,

$$\Delta(H) \leqslant 4p\Delta(G) \leqslant 4pn^{\varepsilon_i} \underset{(c_i \leqslant 1)}{\leqslant} |H|^{\varepsilon_{i+1}}.$$

This shows that H is as required and thus completes the construction of $\mathcal{P}_{i+1}$. □

The following observation will be crucial to the hierarchies to be studied below. Let GRID denote the class of all $n \times n$ grids ($n \in \mathbb{N}$).

Proposition 3.7. *Every graph has connected* GRID-*width at most 2.*

Proof. Given any graph G, let H be the $|G| \times |G|$ grid whose rows and columns are each labelled by the vertices of G. For the vertex $h \in H$ in row g and column g' let $G_h := \{ g, g' \}$; then H_g is the 'cross' in H consisting of row g and column g. By (C1)–(C3), $(G_h)_{h \in H}$ is a connected H-decomposition. □

We finish this section by pointing out a connection between classical tree-width and general decompositions. The proof assumes familiarity with some facts and concepts of standard graph minor theory; these are all explained in [2], and so we do not repeat them here.

Theorem 3.8. *A graph property $\mathcal{P}$ has unbounded tree-width if and only if $\mathcal{K}$ has bounded connected $\mathcal{P}$-width.*

Proof. Suppose first that $\mathcal{P}$ has unbounded tree-width. By the Robertson-Seymour grid theorem [9; 5, 2], this implies GRID $\preccurlyeq \mathcal{P}$. By Lemma 3.1 (ii), therefore, the grids have connected $\mathcal{P}$-width at most 1, while by Proposition 3.7 the complete graphs have connected GRID-width at most 2. By Lemma 3.2, the complete graphs thus have connected $\mathcal{P}$-width at most 2.

Conversely, assume that all the graphs in $\mathcal{K}$ have connected $\mathcal{P}$-width at most $k \in \mathbb{N}$. We show that the graphs in $\mathcal{P}$ contain brambles of unbounded order, which implies that $\mathcal{P}$ has unbounded tree-width. (This is the easy direction of the 'tree-width duality theorem'.) Let $r \in \mathbb{N}$ be given; we show that every $G \in \mathcal{P}$ such that $K = K^{kr}$ has a connected G-decomposition $\mathcal{D} = (K_g)_{g \in G}$ of width at most k contains a bramble of order at least r. Indeed, the co-parts $G_h \subseteq G$ ($h \in K$) of this decomposition are connected (C3) and touch pairwise (C2), so they form a bramble $\mathcal{B}$. But each vertex $g \in G$ lies in at most k co-parts, because $\mathcal{D}$ has width at most k. So to cover all the kr bramble sets G_h we need at least r vertices of G, ie. $\mathcal{B}$ has order at least r. □

4. Hierarchies of graph properties: naive and abstract

Now that we have defined $\mathcal{H}$-decompositions and (connected) $\mathcal{H}$-width for graph properties $\mathcal{H}$ other than the trees, let us return to our original plan and see how the hierarchy of graph properties envisaged in the introduction evolves. Let us begin by taking as our 'universe' the class $\mathbb{C}_0$ of all graph properties (up to $\preccurlyeq$-equivalence) ordered by $\preccurlyeq$, and try to slice $\mathbb{C}_0$ into layers of increasing complexity based on connected width.

We have already seen that $\mathbb{C}_0$ has a unique greatest element, the property $\mathcal{K}$ of complete graphs. Similarly, the set

$$\overline{\mathcal{K}} := \{\, \overline{K^n} \mid n \in \mathbb{N} \,\}$$

of edgeless graphs lies below every other property and thus is the least element of $\mathbb{C}_0$; recall that any graph property $\mathcal{P}$ contains unboundedly large graphs (by definition of 'property') and hence satisfies $\overline{\mathcal{K}} \preccurlyeq \mathcal{P}$.

Our bottom layer $\mathbb{D}_0$ of $\mathbb{C}_0$ thus consists of the properties of bounded connected $\overline{\mathcal{K}}$-width. These are readily identified: a property $\mathcal{P}$ has bounded connected $\overline{\mathcal{K}}$-width if and only if its graphs have bounded components, ie. if there exists a $k \in \mathbb{N}$ such that every component of a graph in $\mathcal{P}$ has order at most k. Indeed, every such graph has a connected $\overline{K}$-decomposition (with $\overline{K} \in \overline{\mathcal{K}}$) of width at most k into its components, while conversely the parts of any connected $\overline{K}$-decomposition of a graph must be unions of its components, and so the order of these components will be bounded together with the order of the parts.

Now consider the remaining class $\mathbb{C}_1 := \mathbb{C}_0 \smallsetminus \mathbb{D}_0$, the properties of unbounded connected $\overline{\mathcal{K}}$-width. By Lemma 1.5 every element of $\mathbb{C}_1$ lies above some minimal element of $\mathbb{C}_1$, so $\mathbb{C}_1$ is determined by its minimal elements in Kuratowski-fashion. These minimal elements are the class STAR of all stars and the class PATH of all paths:

Proposition 4.1. STAR $\neq$ PATH, *and*

$$\mathbb{C}_1 = \{\, \mathcal{P} \in \mathbb{C}_0 \mid \text{STAR} \preccurlyeq \mathcal{P} \text{ or } \text{PATH} \preccurlyeq \mathcal{P} \,\}\,.$$

Proof. Clearly STAR and PATH are incomparable under $\preccurlyeq$, and in particular inequivalent. The displayed characterization of $\mathbb{C}_1$ follows from the fact that the graphs in $\mathbb{C}_1$ contain arbitrarily large components, and every large enough connected graph contains either a star or a path of given order. □

The next layer $\mathbb{D}_1$ of our universe thus consists of the properties from $\mathbb{C}_1$ that have bounded connected $\mathcal{H}$-width for $\mathcal{H}$ = STAR ∪ PATH. These, however, are precisely the properties of bounded connected PATH-width: since stars have connected PATH-width at most 2 (put the centre in every part), Lemma 3.2 implies that the connected PATH-width of any property is at most twice its connected (STAR $\cup$ PATH)-width. But the connected PATH-width of a property differs from its traditional path-width by at most a factor of 2 (cf. the footnote in Section 3), so $\mathbb{D}_1$ consists of the properties in $\mathbb{C}_1$ that have bounded path-width.

This takes us back to those early results of Robertson and Seymour mentioned in the introduction. Let TREE denote the class of all trees.

Proposition 4.2. $\mathbb{C}_2 = \{\, \mathcal{P} \in \mathbb{C}_1 \mid \mathrm{TREE} \preccurlyeq \mathcal{P} \,\}$. $\square$

Proposition 4.3. $\mathbb{C}_3 = \{\, \mathcal{P} \in \mathbb{C}_2 \mid \mathrm{GRID} \preccurlyeq \mathcal{P} \,\}$. $\square$

(Here, $\mathbb{C}_2 := \mathbb{C}_1 \setminus \mathbb{D}_1$, $\mathbb{D}_2 := \{\, \mathcal{P} \in \mathbb{C}_2 \mid \mathcal{P}$ has bounded connected TREE-width $\}$ $= \{\, \mathcal{P} \in \mathbb{C}_2 \mid \mathcal{P}$ has bounded tree-width $\}$, and $\mathbb{C}_3 := \mathbb{C}_2 \setminus \mathbb{D}_2$.)

But now comes the disappointment: by Proposition 3.7, our next layer of $\mathbb{C}_0$, the class $\mathbb{D}_3$ of properties in $\mathbb{C}_3$ of bounded connected GRID-width, contains the entire rest of our universe (Fig. 2)—so our hierarchy stops at the very point where it would go beyond those results of Robertson-Seymour!

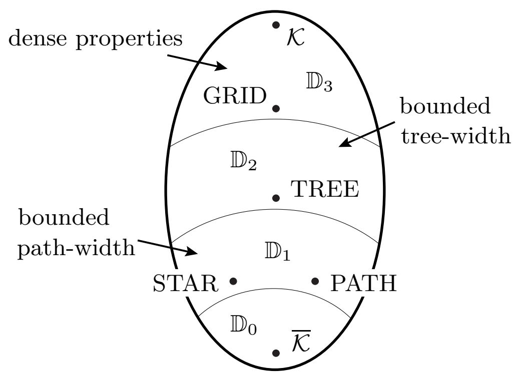

FIGURE 2. The overall hierarchy based on width and unbounded minors

What went wrong? Basically, our use of bounded connected $\mathcal{H}$-width for classes $\mathcal{H}$ other than the trees was naive (and quite counter-intuitive), in that it allowed H-decompositions of G with $|H|$ much bigger than $|G|$. For when H is not a tree then many co-parts H_g in such a decomposition $\mathcal{D}$ can intersect pairwise despite low width (as in the grid example), whereas when H is a tree then these co-parts are subtrees of H and thus have the Helly property, by which they can intersect pairwise only if some $h \in H$ lies in all of them, increasing the order of the corresponding part G_h and thereby the width of $\mathcal{D}$. Thus, the fact that bounding width in a tree-decomposition ensures that the graph decomposed is 'roughly like' the tree along which it was decomposed has more to do with tree structure than one might at first expect.

Rather than try to patch up our definition of $\mathcal{H}$-width, however (eg. by disallowing $|H| > |G|$ in an H-decomposition of G, or similar measures), let us use the rest of this section to lay down in the abstract some minimum requirements for any successful hierarchy, and then review our concrete options more systematically in Section 5.

In any hierarchy of graph properties as above, we shall have two reflexive and transitive relations between properties. The first will be a 'basic' relation $\trianglelefteq$ that expresses which properties we regard as 'smaller' or 'simpler', and which has so far been the unbounded minor relation $\preccurlyeq$. Our objects of study will not be the graph properties themselves but their $\trianglelefteq$-equivalence classes, on which $\trianglelefteq$ defines a partial ordering; any informal use of words like 'above' or 'below' will refer to this ordering $\trianglelefteq$. Since we consider properties only up to $\trianglelefteq$-equivalence, our choice of $\trianglelefteq$ also will also set the 'scale of magnification' for our study.

Second, there will be a relation $\mathcal{P} \leqslant \mathcal{Q}$ to express that the graphs in $\mathcal{P}$ have bounded decompositions in some sense over elements of $\mathcal{Q}$. Since $\leqslant$ is intended to be well-defined also on $\trianglelefteq$-equivalence classes (our objects of study), $\trianglelefteq$-equivalence should imply $\leqslant$-equivalence. Moreover, we shall seek to characterize classes of the form 'all properties of bounded $\mathcal{H}$-size', or $\{\, \mathcal{P} \mid \mathcal{P} \leqslant \mathcal{H} \,\}$, in Kuratowski-fashion by the set of $\trianglelefteq$-minimal elements of their complement, so such classes ought to be closed down under $\trianglerighteq$. Thus, $\mathcal{P} \trianglelefteq \mathcal{Q} \leqslant \mathcal{H}$ should imply $\mathcal{P} \leqslant \mathcal{H}$, which is equivalent (consider $\mathcal{H} = \mathcal{Q}$) to asking that

$$\mathcal{P} \trianglelefteq \mathcal{Q} \quad \Rightarrow \quad \mathcal{P} \leqslant \mathcal{Q}\,,$$

ie. that $\trianglelefteq$ be a refinement (subset) of $\leqslant$. We shall refer to this requirement by saying that $\trianglelefteq$ and $\leqslant$ should be *compatible*. Similarly, we require that $\mathcal{P} \subseteq \mathcal{Q} \Rightarrow \mathcal{P} \leqslant \mathcal{Q}$.

Note that the compatibility of $\trianglelefteq$ and $\leqslant$ is necessary but not sufficient for the existence of a Kuratowski-type characterization of classes of the form $\{\, \mathcal{P} \mid \mathcal{P} \leqslant \mathcal{H} \,\}$: it ensures that the complement of this class is closed upwards under $\trianglelefteq$, but if this complement contains infinite decreasing $\triangleright$-chains then not all of it need lie above its minimal elements. Compatibility does, however, imply that $\leqslant$ is well defined on $\trianglelefteq$-equivalence classes.

Given two classes $\mathbb{C}$ and $\mathbb{C}'$ of graph properties, let us write $\mathbb{C} \trianglelefteq \mathbb{C}'$ if for every $\mathcal{P} \in \mathbb{C}$ there is a $\mathcal{Q} \in \mathbb{C}'$ with $\mathcal{P} \trianglelefteq \mathcal{Q}$. We say that $\mathbb{C}'$ *generates* $\mathbb{C}$ if every element of $\mathbb{C}$ lies above some element of $\mathbb{C}'$. A *basis* of $\mathbb{C}$ is a subclass $\mathbb{B}$ of $\mathbb{C}$ that generates $\mathbb{C}$ and satisfies $\mathbb{B} \trianglelefteq \mathbb{C}'$ for every generating subclass $\mathbb{C}'$ of $\mathbb{C}$.

The following lemma shows that any basis of $\mathbb{C}$ will be unique (so that we may speak of 'the' basis of $\mathbb{C}$), and describes its elements:

Lemma 4.4. *For every class $\mathbb{C}$ of graph properties the following assertions are equivalent.*

(i) *$\mathbb{C}$ has a basis.*

(ii) *$\mathbb{C}$ has a basis that is contained in every generating subclass of $\mathbb{C}$.*

(iii) *Every element of $\mathbb{C}$ lies above some minimal element.*

If (i)–(iii) *hold then the basis of* $\mathbb{C}$ *is unique and consists of the minimal elements of* $\mathbb{C}$.

Proof. Suppose first that (iii) holds. Then the class $\mathbb{B}$ of minimal elements of $\mathbb{C}$ generates $\mathbb{C}$. Since any generating subclass $\mathbb{C}'$ of $\mathbb{C}$ clearly contains $\mathbb{B}$, it follows that $\mathbb{B}$ is a basis as in (ii). Moreover if $\mathbb{C}'$ is also a basis then $\mathbb{C}'$ contains no non-minimal elements of $\mathbb{C}$, because $\mathbb{C}' \trianglelefteq \mathbb{B}$. Hence $\mathbb{C}' = \mathbb{B}$, showing the uniqueness of $\mathbb{B}$ as a basis of $\mathbb{C}$.

The implication (ii)$\to$(i) being trivial, it remains to prove (i)$\to$(iii). Let $\mathbb{B}'$ be any basis of $\mathbb{C}$, and suppose that some $\mathcal{P} \in \mathbb{C}$ does not lie above any minimal element of $\mathbb{C}$. Choose $\mathcal{P}' \in \mathbb{B}'$ with $\mathcal{P}' \trianglelefteq \mathcal{P}$. Then $\mathcal{P}'$ is not minimal in $\mathbb{C}$; pick $\mathcal{Q} \triangleleft \mathcal{P}'$ from $\mathbb{C}$, and choose $\mathcal{Q}' \trianglelefteq \mathcal{Q}$ from $\mathbb{B}'$. Now let $\mathbb{B}''$ be obtained from $\mathbb{B}'$ by deleting $\mathcal{P}'$ and any properties above it; thus $\mathbb{B}'' = \{ \mathcal{P}'' \in \mathbb{B}' \mid \mathcal{P}' \ntrianglelefteq \mathcal{P}'' \}$. Then $\mathbb{B}''$ still generates $\mathbb{C}$ (because it contains $\mathcal{Q}' \triangleleft \mathcal{P}'$) but $\mathbb{B}' \ntrianglelefteq \mathbb{B}''$, contradicting the fact that $\mathbb{B}'$ is a basis. $\square$

Now let $\mathbb{C}_0$ be any fixed class of graph properties that is closed under union, ie. such that for every set $\mathbb{C}_0' \subseteq \mathbb{C}_0$ the property $\bigcup \mathbb{C}_0'$ is an element of $\mathbb{C}_0$. We shall think of $\mathbb{C}_0$ as the universe of properties we wish to study—perhaps the class of all graph properties, or just the class of all tree properties (graph properties consisting of trees only), or the class of graph properties of unbounded tree-width etc.—and will look at $\mathbb{C}_0$ in terms of the posets $\mathbb{P}_{\trianglelefteq} = \mathbb{P}_{\trianglelefteq}(\mathbb{C}_0)$ and $\mathbb{P}_{\leqslant} = \mathbb{P}_{\leqslant}(\mathbb{C}_0)$ that our relations $\trianglelefteq$ and $\leqslant$ impose on it.

To this end, let us define recursively for all ordinals α:

$$\begin{aligned}
\mathbb{C}_\alpha &:= \bigcap_{\beta<\alpha} \mathbb{C}_\beta \text{ when } \alpha \text{ is a non-zero limit;}\\
\mathbb{B}_\alpha &\text{ as the basis of } \mathbb{C}_\alpha \text{ if it exists;}\\
\mathcal{B}_\alpha &:= \bigcup \mathbb{B}_\alpha;\\
\mathbb{D}_\alpha &:= \{ \mathcal{P} \in \mathbb{C}_\alpha \mid \mathcal{P} \leqslant \mathcal{B}_\alpha \};\\
\mathbb{C}_{\alpha+1} &:= \mathbb{C}_\alpha \smallsetminus \mathbb{D}_\alpha.
\end{aligned}$$

Thus, we have a strictly descending well-ordered sequence $\mathbb{C}_0 \supset \mathbb{C}_1 \supset \dots$ of classes of graph properties with $\mathcal{B}_\alpha \in \mathbb{C}_\alpha \smallsetminus \mathbb{C}_{\alpha+1}$ (note that $\mathbb{C}_\alpha$ is closed under union; induction on α), which either terminates naturally with $\mathbb{C}_\alpha = \emptyset$ for some α, or comes to an emergency halt when some $\mathbb{C}_\alpha$ has no basis. If it terminates naturally with the empty class, then every graph property $\mathcal{P} \in \mathbb{C}_0$ lies in some $\mathbb{D}_\alpha$: let β be minimal with $\mathcal{P} \notin \mathbb{C}_\beta$, note that β cannot be a limit, and let α be such that $\beta = \alpha + 1$.

Induction on α shows that the classes $\mathbb{C}_\alpha$ are upwards-closed in $\mathbb{C}_0$ under $\leqslant$ (and hence also under $\trianglelefteq$): if $\mathcal{P} \in \mathbb{C}_\alpha$ and $\mathcal{P} \leqslant \mathcal{Q} \in \mathbb{C}_0$ then $\mathcal{Q} \in \mathbb{C}_\alpha$. Hence the 'layers' $\mathbb{D}_\alpha$ of our hierarchy are convex in $\mathbb{P}_{\leqslant}$: if $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3 \in \mathbb{C}_0$ satisfy $\mathcal{P}_1 \leqslant \mathcal{P}_2 \leqslant \mathcal{P}_3$ and $\mathcal{P}_1, \mathcal{P}_3 \in \mathbb{D}_\alpha$ then $\mathcal{P}_2 \in \mathbb{D}_\alpha$. In particular, both $\mathbb{C}_\alpha$ and $\mathbb{D}_\alpha$

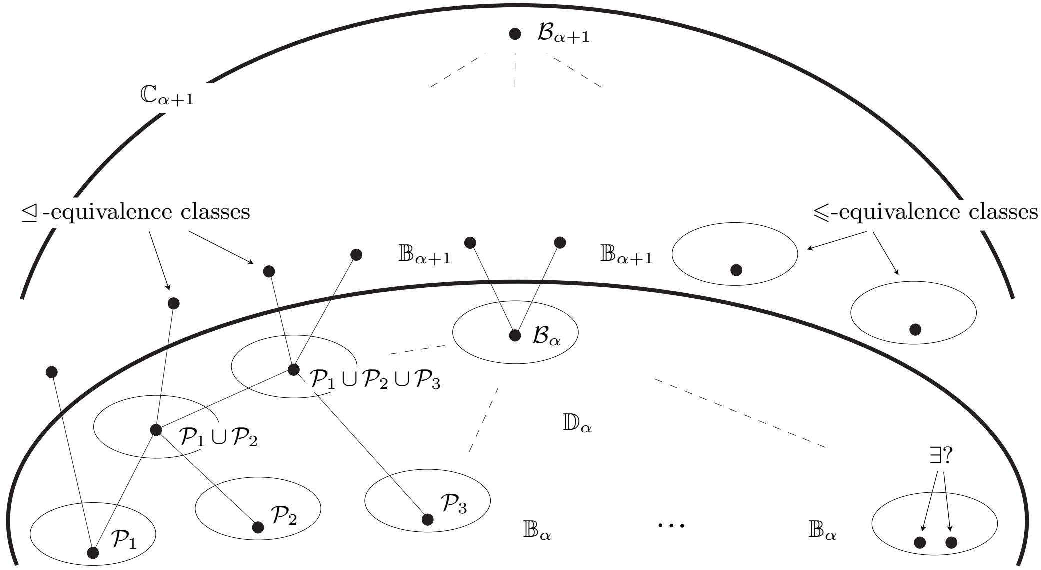

FIGURE 3. Typical posets $\mathbb{P}_\trianglelefteq$ and $\mathbb{P}_\leqslant$

are closed in $\mathbb{C}_0$ under $\leqslant$-equivalence, ie. are unions of $\leqslant$-equivalence classes in $\mathbb{C}_0$ (Fig. 3).

By Lemma 4.4, the properties in $\mathbb{B}_\alpha$ are the minimal elements of $\mathbb{C}_\alpha$. Their union $\mathcal{B}_\alpha$, however, lies higher in $\mathbb{P}_\trianglelefteq$ (unless $|\mathbb{B}_\alpha| = 1$) as soon as $\mathcal{P} \subseteq \mathcal{Q}$ implies $\mathcal{P} \trianglelefteq \mathcal{Q}$ (which it certainly will for all the minor-type relations we shall consider for $\trianglelefteq$): since the elements of $\mathbb{B}_\alpha$ are incomparable, the union of any two of them lies strictly above both in $\mathbb{P}_\trianglelefteq$. If the properties in $\mathbb{B}_\alpha$ are indivisible (cf. Section 2 and Proposition 4.5 below) then more generally by Lemma 2.1 the union of any k of them lies at least $k-1$ levels higher than each. If $\mathbb{B}_\alpha$ is finite (and $\mathcal{P}_1, \mathcal{P}_2 \trianglelefteq \mathcal{Q}$ imply $\mathcal{P}_1 \cup \mathcal{P}_2 \trianglelefteq \mathcal{Q}$, which again will always be the case), then $\mathcal{B}_\alpha$ is the least upper bound for $\mathbb{B}_\alpha$ (and hence for $\mathbb{D}_\alpha$): any property $\mathcal{Q}$ such that $\mathcal{P} \trianglelefteq \mathcal{Q}$ for all $\mathcal{P} \in \mathbb{B}_\alpha$ clearly satisfies $\mathcal{B}_\alpha \trianglelefteq \mathcal{Q}$.

For every fixed universe $\mathbb{C}_0$ and every choice of $\trianglelefteq$ and $\leqslant$, a number of obvious questions arise that cannot be answered in general. The most prominent of these, of course, is whether every $\mathbb{C}_\alpha$ has a basis. If $\trianglelefteq$ is well-founded in $\mathbb{C}_0$ then this is so by Lemma 4.4, but it may well hold in other cases too.

Another question is whether the $\mathcal{B}_\alpha$ form an increasing $\leqslant$-chain, ie. satisfy $\mathcal{B}_\beta \leqslant \mathcal{B}_\alpha$ (and hence $\mathcal{B}_\beta < \mathcal{B}_\alpha$) for all $\beta < \alpha$. While it is possible to construct artificial classes $\mathbb{C}_0$ where this fails—Figure 4 shows a portion of $\mathbb{P}_{\preccurlyeq_*}$ when $\mathbb{C}_0$ consists of all properties of bounded diameter (with $\mathcal{Q}_n$ denoting the trees of height n), the additional property $\mathcal{P} = $ PATH, and all unions of these properties, and bases are taken with respect to $\leqslant_*$; see Section 5 for the definition of $\preccurlyeq_*$ and $\leqslant_*$—it seems to be a feature one would expect for most natural classes $\mathbb{C}_0$. It would certainly have a number of natural consequences:

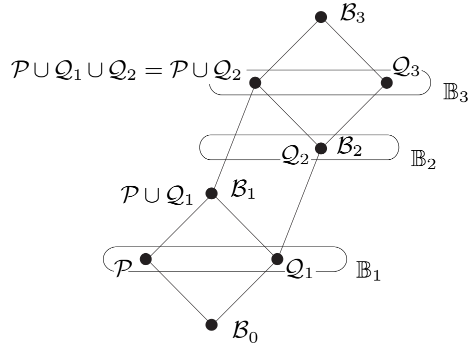

FIGURE 4. A pathological example for $\mathbb{P}_{\preccurlyeq_*}$, where $\mathcal{P}$ has bounded $\mathcal{B}_1$-size but unbounded $\mathcal{B}_2$-size. Here $\mathcal{B}_1 \not\preccurlyeq_* \mathcal{B}_2$, and $\mathcal{P} \cup \mathcal{Q}_2 \in \mathbb{B}_3$ is divisible.

Proposition 4.5. *Let $\mathbb{P}_{\leqslant}$ be such that $\mathcal{P}_2 \smallsetminus \mathcal{P}_1$ is never finite when $\mathcal{P}_1 < \mathcal{P}_2$, and $\mathcal{P}_1, \mathcal{P}_2 \leqslant \mathcal{Q}$ imply $\mathcal{P}_1 \cup \mathcal{P}_2 \leqslant \mathcal{Q}$; for all $\mathcal{P}_1, \mathcal{P}_2, \mathcal{Q} \in \mathbb{C}_0$. If $\mathcal{B}_\beta \leqslant \mathcal{B}_\alpha$ for all $\beta < \alpha$, the following assertions hold for every α:*

(i) $\mathbb{C}_{\alpha+1} = \{ \mathcal{P} \in \mathbb{C}_0 \mid \mathcal{P} \not\leqslant \mathcal{B}_\alpha \}$.

(ii) *The properties $\mathcal{P} \in \mathbb{B}_\alpha$ are indivisible in $\mathbb{C}_0$, ie. are not the union of two properties $\mathcal{P}_i \in \mathbb{C}_0$ with $\mathcal{P}_i \lhd \mathcal{P}$ $(i = 1, 2)$ or the union of one such property and a finite set.*

(The point about (i) above is that while $\mathbb{C}_{\alpha+1} = \{ \mathcal{P} \in \mathbb{C}_\alpha \mid \mathcal{P} \not\leqslant \mathcal{B}_\alpha \}$ holds by definition of $\mathbb{C}_{\alpha+1}$, its characterization as in (i) is no longer subject to the recursive definition of $\mathbb{C}_\alpha$. Thus, once a basis $\mathcal{B}_\alpha$ has been determined, the class $\mathbb{C}_{\alpha+1}$ and its complement $\bigcup_{\beta \leqslant \alpha} \mathbb{D}_\beta$ in $\mathbb{C}_0$ can be written down explicitly.)

Proof. (i) We apply induction on α to show that $\{ \mathcal{P} \in \mathbb{C}_0 \mid \mathcal{P} \not\leqslant \mathcal{B}_\alpha \} \subseteq \{ \mathcal{P} \in \mathbb{C}_\alpha \mid \mathcal{P} \not\leqslant \mathcal{B}_\alpha \} = \mathbb{C}_{\alpha+1}$. This is clear for $\alpha = 0$, so let $\alpha > 0$ and $\mathcal{P} \in \mathbb{C}_0$ with $\mathcal{P} \not\leqslant \mathcal{B}_\alpha$ be given. For every $\beta < \alpha$ we have $\mathcal{P} \not\leqslant \mathcal{B}_\beta$ since $\mathcal{B}_\beta \leqslant \mathcal{B}_\alpha$, so $\mathcal{P} \in \mathbb{C}_{\beta+1}$ by the induction hypothesis. Hence $\mathcal{P} \in \mathbb{C}_\alpha$ both when $\alpha = \beta + 1$ and when α is a limit.

(ii) By Lemma 4.4, the properties $\mathcal{P} \in \mathbb{B}_\alpha$ are the minimal elements of $\mathbb{C}_\alpha$. Therefore any $\mathcal{P}_i$ as above lies in $\mathbb{C}_0 \smallsetminus \mathbb{C}_\alpha$, and hence lies in $\mathbb{D}_{\beta_i}$ for some $\beta_i < \alpha$: let $\gamma \leqslant \alpha$ be minimal with $\mathcal{P}_i \notin \mathbb{C}_\gamma$, note that γ cannot be a limit, and let $\beta_i < \alpha$ be such that $\gamma = \beta_i + 1$. Then if $\beta_1 \leqslant \beta_2$ (say) we have both $\mathcal{P}_1 \leqslant \mathcal{B}_{\beta_1} \leqslant \mathcal{B}_{\beta_2}$ and $\mathcal{P}_2 \leqslant \mathcal{B}_{\beta_2}$. By our assumptions on $\leqslant$ this (and similarly $|\mathcal{P}_1| < \infty$) gives $\mathcal{P}_1 \cup \mathcal{P}_2 \leqslant \mathcal{B}_{\beta_2}$, so $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2$ would imply that $\mathcal{P} \notin \mathbb{C}_{\beta_2+1} \supseteq \mathbb{C}_\alpha$, a contradiction. $\square$

A third general question is how much we can say about the internal $\trianglelefteq$-structure of a single $\leqslant$-equivalence class $\mathbb{C}$. For example, while it is not difficult for most choices of $\trianglelefteq$ and $\leqslant$ to find $\leqslant$-equivalent properties that are $\trianglelefteq$-incomparable, can they both be $\trianglelefteq$-minimal in $\mathbb{C}$? In particular, can $\mathbb{B}_\alpha$

contain more than one element from each $\leqslant$-equivalence class in $\mathbb{C}_\alpha$? Is $\mathbb{C}$ necessarily well-founded by $\trianglelefteq$?

5. Hierarchies of graph properties: some concrete models

Let us set out from our naive hierarchy based on $\preccurlyeq$ and connected width, as described at the start of Section 4. This hierarchy failed to produce new results, because its relation $\leqslant$ was too coarse: we had $\mathcal{K} \leqslant$ GRID in terms of connected width, and so the universe of all graph properties had no more than the four layers shown in Figure 2. How, then, should we sharpen $\leqslant$ to ensure that $\mathcal{K} \not\leqslant$ GRID?

5.1. Bounding spread

Among the many possibilities to achieve this, the most promising appears to be to bound not only the width but also the spread of the decompositions used: when the co-parts H_g of an H-decomposition of G are too large and H is not tree-like, they can touch in many ways without by their connectedness forcing any of the parts G_h to become large (which would happen if H were a tree); hence even a complete graph can have an H-decomposition of small connected width. So let us use $G \leqslant_k H$ to express that G has a connected H-decomposition of *size* at most k, put $\mathcal{P} \leqslant_k \mathcal{Q}$ if for every $G \in \mathcal{P}$ there is an $H \in \mathcal{Q}$ with $G \leqslant_k H$, and write $\mathcal{P} \leqslant_* \mathcal{Q}$ when $\mathcal{P} \leqslant_k \mathcal{Q}$ for some k, ie. when $\mathcal{P}$ has bounded connected $\mathcal{Q}$-size. Note that $\leqslant_*$ is reflexive (because every graph G has the trivial G-decomposition into singletons), and it is transitive by Lemma 3.2.

The following lemma gives a taste of this relation.

Lemma 5.1. *Let $\mathcal{P} \leqslant_* \mathcal{Q}$ be two graph properties.*

(i) *If the vertices of the graphs in $\mathcal{Q}$ have bounded degree then so do those of the graphs in $\mathcal{P}$.*

(ii) *If the paths in the graphs in $\mathcal{Q}$ have bounded length then so do those in the graphs in $\mathcal{P}$.*

Proof. Let $G \in \mathcal{P}$ and $H \in \mathcal{Q}$, and let $(G_h)_{h \in H}$ be a connected H-decomposition of G of size at most k.*

(i) We assume that $\Delta(H) \leqslant d$ and show that $\Delta(G)$ is bounded in terms of d and k. Consider a vertex $g \in G$. Every $H_{g'}$ that touches H_g contains either one of the at most k vertices of H_g or one of its at most dk neighbours. But each of those at most $k(d+1)$ vertices of H lies in $H_{g'}$ for at most k different g', so by (C2) g has no more than $k^2(d+1)$ neighbours.

(ii) We assume that H contains no path of length $\ell \in \mathbb{N}$ and show that the length of the paths in G is bounded in terms of k and ℓ. Let $P \subseteq G$ be any

* In the proof of (i) we shall not use that the decomposition is connected.

path, and let H' be a minimal subgraph of H (not necessarily induced) such that $(P_h)_{h \in H'}$ is a connected H'-decomposition of P, where $P_h := G_h \cap P$. Then H' is connected (because P is; cf. (C2) and (C3)) and has order at least $|P|/k$, so as H' contains no path of length ℓ it has a vertex h of degree at least $k^2 + k + 1$ if $|P|$ is large enough in terms of k and ℓ. We show that h cannot exist, and hence that $|P|$ is bounded as desired.

By the minimality of H', every neighbour h' of h in H' is of one of two types: either the edge hh' is needed for (C3) and thus lies in $H'_g = \{ h'' \in H' \mid g \in P_{h''} \}$ for some $g \in P$, or it is needed for (C2) and joins H'_g to $H'_{g'}$ for some edge $gg' \in P$. If h' is of the first type, then g is one of the at most k vertices of P_h (since $h \in H'_g$), and h' is one of the at most $k-1$ neighbours of h in H'_g. So at most $k(k-1)$ of the neighbours of h are of the first type. Similarly, if h' is of the second type (but not the first), then g is one of the at most k vertices in P_h and g' is one of the at most two neighbours of g on P; thus, h has at most $2k$ neighbours of the second type. (By the minimality of H', every edge gg' of P gives rise to at most one edge hh' by (C2).) Hence h has at most $k(k-1) + 2k$ neighbours in H', contradicting the choice of h. □

However, $\leqslant_*$ is no longer compatible with the unbounded minor relation $\preccurlyeq$. Indeed, if $\mathcal{F}$ denotes the property of being a *fan* (a path plus a new vertex joined to every vertex of the path), and $\mathcal{L}$ is the property of being a *ladder* (a $2 \times n$ grid), then $\mathcal{F} \preccurlyeq \mathcal{L}$ but $\mathcal{F} \nleqslant_* \mathcal{L}$ by Lemma 5.1 (i).

So let us instead use bounded minors as our basic relation $\trianglelefteq$: from Corollary 3.3 (ii) we know that $\mathcal{P} \preccurlyeq_* \mathcal{Q}$ implies $\mathcal{P} \leqslant_* \mathcal{Q}$, so $\preccurlyeq_*$ and $\leqslant_*$ are indeed compatible. We remark in passing that, conversely, even $\mathcal{P} <_* \mathcal{Q}$ does not imply $\mathcal{P} \preccurlyeq_* \mathcal{Q}$: in our fan/ladder example we clearly have $\mathcal{L} \not\preccurlyeq_* \mathcal{F}$, while $\mathcal{L} \leqslant_* \mathcal{F}$ (just use the path in the fan to decompose along) but $\mathcal{F} \nleqslant_* \mathcal{L}$, and hence $\mathcal{L} <_* \mathcal{F}$.

The posets $\mathbb{P}_{\preccurlyeq_*}$ and $\mathbb{P}_{\leqslant_*}$ may be well worth studying for suitable 'small' universes $\mathbb{C}_0$: even for tree properties one can obtain non-trivial results [6]. As regards the overall universe of all graph properties, however, this would be a hopeless task: just as the poset $\mathbb{P}_{\leqslant}$ based on connected width was too coarse to be interesting, $\mathbb{P}_{\leqslant_*}$ is too fine. Even for trees it distinguishes properties (such as trees of different maximum height, or with different degrees of subdivision) that are hardly worth the effort of telling them apart. And besides, there are infinite strictly decreasing chains in both $\mathbb{P}_{\preccurlyeq_*}$ and $\mathbb{P}_{\leqslant_*}$ (Theorem 3.5) that make the hierarchy complicated without adding much insight. (Figure 5 shows some of the results from [6] on the hierarchy of tree properties based on $\preccurlyeq_*$ and $\leqslant_*$. A *comb* is a graph obtained by joining a set of isolated vertices to a given path by disjoint paths; an n-*bush* is a tree of height at most n, the graphs in PATH & STAR are disjoint unions of one path and one star, and n-BUSHES consists of disjoint unions of n-bushes.)

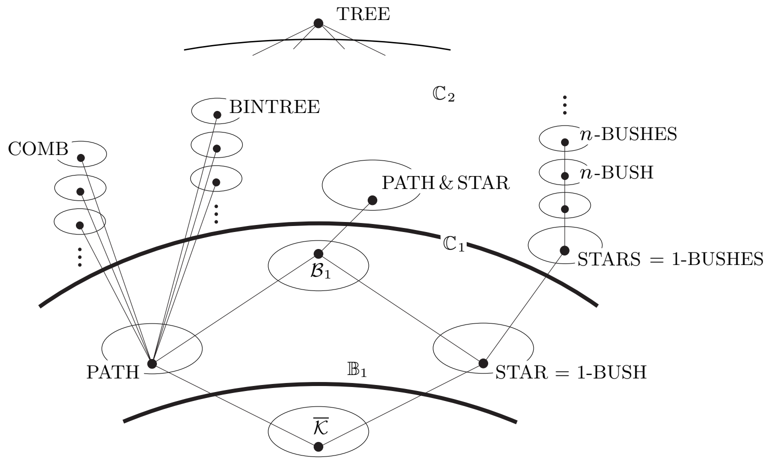

FIGURE 5. A portion of the $(\preccurlyeq_*, \leqslant_*)$-hierarchy of tree properties

5.2. Allowing subdivisions

How could we modify $\preccurlyeq_*$ and $\leqslant_*$ so as to get rid of at least some of the infinite decreasing chains in $\mathbb{P}_{\preccurlyeq_*}$ and $\mathbb{P}_{\leqslant_*}$? Since the most obviously irritating such chains arise just by subdividing—recall the example after Lemma 1.5—the first measure to take would be to include topological minors as special cases of $\trianglelefteq$ and, to preserve compatibility, also of $\leqslant$. One might even go a step further and also allow the converse operation, the suppression of vertices of degree 2. Allowing this as part of $\trianglelefteq$ (and for compatibility then also of $\leqslant$) would indicate that we wish to regard homeomorphic graphs as essentially the same. Allowing the suppression of degree-2 vertices only for $\leqslant$ would indicate that we consider a subdivision of a graph as 'larger but not much'; recall that the intuition behind $\mathcal{P} \leqslant \mathcal{Q}$ was that the graphs in $\mathcal{P}$ should be 'not much larger' than those in $\mathcal{Q}$.

Note that, in order to ensure transitivity, we may have to allow mixed sequences of the various relations we want to admit. For example, we might define $\mathcal{P} \trianglelefteq \mathcal{Q}$ as "*there exist properties* $\mathcal{P}_1, \dots, \mathcal{P}_k$ *with* $\mathcal{P}_1 = \mathcal{P}$ *and* $\mathcal{P}_k = \mathcal{Q}$ *such that, for every* $i = 1, \dots, k-1$, *either* $\mathcal{P}_i \preccurlyeq_* \mathcal{P}_{i+1}$ *or* $\mathcal{P}_i \preccurlyeq_{\text{top}} \mathcal{P}_{i+1}$" (and likewise allow sequences of $\leqslant_*$ and $\preccurlyeq_{\text{top}}$ for $\leqslant$, where $\preccurlyeq_{\text{top}}$ stands for topological minors).

Allowing topological minors, or more generally homeomorphic equivalence, as part of $\leqslant$ will not result in $\mathcal{K} \leqslant$ GRID, since Lemma 5.1 (i) continues to apply. (By contrast, allowing sequences of $\preccurlyeq$ and $\leqslant_*$ in the definition of $\leqslant$ does result in $\mathcal{K} \leqslant$ GRID, so that is not an option.) However, we do get $\mathcal{K} \leqslant$ PLANAR, where PLANAR denotes the class of planar graphs: just draw an arbitrary complete graph K with crossings (of two edges at a time), turn it into a TK by inserting a pair of subdividing vertices at each crossing on the two edges involved, and observe that $K \preccurlyeq_{\text{top}} TK \leqslant_2 H$ for the plane graph

H that arises from the drawing of K by identifying the pairs of subdividing vertices.

Even with these relaxations, hierarchies for large universes may still show more diversity than one will be able to survey. For example, one can construct quite a rich universe just from properties combining trees of bounded height (but unbounded degree) with trees of bounded degree (but unbounded height) to larger graphs in various ways. It therefore seems desirable to look for further ways of weakening the relation $\leqslant$ (without making the complete graphs equivalent to the grids).

5.3. Bounding essential spread only

Rather than bounding the order of every H_g in a connected H-decomposition of G, one might choose to bound only the number of those vertices in each H_g that are needed to accommodate edges of G (rather than just to make H_g connected). Such a decomposition would then be specified as a pair $(\mathcal{D}, \overline{\mathcal{D}})$ of H-decompositions of G, where $\overline{\mathcal{D}}$ is connected but $\mathcal{D}$ need not be connected, and where $\mathcal{D}$ and $\overline{\mathcal{D}}$ have families $(H_g)_{g \in G}$ and $(\overline{H}_g)_{g \in G}$ of co-parts such that $H_g \subseteq \overline{H}_g$ for all g. Let us call $\max(\mathrm{wd}(\overline{\mathcal{D}}), \mathrm{sp}(\mathcal{D}))$ the *essential size* of this decomposition.

As it turns out, decompositions of bounded essential size form a special case of the first relation $\leqslant$ considered in (5.2) above (allowing subdivisions only): if $(\mathcal{D}, \overline{\mathcal{D}})$ is an H-decomposition of G of essential size k as above, we can find graphs G' and G'' such that $G \leqslant_{2k} G' \preccurlyeq_{\mathrm{top}} G'' \leqslant_k H$. Indeed, for each $g \in G$ let H'_g be a minimal connected subgraph of $\overline{H}_g$ containing $V(H_g)$, and let H''_g be obtained from H'_g by suppressing any vertices of degree 2. Since H'_g is a tree with at most k leaves, H''_g has no more than $2k$ vertices. Now let G' be obtained from the disjoint union $\bigcup_{g \in G} H''_g$ by adding for every edge $gg' \in G$ an edge between a vertex $h \in H''_g$ and a vertex $h' \in H''_{g'}$ such that $h = h'$ (in H) or $hh' \in E(H)$; such h and h' exist by (C2) for $\mathcal{D}$. Contraction of the trees H''_g then yields $G \preccurlyeq_{2k} G'$, and hence $G \leqslant_{2k} G'$ as desired. Now let G'' be obtained from G' by reinserting the suppressed vertices of degree 2 (turning each $H''_g \subseteq G'$ back into a copy of H'_g). Then the natural map from $V(G'')$ to $V(H)$ defines an H-decomposition of G'' of spread 1 and width $\mathrm{wd}(\mathcal{D}) = k$.

But conversely, the relation $\leqslant$ based on bounding essential size excludes some of the undesirable instances of the relation of 'bounded size plus topological minors'. For example, we saw that the complete graphs are topological minors of graphs with bounded-size (connected) decompositions over planar graphs. But as Theorem 5.3 (iii) will show, their essential size over planar graphs is unbounded.

In fact, those subdivisions that can be realized by decompositions of bounded essential size can be determined precisely:

Lemma 5.2. *Let $\mathcal{P}$ be any graph property, and let $\mathcal{Q}$ be obtained from the graphs in $\mathcal{P}$ by subdividing every edge at least once. Then $\mathcal{P}$ has bounded* *essential $\mathcal{Q}$-size if and only if $\mathcal{P}' = \{ G' \subseteq G \mid G \in \mathcal{P} \}$ is sparse.*

Proof. Note first that every subgraph of a graph in $\mathcal{Q}$ has average degree at most 4, because its vertices of degree at most 2 cover all its edges. Now if $\mathcal{P}$ has essential $\mathcal{Q}$-size at most k, then every $G' \subseteq G \in \mathcal{P}$ has an H-decomposition (not necessarily connected) with $H \in \mathcal{Q}$ of size at most k. By Lemma 3.4, therefore, the average degree of G' is bounded in terms of k.

Conversely, assume that $\mathcal{P}' = \{ G' \subseteq G \mid G \in \mathcal{P} \}$ is sparse and orient the edges of the graphs $G \in \mathcal{P}$ as in Proposition 1.1 (iv). Thus, all their vertices have out-degrees $< k$, say. For each $G \in \mathcal{P}$, pick a subdivision $H \in \mathcal{Q}$; we shall define a pair $(\mathcal{D}, \overline{\mathcal{D}})$ of H-decompositions of G (where $\overline{\mathcal{D}}$ is connected) in terms of their co-parts H_g and $\overline{H}_g \supseteq H_g$ so that $\mathrm{wd}(\overline{\mathcal{D}}) = 1$ and $\mathrm{sp}(\mathcal{D}) \leqslant k$. As H_g we take the bounded (but possibly disconnected) set consisting of g and the farthest subdividing vertex on each of the $< k$ edges at g that are oriented away from g. As $\overline{H}_g$ we take the unique minimal connected subgraph of the union of these subdivided edges that contains H_g (a subdivided star with centre g). Then the H_g satisfy (C1) and (C2), the $\overline{H}_g$ satisfy (C3), and the width of $\overline{\mathcal{D}}$ and spread of $\mathcal{D}$ are as desired. □

5.4. Disconnected decompositions

The decompositions of (5.3) raise the question of what happens if we discard $\overline{\mathcal{D}}$ altogether and bound the width of $\mathcal{D}$ instead, ie. simply consider decompositions that are not necessarily connected.

Here is a tempting reason for doing so. In a standard tree-decomposition $(G_t)_{t \in T}$, the requirement (T3) that the co-parts T_g be connected has the effect of curbing the potential of large spread for accommodating too many edges of G implicitly through width. (Recall that, by the Helly property of the subtrees of a tree, many co-parts T_g can touch pairwise only if some $t \in T$ lies in all of them, in which case $g \in G_t$ for all those g.) Might this be the *only* reason for requiring (T3)? In other words, now that we have to bound spread explicitly anyhow, could we do without the axiom (C3)?

To look at the corresponding hierarchies in more detail, let us define $\mathcal{P} \leqslant_*^- \mathcal{Q}$ to mean that the graphs in $\mathcal{P}$ have (not necessarily connected) H-decompositions with $H \in \mathcal{Q}$ of bounded size. We may then also relax the bounded minor relation between properties to bounded preminors without losing compatibility, ie. take $\sqsubseteq_*$ instead of $\preccurlyeq_*$ as our basic relation $\trianglelefteq$. The hierarchy of all graph properties then begins as follows.

Theorem 5.3. *The first three layers $\mathbb{D}_0$, $\mathbb{D}_1$, $\mathbb{D}_2$ of the $(\sqsubseteq_*, \leqslant_*^-)$-hierarchy of the class $\mathbb{C}_0$ of all graph properties are characterized as follows.*

(i) $\mathbb{B}_0 = \{ \overline{\mathcal{K}} \}$. *The class $\mathbb{D}_0$ consists of the properties of bounded maximum degree.*

(ii) $\mathbb{B}_1 = \{ \mathrm{STAR} \}$. *The class $\mathbb{D}_1$ consists of those properties in $\mathbb{C}_1$ in whose graphs only boundedly many vertices have unbounded degree. (Formally:*

$\mathcal{P} \in \mathbb{C}_1$ lies in $\mathbb{D}_1$ if and only if $(\exists k)(\forall G \in \mathcal{P}) \,|\{ v \in G : d(v) \geqslant k \}| \leqslant k$.)

(iii) $\mathbb{B}_2 = \{ \text{STARS} \}$, *where* STARS *is the property of being a disjoint union of stars. The class $\mathbb{D}_2$ consists of those properties $\mathcal{P} \in \mathbb{C}_2$ for which $P' = \{ G' \subseteq G \mid G \in \mathcal{P} \}$ is sparse.*

(The properties occuring in (iii) were characterized in Proposition 1.1.)

Proof. (i) Clearly, $\overline{\mathcal{K}}$ is the least element in $\mathbb{C}_0$ under $\sqsubseteq_*$. For the second statement, consider a property $\mathcal{P} \in \mathbb{D}_0$, ie. $\mathcal{P} \leqslant_*^- \overline{\mathcal{K}}$. Since $\overline{\mathcal{K}}$ has bounded degrees and Lemma 5.1 holds also for disconnected decompositions, $\mathcal{P}$ has bounded degrees too. Conversely, suppose that $\Delta(G) < k$ for every $G \in \mathcal{P}$. Let $K \in \overline{\mathcal{K}}$ be the edgeless graph on $V(G)$, and for every $g \in G$ let K_g consist of g and its neighbours. These K_g satisfy (C1) and (C2), so $(K_g)_{g \in G}$ is the family of co-parts of a K-decomposition $\mathcal{D}$ of G. Then $\mathrm{sp}(\mathcal{D}) \leqslant k$ since $|K_g| \leqslant k$ for all g, and $\mathrm{wd}(\mathcal{D}) \leqslant k$ because each vertex h of K lies in at most k co-parts K_g, those with $g = h$ or $g \in N(h)$. Thus $\mathcal{P} \leqslant_k^- \overline{\mathcal{K}}$, ie. $\mathcal{P} \in \mathbb{D}_0$ as required.

(ii) By (i), $\mathbb{C}_1$ consists of the properties of unbounded maximum degree. Hence STAR $\Subset \mathcal{P}$ for every $\mathcal{P} \in \mathbb{C}_1$, so STAR is the least element of $\mathbb{C}_1$ also under $\sqsubseteq_*$. Now consider $\mathcal{P} \in \mathbb{D}_1$. Then all the graphs in $\mathcal{P}$ have bounded-size STAR-decompositions. Deleting their central parts leaves graphs that have bounded-size K-decompositions with $K \in \overline{\mathcal{K}}$, and hence have bounded degree by (i). Conversely, if G has at most k vertices of degree $\geqslant k$, then $G \leqslant_k^- S$ for every sufficiently large star S: put all the vertices of degree $\geqslant k$ in the central part, and add a K-decomposition of size at most k of the rest of G as in (i), where K is the set of leaves of S.

(iii) The proof of $\{ \mathbb{B}_2 \} =$ STARS is again clear from (ii). For the second assertion, it suffices by Proposition 1.1 to prove that a property $\mathcal{P} \in \mathbb{C}_2$ lies in $\mathbb{D}_2$ if and only if its graphs have orientations of bounded out-degree.

For the forward implication, let G be a graph with a U-decomposition $\mathcal{D}$ of size k, where U is a disjoint union of stars. Consider an edge $e = gg' \in G$. If g and g' lie in some common part of $\mathcal{D}$, orient e arbitrarily. If not, then g (say) lies in a part G_u such that u is a leaf of one of the stars in U, and we orient e from g towards g'. Now consider a fixed vertex g, and let us count the edges gg' oriented away from g. For every such edge, either g' lies in the same part as g, or g lies in a leaf part and g' lies in the unique adjacent centre part. Since g lies in at most k parts and for each of these there are at most $2k$ choices for g', the out-degree of g is at most $2k^2$.

Conversely, let G be a graph with an orientation with all out-degrees less than k. For each vertex $g \in G$ take a star $S(g)$ of order $|G|$ with centre $s(g)$, and let U be the disjoint union of those stars. To define a U-decomposition of G, let the co-part U_g for $g \in G$ consist of $s(g)$ and one leaf from every star $S(g')$ such that g' is an out-neighbour of g; let these leaves of $S(g')$ be chosen distinct for different g. These U_g satisfy (C1) and (C2) as required. Clearly, this decomposition has width 1 and spread at most k. □

Theorem 5.3 describes the complete $(\sqsubseteq_*, \leqslant_*^-)$-hierarchy of those graph properties $\mathcal{P}$ that are 'essentially' sparse, in the sense that not only $\mathcal{P}$ itself but also $\mathcal{P}'$ is sparse. What about the remaining graph properties?

Lemma 3.4 and Theorem 3.5 suggest that interesting universes for $(\sqsubseteq_*, \leqslant_*^-)$-hierarchies of such other properties (again closed under taking subgraphs) would each lie within a given range of average degree $d = d(n)$, prescribed up to a multiplicative constant. Since 'sparse' means 'with average degree bounded by a constant', Theorem 5.3 is an example of such a hierarchy for the constant function $d = 1$.

References

[1] A. Bondy & M. Simonovits, Cycles of even length in graphs, *J. Combin. Theory B* **16** (1974), 97–105.

[2] R. Diestel, *Graph theory*, 2nd edition, Springer-Verlag 2000 and `http://www.math.uni-hamburg.de/home/diestel/books/graph.theory/download.html`

[3] R. Diestel, Graph Minors I: a short proof of the path width theorem, *Combinatorics, Probability and Computing* **4** (1995), 27–30.

[4] R. Diestel, Relating subsets of a poset, and a partition theorem for WQOs, *Order* **18** (2001), 275–279.

[5] R. Diestel, K.Y. Gorbunov, T.R. Jensen & C. Thomassen, Highly connected sets and the excluded grid theorem, *J. Combin. Theory B* **75** (1999), 61–73.

[6] R. Diestel & D. Kühn, Tree hierarchies, manuscripts 2000.

[7] P. Erdős, Extremal problems in graph theory, in (M. Fiedler, ed.) Theory of graphs and its applications, *Proc. Symp. Smolenice 1963*, Academic Press, New York 1965, 29–36.

[8] N. Robertson & P.D. Seymour, Graph minors. I. Excluding a forest, *J. Combin. Theory B* **35** (1983), 39–61.

[9] N. Robertson & P.D. Seymour, Graph minors. V. Excluding a planar graph, *J. Combin. Theory B* **41** (1986), 92–114.