Cycles in Unions of Transitive Tournaments

Bogdan Alecu
University of Leeds, UK

Pedro Bureo Villafana
University of Leeds, UK

Vadim Lozin
University of Warwick, Coventry, UK

Abstract

A *tournament* is an orientation of a complete graph. A *bitournament* is an orientation of a complete bipartite graph. Let G be an oriented graph which is the (not necessarily disjoint) union of a bounded number of transitive tournaments. When is it possible to remove the arc sets of a bounded number of transitive bitournaments from G in order to make G acyclic?

In this paper, we begin investigating this question and its ties to *lettericity* of graphs and *geometric griddability* of permutations, two independently developed notions that highlight very similar structural features in their respective objects. We explore the problem through the lens of "minimal obstructions", i.e. minimal classes of directed graphs for which this is not possible, and identify an infinite collection of such classes.

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

Keywords and phrases Structural graph theory, directed graphs, permutation patterns, graph lettericity, minimum feedback arc set problem, well-quasi-orderability, transitive tournaments

Digital Object Identifier 10.4230/LIPIcs.WG.2026.3

1 Introduction and motivation

As far as we are aware, the question from the abstract has never been studied directly. However, its motivation can be traced to two independent lines of research which, in different terminologies, ask questions which are structurally analogous to it.

The first of these lines concerns the study of a graph parameter known as *lettericity*. This parameter is located strictly between two other more thoroughly studied parameters, neighbourhood diversity and linear clique-width, in the sense that bounded neighbourhood diversity implies bounded lettericity, which in turn implies bounded linear clique-width (but the converse implications do not hold). The notion of graph lettericity was introduced by Petkovšek in [18] in order to study well-quasi-orderability under the induced subgraph relation. Independently, a similar notion, named *geometric griddability* [1], was introduced in the area of permutation patterns, where it was applied to study enumerative [4], structural [15] and decidability [7] questions. Despite their independent origins, the notions of lettericity of graphs and geometric griddability of permutations share a connection: they highlight very similar structural features in their respective objects. More concretely, bounded lettericity restricted to permutation graphs becomes equivalent to geometric griddability of the corresponding permutations, as was shown in [3]. From these studies arose natural questions: when does a hereditary class of graphs have bounded lettericity; and when is a class of permutations geometrically griddable?

We wish to approach these questions by identifying minimal classes not satisfying the respective properties. This can be viewed as a "Ramsey-type" approach: indeed, Ramsey's Theorem implies that there exist precisely two minimal hereditary classes of graphs with arbitrarily many vertices, namely complete graphs and edgeless graphs. This provides a characterisation of the boundedness of the (very basic) parameter counting a graph's number of vertices: a hereditary class of graphs has unbounded vertex number if and only if it contains one of the two classes given above. Similar Ramsey-type characterisations are available for various other parameters, such as maximum vertex degree, matching number, neighbourhood diversity [14], tree-depth [17], splitness [9], h-index [2]. In particular, there exist precisely three minimal hereditary classes of graphs with arbitrarily large matching number: complete graphs, complete bipartite graphs and graphs of vertex degree at most 1. For the neighbourhood diversity, the number of minimal hereditary classes where this parameter is unbounded is larger, but still finite. On the other hand, there exist infinitely [10], and even uncountably many [8] minimal hereditary classes of unbounded clique-width and linear clique-width. We remark that in general, identifying all minimal hereditary classes where a parameter is unbounded does not necessarily lead to a characterisation of all hereditary classes where the parameter is unbounded, since there may exist classes not containing a minimal class (as hereditary class containment is not a well-founded relation). As we will see, this is not an obstacle in the case of lettericity.

Our contribution

We introduce a new digraph parameter β counting how many arc sets of transitive orientations of complete bipartite graphs need to be removed in order to obtain an acyclic digraph. We sketch how to reduce two open problems from the literature (namely characterising geometric griddability of permutation classes and boundedness of lettericity of hereditary graph classes) to the problem of characterising boundedness of β on hereditary classes of oriented graphs whose arc set is the union of a bounded number of transitive tournaments. While in this work we only explain how to do this in principle, the reduction could straightforwardly be made concrete (that is, with some technical, but not conceptual difficulties) by using the machinery developed in [3]. We solve the special case of the problem where the intersection graph of the transitive tournaments is a cycle. Finally, as a consistency check, we show that any class of unbounded lettericity contains a minimal such class, showing lettericity really can be characterised in this way.

Our methods

Our main result, the characterisation from Theorem 6, is in the form of a specific structure forced by large β in the classes under consideration. The structure, which we call a *chain ring*, is illustrated in Figure 1, and consists of transitive tournaments intersecting in an "interlaced" way and going around in a cycle.

Contrapositively, if we are in a class of digraphs partitionable into transitive tournaments intersecting in a cyclic manner, then forbidding a chain ring puts a bound on β which depends only on the thickness of the forbidden chain ring (and conversely, β increases with the thickness of the structure).

Our proof is by induction, and uses an argument with a topological flavour: starting with a digraph, we wind around the (cyclic) intersection graph of the tournaments following a simple rule for as long as possible. This allows us to partition the digraph into three regions: the two sides, on which we can apply an appropriate induction hypothesis to obtain small β, and a middle region. If the amount of winding we did was small, then we have small β in the middle region as well; and if it was large, then we are able to exhibit a large chain ring.

Structure of the paper

Section 2 covers the relevant preliminaries. In Section 3, we provide some motivation for the question stated in the abstract, which we formulate precisely. In Section 4, we solve a special case of the question. In Section 5, we provide some additional background, and describe how problems from two areas of ongoing research relate to our question. Finally, we discuss future directions in Section 6.

2 Preliminaries

For $n \in \mathbb{N}$, we write $[n] := \{1, \dots, n\}$.

We will primarily work with *directed graphs*, or *digraphs* for short. The vertex set and the arc set of a digraph G are denoted by $V(G)$, respectively $\vec{E}(G)$ (the arrow is purely decorative, to remind the reader that the arcs have a direction – they are ordered pairs of vertices). We write uv to mean (u, v) ("the arc from u to v"). We will use standard terminology and notation such as (in- and out-)neighbourhood, . . . (see e.g. [11]).

An *oriented graph* is a directed graph with no directed cycles of length at most 2. A standard way of describing oriented graphs is by beginning with an undirected graph $G = (V, E)$, and specifying an *orientation*, that is, a choice of direction for each edge of G. We use the term "orientation" to also refer to the resulting oriented graph. Note that when talking about undirected graphs, we will use the term "edges" rather than "arcs", to mitigate any ambiguities.

A *tournament* is an orientation of the complete graph K_n (for some $n \in \mathbb{N}$), and a *bitournament* is an orientation of the complete bipartite graph $K_{s,t}$ (for some $s, t \in \mathbb{N}$).

Given $k \in \mathbb{N}$ and a family $\{X_i\}_{i\in[k]}$ of sets, its *intersection graph* is the graph with vertex set $[k]$, and with ij an edge if and only if $X_i \cap X_j \neq \varnothing$. We will abuse notation and refer to "the intersection graph of the X_i" to mean the intersection graph of $\{X_i\}_{i\in[k]}$.

A directed graph G is *transitive* if for every $u, v, w \in V(G)$, we have that $uv, vw \in \vec{E}(G) \implies uw \in \vec{E}(G)$. It is acyclic if it contains no directed cycles. We remind the reader the folklore fact that any transitive tournament G can be identified with a strict total order on $V(G)$ (explicitly, a transitive tournament is acyclic, so we can produce a *topological ordering* of its vertices, in which all arcs point "in the same direction"). We also remark that a transitive bitournament has all arcs oriented from one part of the bipartition to the other.

Two kinds of graph containment will play important roles: the *subgraph* containment, and the *induced subgraph* containment. A (di)graph H is a subgraph of a (di)graph G if it can be obtained from G by removing vertices and arcs/edges. It is an induced subgraph of G if it can be obtained from G by only removing vertices. The subgraph of G induced by a set $U \subseteq V(G)$ (that is, the (di)graph obtained from G by removing every vertex not in U) is denoted by $G[U]$. If G does not contain an induced subgraph isomorphic to a (di)graph H, we say that G is H-free, or that G excludes H, or that H is a forbidden induced subgraph for G.

A *class* (or *family* or *property*) of (di)graphs is a collection of (di)graphs closed under isomorphisms. It is *hereditary* if it is closed under taking induced subgraphs. From general theory, any hereditary class $\mathcal{X}$ can be uniquely described by its set of minimal forbidden induced subgraphs (that is, the (di)graphs not in $\mathcal{X}$ minimal under the induced subgraph relation). The class of all (di)graphs excluding (di)graphs in a set S is denoted by $\mathrm{Free}(S)$ (and we may omit the set brackets).

For succinctness, we will sometimes abuse terminology and identify oriented graphs with their arc sets, so that e.g. "removing a transitive bitournament" means "removing the arc set of a transitive bitournament"; this is not problematic, since in our setting, we do not lose generality by assuming that the oriented graphs we consider have no isolated vertices.

We also include here a reminder on terminology from order theory. Given a poset $(X, \preceq)$, two elements $x, y \in X$ are *comparable* if $x \preceq y$ or $y \preceq x$, and *incomparable* otherwise. A set of pairwise incomparable elements is called an *antichain*. A (quasi-)ordered poset $(X, \preceq)$ is said to be *well-quasi-ordered* ("wqo" for short) if there are no infinite antichains, and no infinite sequences $x_0 \succ x_1 \succ x_2 \succ \dots$ (this second property is called *well-foundedness*, and is satisfied trivially for the orders we will consider in this paper).

3 A question about oriented graphs

Let us begin by providing some additional perspective on the question stated in the abstract. Let us fix a number $k \in \mathbb{N}$. Suppose that we are given a set X, and that we are given k strict total orders defined "locally" on subsets of X and which are pairwise compatible, in the sense that they agree whenever they compare the same elements of X. Explicitly, let $X_1, \dots, X_k$ be subsets of X; for $i \in [k]$, let $\prec_i$ be a strict total order on X_i, and suppose that, for every $i, j \in [k]$, if $x, y \in X_i \cap X_j$, then $x \prec_i y$ if and only if $x \prec_j y$. We aim to study when it is possible to extend these $\prec_i$ to a common "global" total order $\prec$ on X.

But the answer is easy: we may interpret each $\prec_i$ as the arc set of a transitive tournament, and the compatibility condition says precisely that the union of these arc sets yields an oriented graph. We call this union $G = (X, \bigcup_i \prec_i)$ the *conflict graph* associated to an instance of our problem. It is a standard result that the $\prec_i$ can be extended to a common total order on X precisely when the conflict graph is acyclic. So suppose instead that we are after a more fine-grained measure quantifying *how close* we are to being able to extend the $\prec_i$ to a common total order on X.

One quickly notes a connection between this question and the Minimum Feedback Arc Set problem, in which we are looking for a minimum set of arcs which, upon removal, yields an acyclic graph. However, removing an individual arc does not have a satisfying interpretation in our setting – one which meaningfully preserves the structure of an instance of our locally-defined-orders problem. Instead, we will focus on another operation which we argue is much more natural. Given an instance $(X_i, \prec_i)_{i\in[k]}$ of our question, a *split* consists of choosing a $j \in [k]$ and replacing $(X_j, \prec_j)$ with two totally ordered sets $(X_j^-, \prec_j^-)$ and $(X_j^+, \prec_j^+)$, where X_j^- is an initial segment of $(X_j, \prec_j)$, $X_j^+ := X_j \setminus X_j^-$, and $\prec_j^-, \prec_j^+$ are the restrictions of $\prec_j$ to X_j^-, respectively X_j^+. Intuitively, a split consists of picking one of the total orders, and "splitting" it into two by declaring that everything smaller than or equal to some element becomes incomparable to everything greater than that element. If X_j is the order chosen in the split operation, then we say that the split is "at X_j", or that we are "splitting X_j". Naturally, we obtain the following problem ("MFSS"):

Minimum Feedback Split Set

Instance: A ground set X, and a finite set $\{(X_i, \prec_i)_{i\in[k]}\}$ of strictly totally ordered subsets of X whose orders are pairwise compatible.

Question: What is the smallest number of splits required to obtain an instance for which $\bigcup_i \prec_i$ is acyclic?

It is more convenient to state our results in the language of oriented graphs. Indeed, we note that each $\prec_i$ from an instance of Minimum Feedback Split Set can be interpreted as a transitive tournament, and the union of the $\prec_i$ is then the arc set of an oriented graph. Moreover, in the conflict graph, a split at X_j corresponds to removing a transitive bitournament whose vertex set spans X_j – and we may identify the split with (the arc set of) that transitive bitournament. This leads to the following related problem ("MFTBS"):

Minimum Feedback Transitive Bitournament Set

Instance: An oriented graph G.

Question: What is the smallest number of transitive bitournament subgraphs of G whose removal yields an acyclic graph?

While the algorithmic aspects of the problems are interesting in their own right, in this work we will focus on their structural features. Specifically, we will consider classes of instances, and we will investigate the qualitative boundary between when the number of transitive bitournaments to be removed is bounded across all instances in the class, versus when it is not.

We also note one slight difference between MFSS and MFTBS: in MFSS, we exclusively allow the transitive bitournaments arising from splits as candidates for removal, while in MFTBS, we allow all transitive bitournaments. However, it is routine to check that the intersection of a transitive bitournament with a transitive tournament is contained in a split at that tournament. In particular, for an instance of MFSS with k strictly totally ordered subsets, the number of splits required to obtain an acyclic graph is within a factor of k of the number of transitive bitournaments required. Given that we are primarily only interested in boundedness, it is thus natural to fix k – that is, to study the restriction of MFTBS to oriented graphs which are the union of a bounded number of transitive tournaments. This restriction turns out to still be general enough to yield results towards our motivating problems.

Formally, for an oriented graph $G = (V, \vec{E})$, we let $\tau(G)$ denote the smallest number k of transitive tournaments $\{(X_i, \prec_i)\}_{i\in[k]}$ with $\vec{E} = \bigcup_{i\in[k]} \prec_i$, and we let $\beta(G)$ denote the smallest number ℓ of transitive bitournaments $\{(Y_i \cup Z_i, Y_i \times Z_i)\}_{i\in[\ell]}$ with $Y_i \cap Z_i = \varnothing$ and $Y_i \times Z_i \subseteq \vec{E}(G)$ for all $i \in [\ell]$ such that $\left(V, \vec{E} \setminus \left(\bigcup_{i\in[\ell]} Y_i \times Z_i\right)\right)$ is acyclic. With this terminology, we are ready to formally state a first variant of our question.

▶ **Question 1.** *Let $\mathcal{C}$ be a hereditary class of oriented graphs of bounded τ. When is β bounded in $\mathcal{C}$?*

4 Infinitely many minimal obstructions

We first introduce some useful terminology and notation.

▶ **Definition 2.** *Let $G = (V, \vec{E})$ be an oriented graph. We call a set $S = \{(X_i, \prec_i)_{i\in[k]}\}$ (with $k \in \mathbb{N}$) of transitive tournaments a* transitive tournament cover *of G if $\vec{E} = \bigcup_{i\in[k]} \prec_i$. The* template *of S is the (undirected) intersection graph of the X_i. We call a set $R = \{B_i\}_{i\in[\ell]}$ of transitive bitournaments a* feedback bitournament set *of G if every cycle of G intersects some B_i.*

With this definition, we have

$$\tau(G) = \min\{|S| : S \text{ is a transitive tournament cover of } G\},$$

and

$$\beta(G) = \min\{|R| : R \text{ is a feedback bitournament set of } G\}.$$

▶ **Notation 3.** *Let $G = (V, \vec{E})$ be an oriented graph, and $S = \{(X_i, \prec_i)_{i\in[k]}\}$ a transitive tournament cover of G. For $i \in [k]$ and $y \in X_i$, let*

$$B_i^y := (X_i, \{x \in X_i : x \preceq_i y\} \times \{x \in X_i : x \succ_i y\}).$$

We will look for obstructions to bounded β, in the form of hereditary classes of oriented graphs where the parameter is unbounded. We note that a trivial obstruction is the (hereditary closure) of graphs which are the disjoint union of cycles; but that class has unbounded τ. In what follows, we describe infinitely many obstructions to bounded β with bounded τ, and show that they are minimal.

Fix $k \in \mathbb{N}$ with $k \geq 3$, and let $\mathcal{X}_k$ consist of all oriented graphs which admit a transitive tournament cover whose template is a subgraph of the cycle on k vertices. By definition, $\tau(G) \leq k$ for every $G \in \mathcal{X}_k$. We define a subclass $\mathcal{C}_k$ of $\mathcal{X}_k$ as follows.

For $t \in \mathbb{N}$, let $X := \{x_i^j : i \in [k], j \in [t]\}$. For $i \in [k]$, indices modulo k, let $X_i := \{x_\ell^j : \ell \in \{i, i+1\}, j \in [t]\}$, and define $\preceq_i$ on X_i as follows:

- $x_\ell^{j_1} \preceq_i x_\ell^{j_2}$, whenever $j_1 \leq j_2$ and $\ell \in \{i, i+1\}$,
- $x_i^{j_1} \preceq_i x_{i+1}^{j_2}$ whenever $j_1 \leq j_2$,
- $x_{i+1}^{j_1} \preceq_i x_i^{j_2}$ whenever $j_1 < j_2$.

In particular, $(X_i, \prec_i)$ is a transitive tournament for every $i \in [k]$, and by construction, the oriented graph $C_{k,t} := (X, \bigcup_i \prec_i)$ is in $\mathcal{X}_k$. We call $C_{k,t}$ the *chain ring* of length k and thickness t, and define $\mathcal{C}_k$ as the hereditary closure of $\{C_{k,t}\}_{t\in\mathbb{N}}$. Note that for each $j \leq t$ and $J \subseteq [t]$ with $|J| = j$, the subgraph of $C_{k,t}$ induced by $\{x_\ell^j : (\ell, j) \in [k] \times J\}$ is isomorphic to $C_{k,j}$. We claim that $\mathcal{C}_k$ is minimal of unbounded β.

▶ **Lemma 4.** *If G is a graph in $\mathcal{X}_k$ with transitive tournament cover $(X_i, \prec_i)_{i\in[k]}$ and B is a transitive bitournament, then $|\{i \in [k] : \vec{E}(B) \cap \prec_i \neq \varnothing\}| \leq 4$. Moreover, for every $i \in [k]$, $\vec{E}(B) \cap \prec_i \subseteq B_i^x$ for some $x \in X_i$.*

Proof. The first statement follows immediately from the fact that if $x \in X_i$, then $N(x) \subseteq X_{i-1} \cup X_i \cup X_{i+1}$ (indices modulo k). The second follows from the fact that each $\prec_i$ is a total order. ◀

We note that the bound given by the lemma is tight: when $k = 4$, a transitive bitournament may intersect all $\prec_i$.

▶ **Lemma 5.** *For every $k \geq 3$, the class $\mathcal{C}_k$ has unbounded β.*

Proof. Fix $k \geq 3$ and suppose, for a contradiction, that there exists a natural p with $\beta(C_{k,t}) \leq p$ for every $t \in \mathbb{N}$. In particular, there exists a feedback bitournament set of $C_{k,t}$ of size at most p. By Lemma 4, any transitive bitournament B has nonempty intersection with at most 4 of the $\prec_i$. In particular, there is a feedback bitournament set R of $C_{k,t}$ of size at most $4p$ consisting of transitive bitournaments of the form B_i^y with $i \in [k]$ and $y \in X_i$. Now

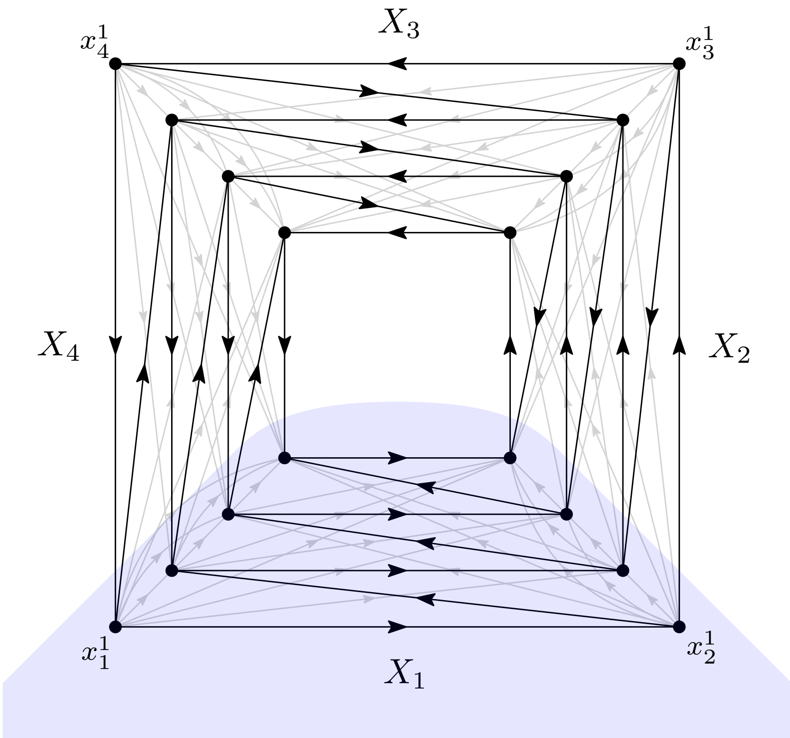

Figure 1 Chain ring $C_{4,4}$. To avoid overloading the figure, we have highlighted the arcs between consecutive elements in each $\prec_i$, and left the others transparent. The set X_1 is shaded.

let $j \in [t]$ and note that, since the vertices x_i^j and x_{i+1}^j are consecutive in $\prec_i$, we have that B_i^y only contains $x_i^j x_{i+1}^j$ if $y = x_i^j$. Since some transitive bitournament of R intersects the cycle $\{x_i^j : i \in [k]\}$, we have $B_i^{x_i^j} \in R$ for some $i \in [k]$. But since j was arbitrary, this gives $|R| \geq t$ – a contradiction, since $t \in \mathbb{N}$ is arbitrary. ◀

It remains to argue that for each $k \geq 3$, the family $\mathcal{C}_k$ is minimal of unbounded β. The way to do this would be to show that $\mathcal{C}_k \cap \mathrm{Free}(C_{k,t})$ has bounded β for every t. We show something stronger, namely that $\mathcal{C}_k$ is the unique obstruction to bounded β in $\mathcal{X}_k$:

▶ **Theorem 6.** *Let $k \geq 3$ and $\mathcal{X} \subseteq \mathcal{X}_k$ be a hereditary class. Then $\mathcal{X}$ has unbounded β if and only if $\mathcal{C}_k \subseteq \mathcal{X}$.*

To prove the theorem, it suffices to show that every graph $G \in \mathcal{X}_k$ which does not contain an induced $C_{k,t}$ has $\beta(G)$ bounded above by some function of t.

▶ **Notation 7.** *For the remainder of this section, we will fix a digraph $G \in \mathcal{X}_k$. By definition of $\mathcal{X}_k$, the digraph G admits a transitive tournament cover $S = \{(X_i, \prec_i)\}_{i \in [k]}$ such that, for $i \neq j$, we have $X_i \cap X_j \neq \varnothing$ only if i and j are consecutive modulo k. For $i \in [k]$, we write $A_i := X_{i-1} \cap X_i$.*

▶ Remark 8. The case where $k = 3$ is special, since it is the only one where some x may belong to three distinct X_i. However, we note that if $G \in \mathcal{X}_3$ and C is a cycle in G, then $V(C) \cap (X_1 \cap X_2 \cap X_3) = \varnothing$. Indeed, if $x \in X_1 \cap X_2 \cap X_3$, then transitivity in each X_i implies that $xy \in \vec{E}(G)$ for each $y \in V(C) \setminus \{x\}$ – a contradiction, since G is oriented, and has at least one in-neighbour in C. Thus if $G \in \mathcal{X}_3$, we have $\beta(G) = \beta(G \setminus (X_1 \cap X_2 \cap X_3))$.

In view of Remark 8, we may – and will, throughout the rest of the section – assume w.l.o.g. that $X_1 \cap X_2 \cap X_3 = \varnothing$ if $k = 3$. In particular, every x belongs to at most two of the X_i.

▶ **Notation 9.** *We say that an arc $uv \in \vec{E}(G)$* crosses *$x \in V(G)$ if there exists $i \in [k]$ with $u \preceq_i x \prec_i v$ (that is, if $uv \in \vec{E}(B_i^x)$). We write S_x for the set of all arcs which cross x. We note S_x is the union of at most two transitive bitournaments (since x belongs to at most two of the X_i). Note that if G has a feedback bitournament set of size k, then by Lemma 4, there exist ℓ vertices $x_1, \ldots, x_\ell$ with $\ell \leq 4k$ such that $\bigcup_{i\in[\ell]} S_{x_i}$ intersects every cycle of G.*

▶ Remark 10. If uv and vw are edges but uw is not, it means that $u \in X_i \setminus X_j$ and $w \in X_j \setminus X_i$ for some $i \in [k]$ and $j \in \{i-1, i+1\}$, while $v \in X_i \cap X_j$. In particular, the internal vertices of a chordless path (of length at least 2) belong to distinct A_i, and a chordless cycle of G intersects each A_i in exactly one vertex (and so if G contains a cycle, then $A_i \neq \varnothing$ for every $i \in [k]$).

▶ **Notation 11.** *In view of Remark 10, we note that chordless cycles and paths of length at least 2 are directed either in order of increasing or decreasing i. In the former case we say the cycle or path is* downwards*, and in the latter, we say it is* upwards*. The following lemma shows that this notion can be extended to all cycles of G, not just chordless ones.*

▶ **Lemma 12.** *Let C be a cycle of G. Then C does not contain both a downwards chordless cycle and an upwards chordless cycle.*

Proof. Suppose, for a contradiction, that C contains both a downwards chordless cycle C_d and an upwards chordless cycle C_u. Using Remark 10, write $V(C_d) = \{x_1, \ldots, x_k\}$ and $V(C_u) = \{y_1, \ldots, y_k\}$, with $x_i, y_i \in A_i$ for $i \in [k]$.

▷ Claim. If $y_i \preceq_i x_i$ for some $i \in [k]$ (indices modulo k), then $y_{i+1} \prec_{i+1} x_{i+1}$. Similarly, if $x_i \preceq_i y_i$, then $x_{i-1} \prec_{i-1} y_{i-1}$. In particular, we have that either $x_i \prec_i y_i$ for every $i \in [k]$, or $y_i \prec_i x_i$ for every $i \in [k]$.

Proof. Suppose that $y_i \preceq_i x_i$. We note that $y_{i+1}y_i$ and x_ix_{i+1} are edges, so by transitivity of $\prec_i$, we have $y_{i+1} \prec_i y_i \preceq_i x_i \prec_i x_{i+1}$ (and in particular $y_{i+1} \prec_{i+1} x_{i+1}$, since $\prec_i$ and $\prec_{i+1}$ agree on A_{i+1}). The second statement follows by symmetry, and the last follows by applying one of the first statements repeatedly. ◁

Using the above claim, assume w.l.o.g. that $x_i \prec_i y_i$ for all $i \in [k]$ (the other case is symmetric). We show that there exists no path from $V(C_u)$ to $V(C_d)$. Indeed, if such a path exists, then we may pick a chordless one, and we note that by transitivity, it cannot have length 1. By Remark 10, it is either downwards or upwards. If it is upwards, we let vv' be the first edge of the path going from a vertex $v \in A_j$ (for some $j \in [k]$) with $x_j \prec_j v$ to a vertex $v' \in A_{j-1}$ with $v' \preceq x_{j-1}$. We note that $x_j \prec_{j-1} v \prec_{j-1} v' \preceq_{j-1} x_{j-1}$, a contradiction (since C_d is downwards). So the path is downwards, in which case we let ww' be the first edge of the path going from a vertex $w \in A_j$ (for some $j \in [k]$) with $y_j \preceq_j w$ to a vertex $w' \in A_{j+1}$ with $w' \prec_{j+1} y_{j+1}$. But then we note that $y_j \preceq_j w \prec_j w' \prec_j y_{j+1}$, a contradiction (since C_u is upwards). ◀

With this in mind, we say a cycle is *downwards* if it contains a downwards chordless cycle, and *upwards* otherwise.

The following technical lemma shows that we do not need to worry about vertices belonging to only one X_i.

▶ **Lemma 13.** *Let* $G' := G[\bigcup_{i\in[k]} A_i]$*. We have* $\beta(G) \leq 24\beta(G')$*.*

Proof. By Lemma 4, there exists a set Y of size ℓ with $\ell \leq 4\beta(G')$ such that $\bigcup_{x\in Y} S_x \cap \vec{E}(G')$ intersects every cycle of G'.

For $x \in Y$, we let $j(x)$ be such that $x \in A_{j(x)}$. We let x^- be the greatest element of $A_{j(x)-1}$ with $x^- \prec_{j(x)-1} x$, and similarly, we let x^+ be the greatest element of $A_{j(x)+1}$ with $x^+ \prec_{j(x)} x$. We let $Y^- := \{x^- : x \in Y\}$, and define Y^+ analogously.

We claim that $\bigcup_{x\in Y^-\cup Y\cup Y^+} S_x$ intersects every cycle of G (and so, since S_x is the union of at most two transitive bitournaments for each x, we have $\beta(G) \leq 6\ell \leq 24\beta(G')$, as required).

To see this, let C be a cycle in G. If $v \in V(C) \setminus V(G')$, then the vertices u and w appearing before and after v in C belong to the same X_i as v; in this case, we construct a shorter cycle by replacing the arcs uv, vw with the arc uw. Repeated applications of this substitution yield a cycle C' in G' with $V(C') = V(C) \cap V(G')$ and which visits its vertices in the same order in which they appear in C. By assumption, some edge uw of C' crosses some $x \in Y$. Let $u = v_1, \ldots, v_r = w$ be the vertices between u and w in C. If u and w belong to different (consecutive) A_i, then some edge in the path $v_1, \ldots, v_r$ crosses x. Otherwise, w.l.o.g. u and w both belong to A_i for some i, we have $v_2, \ldots, v_{r-1} \in X_i \setminus X_{i-1}$, and we have $x \in X_{i-1} \setminus X_i$ (see Figure 2 for an illustration; the other cases are similar). In this case, some edge in the path $v_1, \ldots, v_r$ crosses x^+.

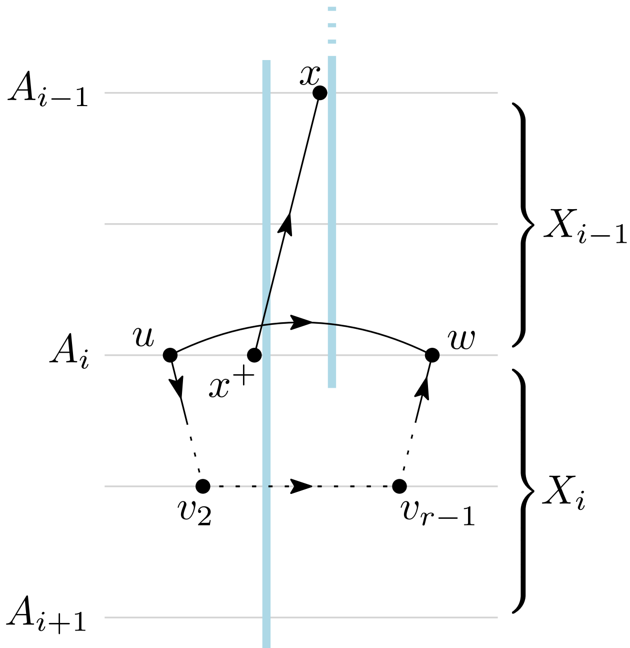

Figure 2 Illustration for the proof of Lemma 13. Note that we may have $x^+ = u$ and $v_2 = v_{r-1}$. The splits at x and x^+ are represented by vertical lines immediately to their right.

◀

In view of Lemma 13, we may – and will, throughout the rest of the section – assume w.l.o.g. that $V(G) = \bigcup_{i\in[k]} A_i$.

▶ **Notation 14.** *We denote by* $\prec$ *the union of the total orders on each* A_i *which are the restrictions of each* $\prec_i$ *to* A_i *(that is,* $x \prec y$ *if there exists* i *with* $x, y \in A_i$ *and* $x \prec_i y$*).*

Given a walk W in G with vertices $x_1, \dots, x_q$ appearing in this order, for each $i \in [q]$, we let $a(i) \in [k]$ be such that $x_i \in A_{a(i)}$. For every $i \in [q-1]$, we let $b(i) \in \{-1, 0, 1\}$ be such that $b(i) = a(i+1) - a(i) \pmod k$. We define the *winding number* of W as

$$\gamma(W) := \frac{1}{k} \sum_{i=1}^{q-1} b(i).$$

Note that for any closed walk, and in particular for any cycle C of G, the winding number is an integer. The $\frac{1}{k}$ factor is a normalisation constant ensuring that a cycle going once in the positive direction of the indices has winding number 1.

▶ **Lemma 15.** *If W is a closed walk in G of length at least 1, then $\gamma(W) \neq 0$.*

Proof. We argue inductively that for any walk W with vertices $x_1, \dots, x_q$ ($q \geq 2$) and with $\gamma(W) = 0$, we have $x_1 \prec x_q$. Indeed, the statement is clear when $q = 2$. Suppose now that W is a walk of length at least 2 with vertices $x_1, \dots, x_q$ and with $\gamma(W) = 0$. Then W visits at least two A_i (otherwise the statement is clear), and since $\gamma(W) = 0$, there exist $i_1, i_2 \in [q] \setminus \{1, q\}$ with $i_1 \leq i_2$ such that $a(i) = a(i')$ for every $i, i' \in \{i_1, \dots, i_2\}$, and

$$(a(i_1) - a(i_1 - 1))(a(i_2 + 1) - a(i_2)) = -1$$

(that is, a point in the walk where the walk "doubles back" after possibly spending some time in some A_i). But then the walk W' obtained from W by replacing the subwalk $x_{i_1-1} x_{i_1} \dots x_{i_2} x_{i_2+1}$ with the arc $x_{i_1-1} x_{i_2+1}$ is a shorter walk from x_1 to x_q with $\gamma(W') = 0$, showing $x_1 \prec x_q$. In particular, W is not a closed walk, as required. ◀

▶ **Lemma 16.** *If G contains no $C_{k,t}$, then $\beta(G)$ is bounded above by a function of t.*

Proof. We proceed by induction on t. The statement we will show by induction for $t \geq 1$ is that if G does not contain a downwards copy of $C_{k,t}$ (that is, a copy of $C_{k,t}$ in which the cycles $\{x_i^j : j \in [t]\}$ are downwards for every $i \in [k]$), then there exists $f(t)$ with the property that there are at most $f(t)$ vertices $x_1, \dots, x_r$ such that every downwards cycle has an edge crossing some x_i. The theorem will then follow by symmetry and by Lemma 4.

Noting that $C_{k,1}$ is a cycle of length k, that any cycle contains a chordless cycle, and that the only chordless cycles of G have length k, the statement is trivially true when $t = 1$.

Fix $t \geq 1$, and suppose the statement is true for t; we aim to show it for $t + 1$, that is, there exists $f(t+1) \in \mathbb{N}$ such that if there is no downwards $C_{k,t+1}$ in G, then there are at most $f(t+1)$ vertices such that every downwards cycle of G has an edge crossing one of them.

The intuition behind our inductive step is straightforward; we present it here before providing the more formal write-up. We may assume that G contains a downwards cycle C of length k intersecting every A_i (otherwise we are in the base case). Starting from the vertex of C belonging to A_i (for some i), we go "in a spiral", repeatedly choosing the least (according to $\prec_i$) vertex of A_{i+1} which is greater than x. This procedure terminates, yielding a downwards cycle $C_{\min}$, which is "less than or equal to" C in every order. An analogous process yields a downwards cycle $C_{\max}$. This splits G into 3 regions: nodes "smaller" than $C_{\min}$, nodes "greater" than $C_{\max}$, and a middle region consisting of the nodes between the two cycles. We check that no downwards cycle of G intersects two of the regions; then the induction hypothesis applies to the first two regions. To deal with cycles in the middle region, we show that every cycle has an edge crossing a vertex of our spiral; thus if our spirals are short, we are done, and if they are long, we find a downwards copy of the forbidden graph $C_{k,t+1}$, a contradiction.

To make this rigorous, we begin by fixing a chordless downwards cycle C of G (if no such cycle exists, then we are in the base case). We assume without loss of generality that C has edges v_iv_{i+1} for $i \in [k]$ (indices are modulo k throughout the rest of this proof), with $v_i \in A_i$ for every $i \in [k]$. Given a vertex $x \in A_i$, we let $s(x)$ denote the least (w.r.t. $\prec_i$) vertex $y \in A_{i+1}$ with $x \prec_i y$ if such a vertex exists, and $p(x)$ the greatest (w.r.t. $\prec_{i-1}$) vertex $y \in A_{i-1}$ with $y \prec_{i-1} x$.

With this notation, we note immediately from transitivity of the $\prec_i$ that $x \preceq y$ implies $s(x) \preceq s(y)$ and $p(x) \preceq p(y)$ whenever these exist. Moreover, $s(v)$ and $p(v)$ exist for every vertex v of C. It follows that the sequences $(s^j(v_1))_{j\in\mathbb{N}}$ and $(p^j(v_1))_{j\in\mathbb{N}}$ (where the superscript denotes the number of applications) exist, and become eventually periodic (since G is finite). The period of the first sequence is a downwards cycle $C_{\min}$ with vertices $u_i \in A_i$ for $i \in [k]$, and the period of the second one is a downwards cycle $C_{\max}$ with vertices $w_i \in A_i$, for $i \in [k]$. See Figure 3 for an illustration. We note that $u_i \preceq v_i \preceq w_i$ for every $i \in [k]$.

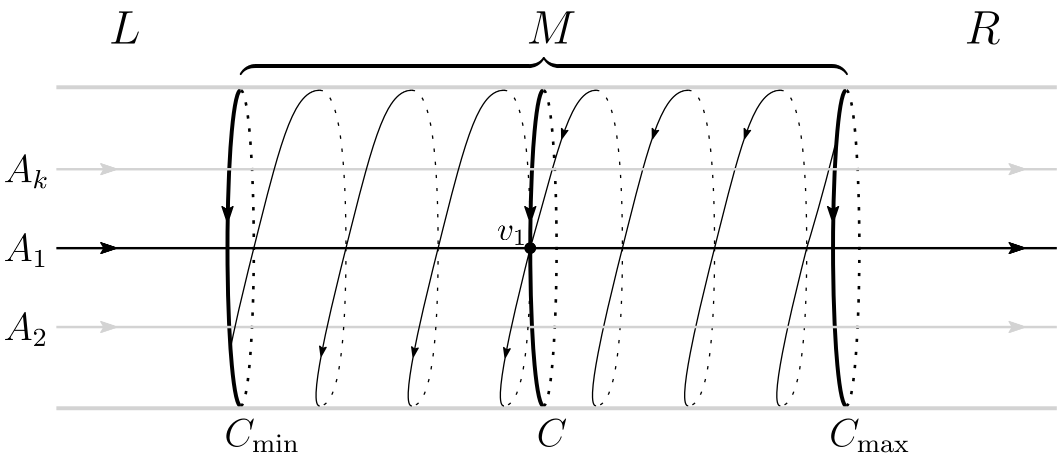

Figure 3 Illustration for the proof of Lemma 16. The essence of the proof lies in showing that any non-trivial closed walk in M must cross the spiral. From a more high-level perspective, consider the local orders on the X_i. They order each A_i from left to right in the illustration. On the cylinder itself, those local orders do not have a common extension to a global order: it is possible to follow them around the cylinder and to end up to the left of where we began (this is effectively what the spiral does). However, if we were to "cut" the cylinder along the spiral (which is effectively what the splits do), we would obtain a space in which the orders *are* globally compatible.

We let $L := \{x \in V(G) : x \prec u \text{ for some } u \in C_{\min}\}$, $R := \{x \in V(G) : x \succ w \text{ for some } w \in C_{\max}\}$, and $M := \{x \in V(G) : u \preceq x \preceq w \text{ for some } u \in C_{\min} \text{ and } w \in C_{\max}\}$. Note that $\{L, M, R\}$ is a partition of $\bigcup_{i\in[k]} A_i = V(G)$.

▷ Claim. If C' is a downwards cycle, then $V(C')$ is contained in one of L, M and R.

Proof. Let C' be a downwards cycle, and note that it intersects L, M or R. Suppose for a contradiction that $V(C') \cap L \neq \varnothing$ and $V(C') \cap (M \cup R) \neq \varnothing$. Then there exists a $z_i \in C' \cap (M \cup R) \cap A_i$ and a z_{i+1} in $C' \cap L \cap A_{i+1}$. But by construction, this is impossible, since then $u_i \preceq z_i$, but $s(z_i) \preceq z_{i+1} \prec u_{i+1} = s(u_i)$. An analogous argument shows any downwards cycle intersecting R is contained in R, and the claim follows. ◁

▷ Claim. $G[L]$ and $G[R]$ do not contain any downwards $C_{k,t}$

Proof. If $G[L]$ contains such an induced $C_{k,t}$, then together with $C_{\min}$, it induces a downwards $C_{k,t+1}$ – a contradiction. An analogous argument applies to $G[R]$. ◁

▷ Claim. $s^{i+k}(v_1) \preceq s^i(v_1)$ for every $i \geq 0$, and $p^{i+k}(v_1) \succeq p^i(v_1)$ for every $i \geq 0$. Moreover, the sequences $(s^j(v_1))_{j\in\mathbb{N}}$ and $(p^j(v_1))_{j\in\mathbb{N}}$ each have fewer than $2k(t+1)$ distinct elements.

Proof. For the first statement, we note that for $i \in \{0,\dots,k-1\}$, we have $s^i(v_1) \preceq v_{i+1}$; in particular, $s^k(v_1) = s(s^{k-1}(v_1)) \preceq v_1 = s^0(v_1)$ (by definition of s and transitivity, since $s^{k-1}(v_1) \preceq v_k$, and v_kv_1 is an edge). Moreover, if $s^{i+k}(v_1) \preceq s^i(v_1)$ for some i, then $s^{i+k+1}(v_1) = s(s^{i+k}(v_1)) \preceq s(s^i(v_1)) = s^{i+1}(v_1)$, so the first statement follows by induction and symmetry. Now write $x_j := s^j(v_1)$, and suppose for a contradiction that the elements $x_1,\dots,x_{2k(t+1)}$ are all distinct. For $i \in \{0,\dots,t\}$, let $D_i := \{x_{2ik+1},\dots,x_{(2i+1)k}\}$; then $G[\bigcup_{i\in\{0,\dots,t\}} D_i]$ is a downwards $C_{k,t+1}$, a contradiction. The case for the other sequence is similar. ◁

The second claim and the induction hypothesis show that the downwards cycles contained in L and R have an edge crossing one of at most $2f(t)$ vertices of G. By the first claim, it remains to deal with downwards cycles contained in M.

We show that every downwards cycle C' contained in M contains an edge crossing $s^j(v_1)$ or $p^j(v_1)$ for some j. By the third claim, this is enough: we may set $f(t+1) = 2f(t)+4k(t+1)$. For ease of notation, write $s^{-j}(v_1) := p^j(v_1)$ (note, however, that s and p are *not* inverses of each other as functions from V to V).

Let C' be a downwards cycle contained in M. By definition of crossing, every edge uv crosses its tail u, so we may assume that C' does not intersect $\{s^j(v_1)\}_{j\in\mathbb{Z}}$. Explicitly, letting x be a vertex of $V(C')\cap A_i$ for some i, we may assume that $s^{(r+1)k+i-1}(v_1) \prec x \prec s^{rk+i-1}(v_1)$ for some $r \in \mathbb{Z}$. Since no edge of C' crosses any vertex in $\{s^j(v_1)\}_{j\in\mathbb{Z}}$, there are three cases for the successor x' of x in C':

- If x' belongs to the same A_i as x, then $s^{(r+1)k+i-1}(v_1) \prec x' \prec s^{rk+i-1}(v_1)$.
- If x' belongs to A_{i-1}, then $s^{(r+1)k+i-2}(v_1) \prec x' \prec s^{rk+i-2}(v_1)$.
- If x' belongs to A_{i+1}, then $s^{(r+1)k+i}(v_1) \prec x' \prec s^{rk+i}(v_1)$.

But by Lemma 15, the winding number of C' is some $r' \neq 0$. Repeated applications of the above three points then yield $s^{(r+r'+1)k+i-1}(v_1) \prec x \prec s^{(r+r')k+i-1}(v_1)$, a contradiction which finishes the proof. ◀

We note that our "spiral" argument, and indeed its whole set-up using winding numbers, is essentially a discretisation of the fact that cutting out a helix from the surface of an infinite cylinder produces a contractible space. It stands to reason that the proof could be rewritten in these terms by considering embeddings of graphs from $\mathcal{X}_k$ into the surface of an infinite cylinder, then arguing that our spiral construction embeds into a helix and that downwards cycles embed into non-null homotopic loops. Finally, we conclude:

Proof of Theorem 6. The "if" direction follows immediately from Lemma 5. The "only if" direction follows from Lemma 16. ◀

5 β, lettericity, and geometric griddability

In this section, we discuss how characterising bounded lettericity and geometric griddability relates to characterising bounded β. We begin with geometric griddability, since it is easier to illustrate; the discussion for lettericity is analogous, and we will highlight specific points as appropriate.

Permutations and griddability

Our parameter β has ties to the study of permutation patterns. We omit the relevant background; the reader who is not familiar with it is encouraged to consult, e.g., [1] and [3]. The essence of it is that there are two successively stronger notions of well-behavedness for a class of permutations, namely *monotone griddability by a partial multiplication matrix*, and *geometric griddability.* The question relevant to us is about the boundary between these two notions:

▶ **Question 17.** *Which classes of permutations monotone griddable by a partial multiplication matrix are minimal non-geometrically griddable?*

The key insight towards reducing Question 17 to Question 1 is that given a permutation π gridded by an $s \times t$ partial multiplication matrix, we can construct an oriented graph G_π on the elements of the permutation with $\tau(G_\pi)$ at most $s+t$, by letting each row and column of the partial multiplication matrix correspond to a transitive tournament ordered according to that row or column's sign. In other words, the monotone gridding defines an oriented graph together with a transitive tournament cover.

We next note that constructing a new gridding of the permutation by adding a vertical or horizontal line to the starting gridding (and considering the ensuing oriented graph and transitive tournament cover) corresponds precisely to the splitting operation described in Section 3.

Finally, if the oriented graph associated to the gridded permutation is acyclic, we can follow the arguments from [1] and use a topological ordering of the graph to produce a geometric gridding of the permutation. This can be done explicitly by defining a distinguished corner in each cell of the monotone gridding, and placing points on the main diagonal/antidiagonal at increasing distance from the cell's distinguished corner, according to the topological ordering; we omit the details.

The above discussion outlines a strategy for characterising geometric griddability of a class $\mathcal{X}$ monotone griddable by a partial multiplication matrix M, given an answer to Question 1:

- Consider the class $G_{\mathcal{X}}$ of oriented graphs G_π (for $\pi \in \mathcal{X}$) defined as above.
- If $G_{\mathcal{X}}$ has bounded β, then for every permutation π, we can add a bounded number of lines to the gridding to make the corresponding oriented graph acyclic.
- This means that the permutation is geometrically griddable by the resulting gridding. Since there is only a bounded number of ways to produce a new gridding by adding a bounded number of lines, this means the class $\mathcal{X}$ is geometrically griddable.
- If $G_{\mathcal{X}}$ has unbounded β, then we identify a subclass $G_{\mathcal{M}}$ minimal of unbounded β using the answer to Question 1.
- We backtrack through our construction of G_π to produce a corresponding subclass $\mathcal{M}$ of $\mathcal{X}$ minimal non-geometrically griddable.

As a special case for the above strategy, Theorem 6 can in principle be used to explicitly describe all minimal non-geometrically griddable permutation classes which are monotonically griddable by a partial multiplication matrix whose cell graph contains a unique cycle. We note that these unicyclic classes have recently received attention in [5]. Indeed, the authors of [5] use the graph G_π described above (under the name "orientation digraph") in order to study wqo in these classes, as well as some other properties. Their analysis reveals similar structure to the one we identified in the previous section: their "coils" are analogous to the spiral construction we use in the proof of Lemma 16. A first notable difference between their work and ours is that they use a weaker assumption (the absence of long coils) in order to derive a weaker conclusion (wqo of the class), while we use a stronger assumption (the absence of thick chain rings) in order to derive a stronger conclusion (geometric griddability of the class). A second notable difference is the focus of the work: [5] keeps things permutation-specific and uses the graph G_π as an auxiliary tool, whereas we attempt to abstract away as much as possible, and study the oriented graphs in their own right. Their differences notwithstanding, the natural continuation of the two approaches is the same: generalising beyond the one-cycle setting. The authors of [5] give arguments and references supporting the claim that lifting wqo results to a polycyclic setting is "an order of magnitude more complex". It is our hope that trying to understand β in the oriented graph setting will make that goal more tractable.

Letter graphs and lettericity

The second notion motivating Question 1 is a graph parameter called *lettericity*, introduced in [18]. Explicitly, we have:

▶ **Question 18.** *When does a hereditary class of graphs have bounded lettericity?*

We will avoid defining the parameter, since for the purposes of our discussion, we only need to know a few facts about it. The curious reader is invited to consult [18]. The most relevant fact is that a permutation class is geometrically griddable if and only if the corresponding class of inversion graphs has bounded lettericity [3]. And indeed, similarly to how one reduces the characterisation of geometric griddability to Question 1, one can in principle relate Question 18 to it. The following outlines a strategy for doing so.

(i) We define a parameter λ for graphs whose boundedness is analogous to monotone griddability by a partial multiplication matrix.
(ii) We identify the obstructions to bounded lettericity when λ is unbounded.
(iii) When λ is bounded, we can define oriented digraphs of bounded τ for which bounded β is equivalent to bounded lettericity.

There is one more reason why discussing lettericity is relevant. In general, there is no guarantee that listing all minimal classes where a parameter is unbounded will lead to a characterisation of when the parameter is unbounded. It could be the case that a parameter p is unbounded in a class $\mathcal{X}$, but $\mathcal{X}$ does not contain a minimal class where p is unbounded – instead, it might contain an infinite chain $\mathcal{X}_0 \supsetneq \mathcal{X}_1 \supsetneq \dots$ whose intersection is of bounded p. We show here that lettericity is well-behaved in this regard. We use the following facts (both shown in [18]):

- If a class $\mathcal{X}$ has bounded lettericity, then $\mathcal{X}$ is wqo by the induced subgraph relation.
- For any graph G and vertex $v \in V(G)$, we have $\operatorname{let}(G) \leq 2\operatorname{let}(G - v) + 1$ (where $\operatorname{let}(G)$ denotes the lettericity of G).

▶ **Theorem 19.** *Any hereditary class $\mathcal{X}$ of unbounded lettericity contains a minimal hereditary class of unbounded lettericity.*

Proof. If $\mathcal{X}$ is a minimal class of unbounded lettericity, we are done. Otherwise, $\mathcal{X}$ contains a graph G such that $\mathcal{X} \cap \operatorname{Free}(G)$ still has unbounded lettericity. Pick a graph $G_0 \in \mathcal{X}_0 := \mathcal{X}$ with this property that has the minimum possible number of vertices (breaking ties arbitrarily), and put $\mathcal{X}_1 := \mathcal{X}_0 \cap \operatorname{Free}(G_0)$. Repeat the process for as long as possible, putting $\mathcal{X}_{k+1} := \mathcal{X}_k \cap \operatorname{Free}(G_k)$, where $G_k \in \mathcal{X}_k$ is a minimum graph such that $\mathcal{X}_{k+1}$ has unbounded lettericity.

If the process stops at some k, then $\mathcal{X}_k$ is the minimal class we were looking for. Otherwise, the process goes on forever, and we get an infinite strictly descending chain $\mathcal{X}_0 \supset \mathcal{X}_1 \supset \dots$ of classes, all of which have unbounded lettericity. Let $\mathcal{X}_{\text{lim}}$ be their intersection, that is, $\mathcal{X}_{\text{lim}} = \mathcal{X} \cap \text{Free}(G_0, G_1, \dots)$.

Note that this limit class cannot have bounded lettericity. Indeed, suppose graphs in $\mathcal{X}_{\text{lim}}$ have lettericity at most t. The graphs $G_0, G_1, \dots$ are pairwise incomparable by the induced subgraph relation, because for each $i < j$, G_j is G_i-free (by construction) and $|G_i| \leq |G_j|$ (by the choice of the graphs). Additionally, for any $i \in \mathbb{N}$ and $v \in G_i$ we have $G_i - v \in \mathcal{X}_{\text{lim}}$, and hence the lettericity of G_i is at most $2t + 1$. Thus, we obtain an infinite antichain of graphs of bounded lettericity, which is a contradiction, since classes of bounded lettericity are well-quasi-ordered.

Moreover, $\mathcal{X}_{\text{lim}}$ is a minimal class of unbounded lettericity. To see this, note first that, by construction, $|G_n| \to \infty$ as $n \to \infty$, since there are only finitely many graphs of a given size. Suppose we can forbid $G \in \mathcal{X}_{\text{lim}}$ with $|G| = k$, and we are still left with a class of unbounded lettericity. But then, by construction, G would have appeared in the sequence (G_i) before any graphs of size at least $k + 1$, contradicting $G \in \mathcal{X}_{\text{lim}} = \mathcal{X} \cap \text{Free}(G_1, G_2, \dots)$. ◀

6 Future directions

In this paper, we have shown that any hereditary class of unbounded lettericity contains a minimal hereditary class of unbounded lettericity and identified infinitely many such classes. However, a complete description of minimal hereditary classes of unbounded lettericity remains a challenging open question. The simplest unsolved instance is the case where the template of the transitive tournament cover is (a subgraph of) two cycles intersecting at one vertex. We note that as far as wqo is concerned, even in this relatively simple setting, the behaviour of objects can be very complex. This is evidenced in [16], where long antichains of permutations are identified which use the two-cycle structure in a non-trivial, aperiodic way. We believe, however, that characterising boundedness of β is more tractable. Forbidding chain rings as induced subgraphs is not enough: one can construct examples of large β where they appear as non-induced subgraphs. In fact, it is possible to generalise Lemma 5 to show that the presence of thick chain rings as subgraphs implies large β (note that this is not immediately obvious, since β may increase when removing arcs). As such, we can formulate a more concrete version of Question 1:

▶ **Question 20.** *Is it true that if an oriented graph G contains no chain ring of thickness t as a subgraph, then $\beta(G)$ is bounded by a function of t?*

Another direction is to try and understand oriented graphs of bounded τ from the perspective of well-quasi-orderability. It is possible to define "chain spirals" analogous to the coils from [5]. A natural question is then whether the results from [5] generalise:

▶ **Question 21.** *Is it true that if graphs from a class $\mathcal{X}$ of bounded τ contain no long chain spirals as subgraphs, then $\mathcal{X}$ is wqo under induced subgraphs?*

We also note that as far as the authors are aware, the parameters τ and β have not been studied before. Given that their origin is closely linked to the well-studied Minimum Feedback Arc Set problem, it would be worth exploring their algorithmic aspects.

Regarding graph lettericity, it is worth observing that this notion can be useful for representing graphs even in classes where lettericity is unbounded. For instance, the class of bipartite permutation graphs is also known as Parikh word representable graphs [19], which is equivalent to representing them as letter graphs over an infinite alphabet. A similar representation is available for all permutation graphs. Both these classes play a critical role in the area of graph parameters: the class of bipartite permutation graphs is a minimal hereditary class of unbounded clique-width [13], while the permutation graphs form a minimal hereditary class of unbounded twin-width [6]. Identifying other critical classes by means of letter graphs with a well-structured infinite decoder is an interesting research direction.

Finally, we ask whether the results obtained in this paper can be adapted to the notion of generalized lettericity, which was recently introduced in [12].

References

1 M.H. Albert, M.D. Atkinson, M. Bouvel, N. Ruškuc, and V. Vatter. Geometric grid classes of permutations. *Trans. Amer. Math. Soc*, 365:5859–5881, 2013.

2 B. Alecu, V. E. Alekseev, A. Atminas, V. Lozin, and V. Zamaraev. Graph parameters, implicit representations and factorial properties. *Discrete Math.*, 346(10):113573, 2023. Paper No. 113573, 17 pp. `doi:10.1016/J.DISC.2023.113573`.

3 B. Alecu, R. Ferguson, M. Kanté, V. Lozin, V. Vatter, and V. Zamaraev. Letter graphs and geometric grid classes of permutations. *SIAM J. Discrete Math.*, 36(4):2774–2797, 2022. `doi:10.1137/21M1449646`.

4 D. Bevan. Growth rates of geometric grid classes of permutations. *Electron. J. Combin.*, 21(4):4.51, 2014. Paper 4.51, 17 pp.

5 D. Bevan, R. Brignall, and N. Ruškuc. On cycles in monotone grid classes of permutations, 2024. To appear in *Combinatorial Theory*. `arXiv:2410.05834`.

6 É. Bonnet, E. J. Kim, S. Thomassé, and R. Watrigant. Twin-width I: Tractable FO model checking. *J. ACM*, 69(1):1–46, 2021. `doi:10.1145/3486655`.

7 S. Braunfeld. Decidability in geometric grid classes of permutations. *Proc. Amer. Math. Soc.*, 153(3):987–1000, 2025.

8 R. Brignall and D. Cocks. Uncountably many minimal hereditary classes of graphs of unbounded clique-width. *Electron. J. Combin.*, 29(1):P1.63, 2022. Paper P1.63. `doi:10.37236/10483`.

9 M. Chudnovsky and P. Seymour. Extending the Gyárfás–Sumner conjecture. *J. Combin. Theory Ser. B*, 105:11–16, 2014. `doi:10.1016/J.JCTB.2013.11.002`.

10 A. Collins, J. Foniok, N. Korpelainen, V. Lozin, and V. Zamaraev. Infinitely many minimal classes of graphs of unbounded clique-width. *Discrete Applied Math.*, 248:145–152, 2018. `doi:10.1016/J.DAM.2017.02.012`.

11 R. Diestel. *Graph Theory*. Springer, 5 edition, 2017.

12 Z. Feng, H. Fernau, K. Mann, I. Raman, and S. C. Sacher. Generalized lettericity of graphs. In *Algorithms and Discrete Applied Mathematics (CALDAM 2025)*, volume 15536 of *Lecture Notes in Computer Science*, pages 134–146. Springer, 2025. `doi:10.1007/978-3-031-83438-7_12`.

13 V. Lozin. Minimal Classes of Graphs of Unbounded Clique-Width. *Ann. Comb.*, 15(4):707–722, 2011.

14 V. Lozin. Hereditary classes of graphs: a parametric approach. *Discrete Applied Math.*, 325:134–151, 2023. `doi:10.1016/J.DAM.2022.10.016`.

15 V. Vatter M.H. Albert, N. Ruškuc. Inflations of geometric grid classes of permutations. *Israel J. Math.*, 205(1):73–108, 2015.

16 M. Murphy and V. Vatter. Profile classes and partial well-order for permutations. *Electron. J. Combin.*, 9(2):R17, 2002. Paper R17. `doi:10.37236/1689`.

17 J. Nešetřil and P. Ossona de Mendez. On low tree-depth decompositions. *Graphs and Combinatorics*, 31:1941–1963, 2015. `doi:10.1007/S00373-015-1569-7`.

18 M. Petkovšek. Letter graphs and well-quasi-order by induced subgraphs. *Discrete Math.*, 244:375–388, 2002. `doi:10.1016/S0012-365X(01)00094-2`.

19 W. C. Teh, Z. C. Ng, M. Javaid, and Z. J. Chern. Parikh word representability of bipartite permutation graphs. *Discrete Appl. Math.*, 282:208–221, 2020. `doi:10.1016/J.DAM.2019.12.005`.