

Manuscript version: Published Version

The version presented in WRAP is the published version (Version of Record).

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/168183>

How to cite:

The repository item page linked to above, will contain details on accessing citation guidance from the publisher.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

LETTER GRAPHS AND GEOMETRIC GRID CLASSES OF PERMUTATIONS*

BOGDAN ALECU[†], ROBERT FERGUSON[‡], MAMADOU MOUSTAPHA KANTÉ[§],
VADIM V. LOZIN[†], VINCENT VATTER[‡], AND VICTOR ZAMARAEV[¶]

Abstract. We uncover a connection between two seemingly unrelated notions: lettericity, from structural graph theory, and geometric griddability, from the world of permutation patterns. Both of these notions capture important structural properties of their respective classes of objects. We prove that these notions are equivalent in the sense that a permutation class is geometrically griddable if and only if the corresponding class of inversion graphs has bounded lettericity.

Key words. inversion graphs, permutation patterns, well-quasi-order

MSC codes. 05C75, 05A05, 68R10

DOI. 10.1137/21M1449646

1. Introduction. Structural graph theory and the study of permutation patterns are two flourishing areas of mathematics. They share a few conceptual similarities; among others, the widespread featuring of forbidden substructures, and the prominent usage of decompositions. Nevertheless, the fields have developed their own separate terminologies, techniques, and goals.

Because of this, it happens periodically that equivalent notions appear in the two areas, independently of each other, and under different names. In most cases, a relationship between such notions can be derived directly from their definitions. Sometimes, however, this link can be much more subtle. This is indeed the case with the two notions studied in the present paper: lettericity on the one hand, and geometric griddability on the other. At a first reading of the definitions, nothing suggests that there should be a connection between these two notions. However, such a connection does exist, and was first observed by Alecu, Lozin, de Werra, and Zamaraev [6, 7]. They showed that if \mathcal{C} is a geometrically griddable permutation class, then $G_{\mathcal{C}}$, the associated class of inversion graphs, has bounded lettericity. They conjectured that the converse to this statement also holds. We prove this conjecture, yielding the following result.

THEOREM 1.1. *The permutation class \mathcal{C} is geometrically griddable if and only if the corresponding graph class $G_{\mathcal{C}}$ has bounded lettericity.*

*Received by the editors September 29, 2021; accepted for publication (in revised form) May 31, 2022; published electronically November 14, 2022.

<https://doi.org/10.1137/21M1449646>

Funding: The first author's research was supported by EPSRC via a Doctoral Training Partnership grant to the University of Warwick. The third author's research was supported by the French Agency for Research under the ASSK project (ANR-18-CE40-0025). The fifth author's research was supported by the Simons Foundation via award 636113.

[†]Mathematics Institute, University of Warwick, Coventry, CV4 7AL, UK (b.alecu@warwick.ac.uk, V.Lozin@warwick.ac.uk).

[‡]Department of Mathematics, University of Florida, Gainesville, FL 32611 USA (rferguson94@ufl.edu, vatter@ufl.edu).

[§]Université Clermont Auvergne, LIMOS, CNRS, Aubière, 63000, France (mamadou.kante@uca.fr).

[¶]Department of Computer Science, University of Liverpool, Liverpool, L69 3BX, UK (Viktor.Zamaraev@liverpool.ac.uk).

Despite their seemingly unrelated definitions, this result reflects the fact that the concepts of lettericity and geometric griddability capture the same structural data of their respective combinatorial objects: a partition of the elements and a linear ordering of those elements, encoded in words, that interacts nicely with the partition. Indeed, this close connection is evidenced by the way the word encodings imply well-quasi-orderability of the relevant classes (via Higman's lemma [20]; in fact both are *labelled* well-quasi-ordered—see Atminas and Lozin [10] and Brignall and Vatter [13]). Further context for Theorem 1.1 can be found in the article [23, section 2] accompanying Lozin's plenary lecture at the conference *LATA* 2019.

We define the terms involved in sections 2, 3, and 4. Then, in section 5, we develop an alternative approach to geometric grid classes that is necessary for our work. The proof of our main result is contained in section 6, and we conclude in section 7. The unification of lettericity and geometric griddability established in this paper marks the starting point of a promising research direction, which we discuss in the conclusion.

2. Graphs, graph classes, and lettericity. Our graphs are all finite, simple, and undirected. Given a graph G , we denote by $V(G)$ its vertex set and by $E(G)$ its edge set. We write $u \sim v$ to denote that the vertices u and v are adjacent, meaning that $uv \in E(G)$. We denote the complement of G by \bar{G} . We denote the *complete graph* or *clique* on n vertices by K_n , so an *independent set* or *co-clique* on n vertices is denoted by \bar{K}_n . Given graphs G and H on disjoint vertex sets, we denote their *disjoint union* by $G \uplus H$. This is the graph defined by $V(G \uplus H) = V(G) \cup V(H)$ and $E(G \uplus H) = E(G) \cup E(H)$. Given a graph G and a natural number m , we denote by mG the disjoint union of m copies of G (that is, the copies are chosen to have disjoint vertex sets).

An *induced subgraph* of the graph G consists of a subset of the vertices of G together with all edges of G connecting vertices of this subset. This is the only notion of subgraph considered in this paper, and it is a partial order on the set of finite graphs. A *class* of graphs is a set of finite graphs that is closed under isomorphism and closed downward under the induced subgraph ordering (these are also frequently called *hereditary properties*). Thus if \mathcal{C} is a graph class, $G \in \mathcal{C}$, and H is isomorphic to an induced subgraph of G , then $H \in \mathcal{C}$. We introduce here the notions and results required, referring to the encyclopedic text of Brandstädt, Le, and Spinrad [11] for further information and context on graph classes.

Letter graphs were introduced in 2002 by Petkovšek [26], and describe a way to build graphs from words. Let Σ be a finite alphabet, and let $D \subseteq \Sigma^2$ be a set of ordered pairs that we call a *decoder*. Equivalently, we may think of the decoder as a digraph with vertex set Σ . For any word $w = w(1)w(2)\cdots w(n) \in \Sigma^*$, the *letter graph* of w with respect to D is defined to be the graph $\Gamma_D(w)$ defined by

$$\begin{aligned} V(\Gamma_D(w)) &= \{1, 2, \dots, n\} \text{ and} \\ E(\Gamma_D(w)) &= \{ij : i < j \text{ and } (w(i), w(j)) \in D\}. \end{aligned}$$

We frequently say that the vertex $i \in V(\Gamma_D(w))$ is *encoded* by the letter $w(i) \in \Sigma$. If Σ is an alphabet of size k , then the graph $\Gamma_D(w)$ is said to be a *k-letter graph*.

It is easy to see that every graph on n vertices is isomorphic to an n -letter graph—fix a labeling of the vertices of G by $\Sigma = \{1, 2, \dots, n\}$, let $w = 12\cdots n$, and set $D = \{(i, j) : i < j \text{ and } i \sim j \text{ in } G\}$. A triple (Σ, D, w) such that G is isomorphic to $\Gamma_D(w)$ is called a *letter graph realization*, or simply *lettering*, of G . The minimum value of k such that G admits a lettering with $|\Sigma| = k$ is called the *lettericity* of the




FIG. 1. A threshold graph with vertex weights indicated. Two vertices are connected by an edge if the sum of their weights is nonnegative.

graph G . Another fact that is easy to see is that if G is isomorphic to the letter graph $\Gamma_D(w)$, then its complement \bar{G} is isomorphic to $\Gamma_{\Sigma^2 \setminus D}(w)$. From this, we immediately obtain the following.

PROPOSITION 2.1 (see Petkovšek [26, Proposition 2 (iii)]). *The lettericity of a graph and its complement are equal.*

It is apparent from the definition that all of the vertices of $\Gamma_D(w)$ that are encoded by a single letter form either a clique or a co-clique—if the letter is $a \in \Sigma$, then these vertices form a clique if $(a, a) \in D$ and they form a co-clique if $(a, a) \notin D$. This implies that the lettericity of a graph is bounded below by its *co-chromatic number*, which is defined to be the minimum number of cliques and co-cliques that the vertices of the graph can be partitioned into. That said, the perfect matchings have co-chromatic number 2 but unbounded lettericity, as we sketch below.

PROPOSITION 2.2. *The set of perfect matchings $\{mK_2 : m \geq 1\}$ has unbounded lettericity.*

Proof. Suppose to the contrary that the lettericity of every matching mK_2 was bounded by some constant, say r . In such an r -lettering of a matching, every edge has only r^2 possible labelings (or letterings). Therefore, the matching $(mr^2 + 1)K_2$ must contain m edges with the same labels. It follows that the matching mK_2 would therefore actually have lettericity 2. However, it is not difficult to verify that $3K_2$ has lettericity 3. \square

In fact, without much additional effort, one can show that the lettericity of mK_2 is precisely m . This was claimed without proof by Petkovšek [26, section 5.3], and proved by Alecu, Lozin, and de Werra [5, Lemma 3]; the approach used above is from Ferguson and Vatter [16].

We define the lettericity of a class of graphs to be the greatest lettericity of any of its members. Obviously, this quantity is often infinite (for example, whenever a class contains all perfect matchings). The quintessential example of a graph class of finite lettericity is the class of threshold graphs introduced by Chvátal and Hammer, first in a 1973 research report and then in a 1977 paper [14], and the subject of a book by Mahadev and Peled [24]. The original definition of this class (which lends the class its name) is that the graph G is a *threshold graph* if there exists an assignment $v \mapsto c_v$ of real-valued weights to the vertices of G such that $u \sim v$ if and only if $c_u + c_v \geq 0$ (in fact, in the original definition this threshold value is arbitrary, but we lose nothing by taking it to be 0). For example, the graph shown in Figure 1 is a threshold graph, with weights as indicated by its labels.

The set of threshold graphs is a graph class, and it can be shown that it is precisely the class of graphs that do not contain induced subgraphs isomorphic to the matching on two edges $2K_2$, the cycle on four vertices $C_4 = 2\bar{K}_2$, or the path on four vertices P_4 .




FIG. 2. Two additional drawings of the threshold graph from Figure 1.




FIG. 3. The containment order on permutations and the induced subgraph order on their corresponding inversion graphs.

More relevant to our concerns, it can also be shown that every threshold graph can be built from the empty graph by successively adding isolated vertices (adjacent to none of the previous vertices) and dominating vertices (adjacent to all of the previous vertices). Figure 2 illustrates how our example from Figure 1 can be drawn in this manner.

In the first of these drawings, we have retained the weights on the vertices for ease of comparison with our previous drawing of this graph, while in the second we have labelled vertices by their order of addition to the graph (as isolated or dominating vertices). From this perspective, it is clear how to encode threshold graphs as two-letter graphs; using the two-letter alphabet $\Sigma = \{i, d\}$ and the decoder $D = \{(i, d), (d, d)\}$, we encode vertices in their order of addition to the graph by i if they are added as isolated vertices or by d if they are added as dominating vertices. Thus with this alphabet and decoder, our example can be seen to be isomorphic to $\Gamma_D(\text{ididid})$. In fact, it is also isomorphic to $\Gamma_D(\text{ddidid})$ because the first letter does not matter in this case, but that is beside the point.

The discussion above demonstrates that the class of threshold graphs has lettericity 2. In this case, every threshold graph can be encoded using the same decoder, but note that our definition of the lettericity of a class does not require us to use the same decoder for every graph in the class. That said, since there are only finitely many possible decoders for a given finite alphabet, if a class has finite lettericity, then we can also find a single decoder D over a larger but still finite alphabet Σ such that all members of the class can be expressed as $\Gamma_D(w)$ for some word $w \in \Sigma^*$.

3. Permutations, permutation classes, and grids. We think of permutations in one-line notation, so a permutation of length n is simply an ordering of the set $\{1, 2, \dots, n\}$. The permutation π contains the permutation σ of length k if it has a subsequence of length k that is order isomorphic to σ , by which we mean that the subsequence has the same pairwise comparisons as σ . For example, $\pi = 372694185$ contains $\sigma = 32514$, as witnessed by its subsequence 32918, but π avoids 54321 because it has no decreasing subsequence of length five. This containment is depicted graphically on the left of Figure 3. As in this figure, we often identify a permutation π with its *plot*: the set of points $\{(i, \pi(i))\}$ in the plane.




FIG. 4. A $\begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$ -gridding of the permutation 524361 (recall that permutations are read from left to right). Here the column divisions are given by $x_1 = 1$, $x_2 = 4$, and $x_3 = 7$, while the row divisions are $y_1 = 1$, $y_2 = 5$, and $y_3 = 7$.

A class of permutations is a set of permutations that is closed downward under the containment order defined above. Thus if \mathcal{C} is a permutation class, $\pi \in \mathcal{C}$, and σ is contained in π , then $\sigma \in \mathcal{C}$. We introduce here the notions and results required, referring to the survey by Vatter [29] for further information and context on permutation classes.

The *inversion graph* of the permutation $\pi = \pi(1) \cdots \pi(n)$ is the graph G_π on the vertices $\{1, \dots, n\}$ in which $i \sim j$ if and only if $\pi(i)$ and $\pi(j)$ form an *inversion*, meaning that either $i < j$ and $\pi(i) > \pi(j)$ or $i > j$ and $\pi(i) < \pi(j)$, or equivalently, if and only if

$$(i - j)(\pi(i) - \pi(j)) < 0.$$

These graphs are commonly called permutation graphs in the literature, but we use the term inversion graph here to emphasize *how* they arise from permutations. As shown on the right of Figure 3, to obtain the inversion graph of a permutation from its plot, we simply add all edges between pairs of entries in which one lies to the northwest of the other.

If the permutation σ is contained in the permutation π , then it is clear that G_σ is an induced subgraph of G_π . The converse, however, does not generally hold; for one example, $G_{2413} \cong G_{3142} \cong P_4$, but neither 2413 nor 3142 is contained in the other. That said, only the increasing and decreasing permutations correspond to cliques and co-cliques, so every clique (resp., co-clique) in G_π corresponds to a decreasing (resp., increasing) subsequence of π .

The monotone grid class of a $0/\pm 1$ matrix consists of all those permutations whose plots can be sliced into rectangles in such a way that each of the resulting rectangular sections contains a monotone subsequence, as specified by the matrix. More precisely, suppose that M is a $0/\pm 1$ matrix of size $t \times u$. In order for the entries of our matrices to align with the plots of permutations, we index matrices in cartesian coordinates, so this means for us that M consists of t columns and u rows. An M -gridding of the permutation π of length n is a choice of *column divisions* $1 = x_1 \leq \cdots \leq x_{t+1} \leq n + 1$ and *row divisions* $1 = y_1 \leq \cdots \leq y_{u+1} \leq n + 1$ such that for all i and j , the subsequence of π with indices in $[x_i, x_{i+1})$ and values in $[y_j, y_{j+1})$ is increasing if $M(i, j) = 1$, decreasing if $M(i, j) = -1$, or empty if $M(i, j) = 0$. The *monotone grid class* of M is the set of all permutations that possess an M -gridding, and is denoted by $\text{Grid}(M)$.

Figure 4 depicts a member of the grid class of $\begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$. This particular grid class admits a simpler definition: it consists of those permutations whose entries can be partitioned into an increasing subsequence and a decreasing subsequence. This is the definition under which these permutations were first studied, by Stankova [27], who called them *skew-merged*. Stankova proved that a permutation is skew-merged if and only if it contains neither 2143 nor 3412.

The permutation class \mathcal{C} is said to be *monotonically griddable* if $\mathcal{C} \subseteq \text{Grid}(M)$ for some $0/\pm 1$ matrix M . Monotone grid classes were first considered in a 2003 paper by Murphy and Vatter [25], although they called them “profile classes.” In their characterization of permutation classes of polynomial growth (see also [21]), Huczynska and Vatter [22] gave monotone grid classes their name and characterized the monotonically griddable permutation classes. Stankova’s characterization of the skew-merged permutations in fact constitutes the base case of the inductive proof of the following result in [22]. For a generalization of this result we refer to Vatter [28, Theorem 3.1].

THEOREM 3.1 (see Huczynska and Vatter [22, Theorem 2.5]). *The permutation class \mathcal{C} is monotonically griddable if and only if there is a natural number m such that \mathcal{C} contains neither $2143 \cdots (2m)(2m-1)$ nor $(2m-1)(2m) \cdots 3412$.*

The graphical analogues of the skew-merged permutations are the *split graphs*: the graphs whose vertices can be partitioned into a clique and a co-clique. This graph class was first studied by Földes and Hammer [17] in 1977, who showed that a graph is split if and only if it contains none of the graphs $2K_2$, $\overline{2K_2} = C_4$, or C_5 as induced subgraphs. We note that inversion graphs never contain induced cycles of length five or more, so a class of inversion graphs is split if and only if it contains neither $2K_2$ nor $\overline{2K_2} = C_4$.

More generally, we observe that

$$G_{2143 \cdots (2m)(2m-1)} = mK_2 \text{ and} \\ G_{(2m-1)(2m) \cdots 3412} = \overline{mK_2}.$$

Just as only the increasing and decreasing permutations correspond to cliques and co-cliques, only the permutation $2143 \cdots (2m)(2m-1)$ corresponds to the matching mK_2 , and only its reverse $(2m-1)(2m) \cdots 3412$ corresponds to $\overline{mK_2}$. This gives us the following consequence of Theorem 3.1, which is somewhat analogous to the main result of our paper, Theorem 1.1.

COROLLARY 3.2. *The permutation class \mathcal{C} is monotonically griddable if and only if there is a natural number m such that the corresponding graph class $G_{\mathcal{C}}$ contains neither mK_2 nor $\overline{mK_2}$.*

Proposition 2.2 shows that the lettericity of the perfect matchings mK_2 is unbounded, and thus by Proposition 2.1, the lettericity of their complements $\overline{mK_2}$ is also unbounded. Thus we have the following consequence of Corollary 3.2.

COROLLARY 3.3. *If the graph class $G_{\mathcal{C}}$ has bounded lettericity, then the permutation class \mathcal{C} is monotonically griddable.*

The converse of Corollary 3.3 does not hold, as there are monotonically griddable permutation classes for which the corresponding graph classes have unbounded lettericity (in fact, we’ve just presented an example: the skew-merged permutations are monotonically griddable, but the corresponding graph class, the split graphs, has unbounded lettericity). We close this section by remarking that the graphical analogues of general monotonically griddable permutation classes have only recently attracted attention, in the work of Atminas [9].

4. Geometric grid classes. From our previous discussion, it follows that Chvátal and Hammer’s threshold graphs are a subclass of Földes and Hammer’s split graphs. This can be seen by either recalling the iterative construction of threshold graphs




FIG. 5. On the left, the standard figure of the matrix $\begin{pmatrix} -1 & 1 & 1 \\ 0 & -1 & -1 \end{pmatrix}$. On the right, a drawing of the permutation 6437251 on this figure.

(by adding isolated or dominating vertices in sequence), or from the forbidden subgraph characterizations of the two classes (graphs that avoid $2K_2$, $\overline{2K_2} = C_4$, and P_4 will necessarily avoid $2K_2$, $\overline{2K_2} = C_4$, and C_5). Moreover, this containment of graph classes is strict because P_4 is a split graph, but is not a threshold graph. The relationship between the class of split graphs and the class of threshold graphs is a graphical analogue of the relationship between monotone grid classes and geometric grid classes. Having defined the former in the previous section, we define the latter here.

Geometric grid classes may be defined in terms of drawings of permutations in the plane \mathbb{R}^2 , which we think of as divided into *cells* whose corners are given by the integer lattice \mathbb{Z}^2 . The *standard figure* of a $0/\pm 1$ matrix M is the point set in \mathbb{R}^2 consisting of the following:

- the increasing open line segment from $(i-1, j-1)$ to (i, j) if $M(i, j) = 1$ or
- the decreasing open line segment from $(i-1, j)$ to $(i, j-1)$ if $M(i, j) = -1$.

An example is shown on the left of Figure 5; we remind the reader again that we index matrices in cartesian coordinates. We further extend this indexing to the cells of the standard figure of M , so that the cell $(2, 1)$ of this standard figure corresponds to the entry $M(2, 1)$ and lies in the second column and first row.

Any subset of n points of the standard figure of M in which no two points lie on a common horizontal or vertical line (such point sets are often called *generic*) corresponds to a permutation by labeling the points from 1 to n from bottom to top and then reading their labels from left to right. We say that the permutation π of length n can be *drawn* on the standard figure of M if it corresponds to such a subset. An example is shown on the right of Figure 5.

The *geometric grid class* of M is the class of all permutations that can be drawn on the standard figure of M in this manner. We denote this class by $\text{Geom}(M)$. The permutation class \mathcal{C} is called *geometrically griddable* if it is contained in a geometric grid class, that is, if $\mathcal{C} \subseteq \text{Geom}(M)$ for some (finite) $0/\pm 1$ matrix M .

It is clear that

$$\text{Geom}(M) \subseteq \text{Grid}(M)$$

for all $0/\pm 1$ matrices M . In some cases (characterized by whether a certain graph determined by the matrix M is a forest—see Vatter [29, Proposition 12.4.12]), this is an equality, but not always. In this work, we use only the fact that geometrically griddable classes are necessarily monotonically griddable. Thus, in particular, Theorem 3.1 implies that every geometrically griddable permutation class must avoid $2143 \cdots (2m)(2m-1)$ and $(2m-1)(2m) \cdots 3412$ for some natural number m . We note that there is, to-date, no characterization such as Theorem 3.1 for geometrically griddable classes themselves.




FIG. 6. The permutation 3142 cannot be drawn on an X. Once we place the 3, 1, and 4 on the X, there is no place for the 2 to lie simultaneously above the 1 and to the right of the 4.

One of the first examples of a geometric grid class to be considered (in fact, before the machinery of geometric grid classes had been established) is the class of permutations that can be drawn on a circle [31]. This class is the same as the geometric grid class of the matrix $\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}$.

More relevant to our previous discussion is the class of permutations that can be drawn on an X, or in other words, the geometric grid class of the matrix $\begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$. As with the permutations that can be drawn on a circle above, this geometric grid class is a strict subset of the corresponding monotone grid class:

$$\text{Geom} \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix} \subsetneq \text{Grid} \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}.$$

An example of why this inclusion is strict is shown in Figure 6. In fact, this geometric grid class consists of those skew-merged permutations that are also separable, meaning that they avoid 2413 and 3142.

If G_π is a threshold graph, then π can be drawn on an X. Indeed, in this case we can go a step further: if G is a threshold graph, then $G \cong G_\pi$ for a permutation π that can be drawn on an X. But we do not need to consider all permutations that can be drawn on an X to get the class of threshold graphs. In [18],¹ Golumbic showed that the threshold graphs can be obtained as inversion graphs of members of the class

$$\text{Geom} \begin{pmatrix} -1 \\ 1 \end{pmatrix},$$

which is the class of permutations that can be drawn on a \succ . Golumbic’s characterization is in terms of the shuffle product, where for words u and v , the set $u \sqcup v$ consists of all words formed by interleaving the letters of u and v in all possible ways. For example,

$$12 \sqcup 543 = \{12543, 15243, 15423, 15432, 51243, 51423, 51432, 54123, 54132, 54312\}.$$

His result then states that the threshold graphs are precisely the inversion graphs of the permutations contained in the shuffle product

$$\left(12 \cdots m\right) \sqcup \left(n(n-1) \cdots (m+1)\right)$$

for some nonnegative integers n and m , which is equivalent to being drawn on a \succ . By symmetry, one can also obtain the threshold graphs as the inversion graphs of permutations that can be drawn on a \vee , a \prec , or a \wedge ; see Figure 7.

¹This theorem can also be found in Golumbic’s textbook *Algorithmic Graph Theory and Perfect Graphs* [19].




FIG. 7. The inversion graphs of all five of these permutations are isomorphic to the threshold graph depicted in Figures 1 and 2.

This example illustrates one difficulty in the proof of Theorem 1.1. Letting \mathcal{X} denote the class of permutations that can be drawn on an \times and \mathcal{V} denote the class of permutations that can be drawn on a \vee , we see that \mathcal{V} is a proper subclass of \mathcal{X} , but that these two classes correspond to the same class of inversion graphs (the threshold graphs). In symbols,

$$\mathcal{V} \subsetneq \mathcal{X} \quad \text{but} \quad G_{\mathcal{V}} = G_{\mathcal{X}}.$$

In particular, this example shows that we may need more cells in our geometric grid classes than letters in the corresponding class of letter graphs—the threshold graphs have lettericity 2 while the permutations that can be drawn on an \times require four nonempty cells. In our proof, we make no effort to conserve cells, and consequently this explosion will be quite dramatic.

5. Encoding geometric grid classes. In the previous section we defined geometric grid classes geometrically, as classes of permutations that can be drawn on the standard figures of their defining matrices. While that description suffices as a definition, it is not very convenient for actually proving theorems about geometric grid classes. Here we give an order-theoretic viewpoint of geometric gridability, and show how one direction of Theorem 1.1 follows readily. We note that this approach essentially follows Vatter and Waton [30], and thus actually predates the definition of geometric grid classes.

The matrix M of size $t \times u$ is a *partial multiplication matrix* if there are *column* and *row signs*

$$c_1, \dots, c_t, r_1, \dots, r_u \in \{1, -1\}$$

such that every nonzero entry $M(k, \ell)$ is equal to $c_k r_\ell$. Partial multiplication matrices are necessarily $0/\pm 1$ matrices. However, not every $0/\pm 1$ matrix is a partial multiplication matrix, as can be seen by considering the matrix $\begin{pmatrix} 1 & -1 \\ 1 & -1 \end{pmatrix}$. It is nevertheless true that every geometric grid class can be viewed as the geometric grid class of a partial multiplication matrix. To see this, let M be an arbitrary $0/\pm 1$ matrix, and form a new matrix $M^{\times 2}$ by replacing each entry of M by an appropriate 2×2 block, according to the substitution rules

$$0 \leftarrow \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, \quad 1 \leftarrow \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \text{and} \quad -1 \leftarrow \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}.$$

An example of this process is shown in Figure 8. As indicated by that figure, the standard figure of $M^{\times 2}$ is simply a scaled copy of the standard figure of M , and thus

$$\text{Geom}(M^{\times 2}) = \text{Geom}(M).$$

By definition, $M^{\times 2}$ is comprised of 2×2 blocks, each obtained by replacing two or four entries of the matrix $\begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$ by 0. As the matrix $\begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}$ is a partial multiplication matrix, it follows that $M^{\times 2}$ is a partial multiplication matrix. In fact, remembering




FIG. 8. The standard figures of the matrices $M = \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$ and $M^{\times 2} = \begin{pmatrix} 0 & 1 & -1 & 0 \\ 1 & 0 & 0 & -1 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{pmatrix}$.

that we index matrices in cartesian coordinates, we can use the same column and row signs, $c_k = (-1)^k$ and $r_\ell = (-1)^\ell$, for every matrix of the form $M^{\times 2}$ (although we do not assume this). We may therefore restrict our attention to partial multiplication matrices, as recorded below.

PROPOSITION 5.1 (see Albert et al. [1, Proposition 4.2]). *Every geometric grid class is the geometric grid class of a partial multiplication matrix.*

With this result in hand, suppose that M is a partial multiplication matrix with column and row signs (c_k) and (r_ℓ) , respectively, and that $\pi \in \text{Grid}(M)$. We fix a particular M -gridding of π and denote this *gridded permutation* (the permutation π together with an M -gridding) by π^\sharp . While π^\sharp is technically a permutation together with a choice of column and row divisions, in practice we often think of it as simply a permutation together with an assignment of its entries into cells. We now use the column and row signs to orient the entries of π^\sharp in two different manners, one geometric and the other order-theoretic. We begin with the geometric interpretation.

We interpret the column sign c_k as specifying whether the entries of π^\sharp in column k are oriented from left to right (if $c_k = 1$) or from right to left (if $c_k = -1$). We similarly interpret the row sign r_ℓ as specifying whether the entries of π^\sharp in row ℓ are oriented from bottom to top (if $r_\ell = 1$) or from top to bottom (if $r_\ell = -1$). We depict these orientations with arrows to the left of and below the corresponding rows and columns of the standard figure of the matrix M , as shown in Figure 9.

These orientations are consistent with the contents of each cell in the sense that $M(k, \ell) = c_k r_\ell$ if it is nonzero, because M is a partial multiplication matrix. The column and row signs allow us to further define a *base point* in each nonempty cell of the standard figure, which we take to be the point from which all the arrows of that cell originate, and we further depict the line segment in each cell (upon which the points lie) as a ray pointing away from the base point.

This is the geometric orientation of the entries of π^\sharp . In the order-theoretic viewpoint, the column and row signs determine a linear orders on the entries of π^\sharp in each column and row.

For each index k , we define a linear *column order* \leq_k^{col} on all of the entries of π that lie in the k th column of π^\sharp . The orientation of \leq_k^{col} is dictated by c_k , in the sense that we order these entries from left to right if $c_k = 1$ and right to left if $c_k = -1$. Thus if both $\pi(i)$ and $\pi(j)$ lie in column k of π^\sharp , then

$$i \leq_k^{\text{col}} j \iff \begin{cases} i \leq j & \text{if } c_k = 1 \text{ or} \\ i \geq j & \text{if } c_k = -1. \end{cases}$$

Analogously, for each index ℓ , we define a linear *row order* \leq_ℓ^{row} on all of the entries




FIG. 9. Two drawings of the geometric gridding of $\pi(1)\pi(2)\cdots\pi(7) = 6437251$ from Figure 5, with arrows to indicate column and row signs. In the drawing on the left, the points are labelled by their vertical positions (that is, the values $\pi(i)$ of the corresponding entries of the permutation), while in the drawing on the right, they are labelled by their horizontal positions (that is, the indices of the corresponding entries of the permutation).




FIG. 10. Hasse diagrams showing the local orders for the gridded permutation shown in Figure 9, together with the Hasse diagram of the transitive closure of their union on the far right.

of π that lie in the ℓ th row of $\pi^\#$. This order is oriented from bottom to top if $r_\ell = 1$ and from top to bottom if $r_\ell = -1$, and thus if both $\pi(i)$ and $\pi(j)$ lie in row ℓ of $\pi^\#$, then

$$i \trianglelefteq_\ell^{\text{row}} j \iff \begin{cases} \pi(i) \leq \pi(j) & \text{if } r_\ell = 1 \text{ or} \\ \pi(i) \geq \pi(j) & \text{if } r_\ell = -1. \end{cases}$$

We refer to the column and row orders together as the *local orders* of $\pi^\#$.

We pause to briefly consider these local orders for the gridding of our example in Figure 9, with column and row signs indicated by the picture. Note that we typically discuss entries of permutations by *value*, as on the left of Figure 9, but we have defined the local orders by *index*. This is why we have also supplied the drawing on the right of Figure 9, which is identical to the one on the left, except that points are labelled by the index of their corresponding entry in the permutation.

As M has three columns and two rows, there are five local orders:

$$\begin{aligned} 1 &\trianglelefteq_1^{\text{col}} 2, \\ 4 &\trianglelefteq_2^{\text{col}} 3, \\ 7 &\trianglelefteq_3^{\text{col}} 6 \trianglelefteq_3^{\text{col}} 5, \\ 7 &\trianglelefteq_1^{\text{row}} 5 \trianglelefteq_1^{\text{row}} 3, \\ 4 &\trianglelefteq_2^{\text{row}} 1 \trianglelefteq_2^{\text{row}} 6 \trianglelefteq_2^{\text{row}} 2. \end{aligned}$$

It is not an accident that the transitive closure of the union of the local orders

happens to form a poset. For the gridded permutation from Figure 9, the Hasse diagram of this poset is shown in Figure 10, along with the Hasse diagrams of the local orders. To see that this is the case in general, we note that each local order actually orders the elements of π in increasing order of their distance to the base point of their respective cell; this is a feature of geometric gridings by partial multiplication matrices.

To make this precise, let us say that a *geometric realization* of the gridded permutation $\pi^\#$ is a drawing of the permutation π on the standard figure of M in such a way that every entry of π lies in the cell specified by $\pi^\#$. Suppose that $\pi^\#$ has a geometric realization, and fix one. For each index i , let $d(i)$ denote the (euclidean) distance between the point of this geometric realization that corresponds to $\pi(i)$ and the base point of its cell (thus $d(i)$ lies between 0 and $\sqrt{2}$ for every index i). It is possible that points in independent cells (that is, cells that do not share a row or a column) have the same distance to their base points. If necessary, one can modify the realization to avoid this situation by perturbing the points in question—since the plot of the permutation is a generic set of points, one can easily check that this is always possible. We then have, if the entries $\pi(i)$ and $\pi(j)$ both lie in column k of $\pi^\#$, that

$$i \trianglelefteq_k^{\text{col}} j \iff d(i) \leq d(j),$$

while, analogously, if $\pi(i)$ and $\pi(j)$ both lie in row ℓ of $\pi^\#$, then

$$i \trianglelefteq_\ell^{\text{row}} j \iff d(i) \leq d(j).$$

Recall that a set of linear orders is *consistent* if the transitive closure of their union forms a poset, which happens if and only if their union does not contain a cycle. By our above observation, if $\pi^\#$ has a geometric realization, then the local orders of this gridded permutation must be consistent, because each of them agrees with the distance-to-base-point ordering, which is a linear ordering.

Conversely, suppose that the local orders of the gridding $\pi^\#$ are consistent, and thus their union forms a poset, say \trianglelefteq . Fix a linear extension of this poset, which amounts to choosing a permutation ψ of the indices of π such that

$$i \trianglelefteq j \implies \psi(i) \leq \psi(j).$$

Also, choose an arbitrary set of distances

$$0 < d_1 < d_2 < \cdots < d_n < \sqrt{2}.$$

We may then obtain a geometric realization of this gridding of π by placing, for each index i , a point corresponding to $\pi(i)$ at distance $d_{\psi(i)}$ from the base point in the appropriate cell of the standard figure of M (the cell in which $\pi(i)$ lies in $\pi^\#$).

Returning to our example from Figure 9, we have already drawn the Hasse diagram of the transitive closure of the union of the local orders. To obtain a linear extension of this poset, we may choose to place the points in the distance-to-base-point order 4, 1, 7, 6, 5, 2, 3. This corresponds to choosing

$$\psi = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 2 & 6 & 7 & 1 & 5 & 4 & 3 \end{pmatrix},$$

because then

$$\psi^{-1} = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 4 & 1 & 7 & 6 & 5 & 2 & 3 \end{pmatrix}.$$

(In general, $\psi(i)$ tells us where i lies in the linear extension, so reading ψ^{-1} from left to right gives us the linear extension itself.) We have not specified how to choose the distances d_i , but it is not hard to see that this choice is immaterial to the gridded permutation we are trying to produce.

In general, we have established the following result.

PROPOSITION 5.2. *Let M be a partial multiplication matrix with fixed column and row signs, and let $\pi^\#$ be an M -gridding of a permutation $\pi \in \text{Grid}(M)$. Then $\pi^\#$ has a geometric realization if and only if its local orders are consistent.*

This approach underlies the encoding of geometric grid classes by words, first presented in [1], and described briefly here. Define the *cell alphabet* of the matrix M to be the set

$$\Sigma = \{a_{k\ell} : M(k, \ell) \neq 0\},$$

so that the letters $a_{k\ell}$ correspond to the nonzero entries of M . We now describe a mapping $\varphi^\#$ from the language of words Σ^* to the set of geometrically gridded members of $\text{Geom}(M)$. For each word $w = w(1)w(2)\cdots w(n) \in \Sigma^*$, we construct a permutation $\pi \in \text{Geom}(M)$ as follows: choose a set $0 < d_1 < d_2 < \cdots < d_n < \sqrt{2}$ of distances, and then for $1 \leq i \leq n$ if $w(i) = a_{k\ell}$, place a point on the diagonal of the cell (k, ℓ) of the standard figure of M at (euclidean) distance d_i from the base point of that cell.

As in our discussion above, the choice of distances is also immaterial here, and so the mapping $\varphi^\#$ is well defined. Moreover, our previous discussion shows that $\varphi^\#$ is onto, by the following argument. Given any gridding $\pi^\#$ that has a geometric realization, Proposition 5.2 shows that the local orders must be consistent. This means that the transitive closure of their union will have a linear extension, which we can take to define the permutation ψ as before. We may then define the word w by the rule that

$$w(\psi(i)) = a_{k\ell} \iff \pi(i) \text{ lies in cell } (k, \ell) \text{ in the gridding } \pi^\#,$$

and we see that $\varphi^\#(w) = \pi^\#$.

By simply removing the gridding from the gridded permutation $\varphi^\#(w)$, we obtain a permutation $\varphi(w) \in \text{Geom}(M)$. It is easy to see that the mapping φ is *order-preserving*, in the sense that if u is a not-necessarily-consecutive subword of w , then $\varphi(u) \leq \varphi(w)$ as permutations. This observation and Higman's lemma [20] thereby imply that all geometrically griddable classes are well-quasi-ordered, a result of Albert, Atkinson, Bouvel, Ruškuc, and Vatter [1, Theorem 6.1].

Our discussion shows that every different linear extension ψ gives rise to a different word w that encodes the same gridded permutation $\pi^\#$. Indeed, part of the analysis of [1] is the observation that changing the relative order in the linear extension ψ of elements of independent cells is reflected, on the word level, by the set of encodings of a given gridded permutation $\pi^\#$ being described as a partially commutative monoid (also known as a trace monoid, such objects have been extensively studied; see Diekert's text *Combinatorics on Traces* [15], for example). That said, this further formalization is not necessary for our arguments.

Another detail that does not concern us in the present work is that members of $\text{Geom}(M)$ may have many different M -griddings with geometric realizations. For example, if we were to ask only that our words encode the *permutation* shown in Figure 9, then we would have many more, as it has several alternative griddings (for example, the point corresponding to $\pi(2)$ could shift to any of two other cells, or if that

point stayed put, then the point corresponding to $\pi(3)$ could shift cells itself). Albert, Atkinson, Bouvel, Ruškuc, and Vatter [1, Theorem 8.1] show how to identify—in a regular manner—a distinguished geometric gridding for every member of $\text{Geom}(M)$, thus establishing that every geometrically griddable class is in bijection with a regular language and thus has a rational generating function.

Our task here is different. Instead of trying to choose one geometric gridding for each member of a geometric grid class, we seek only to show that certain gridings *are* geometric. Returning to this goal, the easier direction of Theorem 1.1—which states that the graph class associated to a geometrically griddable permutation class has bounded lettericity—has already been proved in [7, 6], but we include another proof here, both for completeness and because it further illustrates some ideas used in the proof of the converse. Note that in the proof of this result, we are able to use the same word to encode both the permutation $\pi \in \text{Geom}(M)$ and the letter graph satisfying $\Gamma_D(w) \cong G_\pi$.

THEOREM 5.3 (see Alecu et al. [7, 6]). *If the permutation class \mathcal{C} is geometrically griddable, then the corresponding graph class $G_{\mathcal{C}}$ has bounded lettericity.*

Proof. Suppose that $\mathcal{C} \subseteq \text{Geom}(M)$, where M is a partial multiplication matrix with column signs (c_k) and row signs (r_ℓ) , and denote the cell alphabet of M by $\Sigma = \{a_{k\ell} : M(k, \ell) \neq 0\}$. Let $\pi \in \mathcal{C}$ be arbitrary, and suppose that $\pi = \varphi(w)$, where $w = w(1) \cdots w(n) \in \Sigma^*$. We prove the result by showing how to construct a decoder $D \subseteq \Sigma^2$ such that $\Gamma_D(w) \cong G_\pi$.

First, if we have $M(k, \ell) = 1$, then the entries in cell (k, ℓ) are increasing, and so we set $(a_{k\ell}, a_{k\ell}) \notin D$. Otherwise if $M(k, \ell) = -1$, then the entries in cell (k, ℓ) are decreasing, and we let $(a_{k\ell}, a_{k\ell}) \in D$.

Next, if the cells (k, ℓ) and (k', ℓ') share neither a row nor a column (that is, if $k \neq k'$ and $\ell \neq \ell'$), then either every pair of entries with one from each of these two cells forms an inversion, in which case we set $(a_{k\ell}, a_{k'\ell'}), (a_{k'\ell'}, a_{k\ell}) \in D$, or no such pair forms an inversion, in which case we set $(a_{k\ell}, a_{k'\ell'}), (a_{k'\ell'}, a_{k\ell}) \notin D$.

Finally, we must consider cells that share either a row or a column but not both. As the case where the cells share a row is symmetric, we consider entries lying in the cells (k, ℓ) and (k, ℓ') , which share a column but lie in different rows. We may suppose without loss of generality that $\ell < \ell'$, so the entries of the cell (k, ℓ) lie below those of the cell (k, ℓ') . Thus the only inversions between a pair of such entries consist of an entry in the cell (k, ℓ') lying to the left of an entry from the cell (k, ℓ) . If $c_k = 1$, then the entries of this column are encoded from left to right, so we set $(a_{k\ell'}, a_{k\ell}) \in D$ and $(a_{k\ell}, a_{k\ell'}) \notin D$. Otherwise, if $c_k = -1$, then these entries are encoded from right to left, so we set $(a_{k\ell}, a_{k\ell'}) \in D$ and $(a_{k\ell'}, a_{k\ell}) \notin D$. \square

6. Proof of the main result. In this section we prove the other direction of Theorem 1.1. The bulk of this proof is conducted on the level of an individual permutation, and at the heart of our approach to individual permutations is the following fact, formalized in Proposition 6.7: if a permutation π has a monotone gridding by a matrix M , then π is geometrically griddable by a matrix M^* whose size only depends on the size of M and the lettericity of its inversion graph G_π .

The proof of this fact is somewhat technical, but the concept behind it is quite transparent. A large amount of structural information about the permutation π is available to us. This information comes from two sources: the monotone gridding by M on the one hand, and the letter graph representation of G_π on the other. The problem is that it is not a priori clear how to interpret this information simultaneously.




FIG. 11. In all four cases, the entry $\pi(x)$ separates $\pi(i)$ from $\pi(j)$.

There is, however, a straightforward solution: we take a “common refinement” of the available information, and use the word that encodes G_π as a letter graph to define a linear extension ψ on the entries of π as in the previous section. In practice, the way this is achieved is by first refining the alphabet and the decoder of the letter graph to ensure that each letter is used in only one cell. We then refine the gridding by adding vertical and horizontal lines surrounding the points encoded by each letter. The bulk of the work then consists of showing that this refinement behaves as we would like it to.

We begin our proof by fixing a bijection between the inversion graph of π and an isomorphic letter graph, and using this to define a linear extension ψ on the entries of π . Our initial results, Propositions 6.1 and 6.2 in particular, establish that ψ behaves as we would like it to “locally.” Later results then broaden the scope until Proposition 6.7 concludes our analysis of a single permutation, at which point the proof of our main result is quickly completed in section 6.5.

Suppose that π is a permutation of length n and that

$$G_\pi \cong \Gamma_D(w)$$

for some finite alphabet Σ , some decoder $D \subseteq \Sigma^2$, and some word $w \in \Sigma^n$. We consider a particular isomorphism between G_π and $\Gamma_D(w)$, which we denote by

$$\psi : V(G_\pi) \rightarrow V(\Gamma_D(w)).$$

One may notice that we have given this isomorphism the same name we gave the linear extension in section 5. This is intentional, as the ultimate result of this section states that they play the same role.² We say that the entry $\pi(i)$ *corresponds* to the vertex $\psi(i)$ of $\Gamma_D(w)$ and is *encoded* by the letter $w(\psi(i))$, which we denote symbolically by writing

$$\pi(i) \longleftrightarrow \text{vertex } \psi(i) \text{ in } \Gamma_D(w) \longleftrightarrow w(\psi(i)).$$

As shown in Figure 11, we say that the entries $\pi(i)$ and $\pi(j)$ of the permutation π are *separated* if π contains a third entry lying between them either horizontally or vertically, but not both, and we say that this third entry separates $\pi(i)$ from $\pi(j)$. The analogous concept in the graph context is that two vertices are *distinguished* if there is a third vertex that is adjacent to one but not both of them. It is easy to see that the entries $\pi(i)$ and $\pi(j)$ are separated in π if and only if the vertices i and j are distinguished in G_π .

The connection between separation and lettering is given below.

²This is also the same role as played by the index correspondence used in Albert, Ruškuc, and Vatter [2, section 3].




FIG. 12. The situation to which Proposition 6.2 applies.

PROPOSITION 6.1. *Suppose that $\psi : V(G_\pi) \rightarrow V(\Gamma_D(w))$ is an isomorphism. If $\pi(i_1)$ and $\pi(i_2)$ are encoded by the same letter and are separated by $\pi(x)$, then $\psi(i_1) < \psi(x) < \psi(i_2)$ or the reverse.*

Proof. Since $\pi(x)$ separates $\pi(i_1)$ and $\pi(i_2)$, the vertex $\psi(x)$ distinguishes the vertices $\psi(i_1)$ and $\psi(i_2)$. We assume that $\psi(i_1) < \psi(i_2)$; if instead the reverse is true, then the result follows by a symmetrical argument. If $\psi(x) < \psi(i_1) < \psi(i_2)$ or $\psi(i_1) < \psi(i_2) < \psi(x)$, then the vertex $\psi(x)$ would have to be adjacent to either both vertices $\psi(i_1)$ and $\psi(i_2)$ or neither, because they are encoded by the same letter. \square

6.1. An initial gridding. We now suppose that $\pi \in \text{Grid}(M)$ for some $0/\pm 1$ matrix M of size $t \times u$, and we denote by $\pi^\#$ the permutation π together with a fixed M -gridding. Note that $\pi^\#$ is only assumed to be a monotone gridding (if it is in fact geometric, then we are already done).

The following result, which follows routinely from Proposition 6.1, applies to the arrangement shown in Figure 12. This result has a symmetric version that applies when the four entries alternate when read by value (in the version stated below, the entries alternate when read by position).

PROPOSITION 6.2. *Suppose that $\psi : V(G_\pi) \rightarrow V(\Gamma_D(w))$ is an isomorphism, that $\pi^\#$ is a monotonically gridded permutation, and that we have the following:*

- $i_1 < i_2 < i_3 < i_4$,
- $\pi(i_1)$ and $\pi(i_3)$ are encoded by the same letter and lie together in one cell of $\pi^\#$, and
- $\pi(i_2)$ and $\pi(i_4)$ are encoded by the same letter and lie together in a different cell of $\pi^\#$.

Then $\psi(i_1) < \psi(i_2) < \psi(i_3) < \psi(i_4)$ or the reverse.

Proof. Because $i_1 < i_2 < i_3$ and $\pi(i_2)$ lies in a different cell from $\pi(i_1)$ and $\pi(i_3)$, it separates these entries. Proposition 6.1 then implies that either $\psi(i_1) < \psi(i_2) < \psi(i_3)$ or the reverse; as in our previous proof, if the reverse holds, then the result follows by a symmetrical argument.

Similarly, because $i_2 < i_3 < i_4$ and $\pi(i_3)$ lies in a different cell from $\pi(i_2)$ and $\pi(i_4)$, $\pi(i_3)$ separates these entries. Therefore, Proposition 6.1 implies that $\psi(i_2) < \psi(i_3) < \psi(i_4)$, because the reverse would be a contradiction to the fact that $\psi(i_2) < \psi(i_3)$. It follows that $\psi(i_1) < \psi(i_2) < \psi(i_3) < \psi(i_4)$, as claimed. \square

6.2. Monotone intervals. There may be many choices for ψ . Indeed, if we consider the identity $\pi = 12 \cdots n$, the constant word $w = \mathbf{aa} \cdots \mathbf{a}$, and an empty decoder $D = \emptyset$, then every bijection $\psi : [n] \rightarrow [n]$ is an isomorphism between $G_{12 \cdots n} \cong \bar{K}_n$ and $\Gamma_D(w) \cong \bar{K}_n$. Thus we need to impose some conditions on π in order to get better control of the bijection ψ .

An *increasing interval* in a permutation is a sequence of consecutive entries such




FIG. 13. An example of the situation in Proposition 6.4.

that each of the entries is 1 more than the previous entry, while a *decreasing interval* is, analogously, a sequence of consecutive entries such that each is 1 less than the previous entry. A *monotone interval* is then either an increasing or a decreasing interval. The empty set and single entries by themselves are considered to be *trivial* monotone intervals, while larger monotone intervals are *nontrivial*. If the permutation π has a nontrivial monotone interval, then we can *contract* this monotone interval into a single entry to obtain a permutation σ . Conversely, we can *inflate* σ (replacing the appropriate entry by a monotone interval) to obtain π .

Suppose that $\pi \in \text{Grid}(M)$ were to have one or more nontrivial monotone intervals, and fix a particular M -gridding π^\sharp of it. We may contract its monotone intervals to obtain a permutation $\sigma \in \text{Grid}(M)$ with corresponding M -gridding σ^\sharp (arising from π^\sharp). If σ^\sharp has a geometric realization, then it follows that π^\sharp has a geometric realization, so $\pi \in \text{Geom}(M)$ whenever $\sigma \in \text{Geom}(M)$. We record this fact below for later reference.

OBSERVATION 6.3. *Suppose that the M -gridded permutation σ^\sharp can be obtained by contracting monotone intervals of the M -gridded permutation π^\sharp . If σ^\sharp has a geometric realization, then π^\sharp also has a geometric realization.*

Because of this observation, we may (and indeed, we need to) insist in what follows that π has only trivial monotone intervals.³ We now begin to relate the position of entries of π to the order of their encodings in w , as in the result below.

PROPOSITION 6.4. *Suppose that $\psi : V(G_\pi) \rightarrow V(\Gamma_D(w))$ is an isomorphism, that π^\sharp is a monotonically gridded permutation with only trivial monotone intervals, and that there are indices $i_1 < i_2 < i_3$ for which the entries $\pi(i_1)$, $\pi(i_2)$, and $\pi(i_3)$ all lie in the same cell of π^\sharp and are all encoded by the same letter. Then $\psi(i_1) < \psi(i_2) < \psi(i_3)$ or the reverse.*

Proof. Because $\pi(i_1)$, $\pi(i_2)$, and $\pi(i_3)$ all lie in the same cell of π^\sharp , they form a monotone subsequence. Because π has only trivial monotone intervals, there must be an entry $\pi(x)$ separating $\pi(i_1)$ and $\pi(i_2)$ and an entry $\pi(y)$ separating $\pi(i_2)$ and $\pi(i_3)$. Thus our situation resembles that shown in Figure 13.

Proposition 6.1 implies that either $\psi(i_1) < \psi(x) < \psi(i_2)$ or the reverse. We assume this inequality holds and prove that $\psi(i_1) < \psi(i_2) < \psi(i_3)$; if the reverse holds instead, then a symmetrical argument establishes that $\psi(i_1) > \psi(i_2) > \psi(i_3)$.

There are three more implications of Proposition 6.1, and Figure 14 shows how each of these implications updates our knowledge of the Hasse diagram of the indices

³We choose not to give such permutations a name, instead repeating this slightly cumbersome phrase, although these permutations have been given names in the literature. Atkinson and Stitt [8] refer to them as *strongly irreducible*, while Brignall, Jelínek, Kynčl, and Marchant [12] call them *adjacency-free*.




FIG. 14. The evolving Hasse diagrams showing what we know about the ordering of the values of ψ in the proof of Proposition 6.4.

$\psi(i_1)$, $\psi(i_2)$, $\psi(i_3)$, $\psi(x)$, and $\psi(y)$. We must have $\psi(i_1) < \psi(x) < \psi(i_3)$ because the reverse would contradict the fact that $\psi(i_1) < \psi(x)$, leading to the second diagram of Figure 14. We must then have $\psi(i_1) < \psi(y) < \psi(i_3)$ because the reverse would contradict the fact that $\psi(i_1) < \psi(i_3)$, leading to the third diagram of Figure 14. Finally, we must have $\psi(i_2) < \psi(y) < \psi(i_3)$ because the reverse would contradict the fact that $\psi(y) < \psi(i_3)$, leading to the fourth diagram of Figure 14.

We started by assuming (by symmetry) that $\psi(i_1) < \psi(i_2)$, and our final string of inequalities shows that $\psi(i_2) < \psi(i_3)$, completing the proof of the result. \square

6.3. Relettering. To briefly recap our situation, we have a monotonically gridded permutation $\pi^\#$ of length n with only trivial monotone intervals, an alphabet Σ , a decoder $D \subseteq \Sigma^2$, and a word $w \in \Sigma^n$ such that $G_\pi \cong \Gamma_D(w)$. Further, we are considering a particular isomorphism $\psi : V(G_\pi) \rightarrow V(\Gamma_D(w))$. We say that this collection of objects (the monotonically gridded permutation, the word, and the isomorphism) has the *one cell per letter property* if for every letter $a \in \Sigma$, all entries encoded by a lie in the same cell of $\pi^\#$. Put another way, this property says that if $w(\psi(i_1)) = w(\psi(i_2))$, then $\pi^\#(i_1)$ and $\pi^\#(i_2)$ lie in the same cell.

We can always obtain the one cell per letter property at the expense of a larger alphabet, a process that we call *relettering*. The first step is to expand our alphabet Σ to a set

$$\Sigma^\# = \{(a, k, \ell) : a \in \Sigma \text{ and } M(k, \ell) \neq 0\}$$

that includes cell information. Note that if $\pi^\#$ is a $t \times u$ gridding, then $|\Sigma^\#| \leq tu |\Sigma|$. Second, we add this information to our word w , producing a new word $w^\# \in (\Sigma^\#)^*$ in which

$$w^\#(\psi(i)) = (a, k, \ell) \iff w(\psi(i)) = a \text{ and } \pi(i) \text{ lies in cell } (k, \ell) \text{ of } \pi^\#.$$

Finally, we update the decoder to essentially ignore the extra information in the new letters, defining

$$D^\# = \{((a, i, j), (b, k, \ell)) : (a, b) \in D, M(i, j) \neq 0, \text{ and } M(k, \ell) \neq 0\}.$$

If we had an isomorphism $\psi : V(G_\pi) \rightarrow V(\Gamma_D(w))$ before the relettering, then after the relettering, by a very slight abuse of notation, we view this also as an isomorphism $\psi : V(G_\pi) \rightarrow V(\Gamma_{D^\#}(w^\#))$. We may do so because the relettering does not affect the letter graph or its labels at all: we have

$$V(\Gamma_{D^\#}(w^\#)) = V(\Gamma_D(w)) = \{1, 2, \dots, n\}$$

and

$$ij \in E(\Gamma_{D^\#}(w^\#)) \iff ij \in E(\Gamma_D(w)).$$

As we are assuming that π has only trivial monotone intervals, once we have done this relettering (to achieve the one cell per letter property), Proposition 6.4 implies that entries of the same letter “read in the same order” in the following sense. For each letter $a \in \Sigma^\#$, the entries encoded by a all lie in a single cell of $\pi^\#$. If these entries have indices $i_1 < i_2 < \dots < i_s$, then we have

$$\psi(i_1) < \psi(i_2) < \dots < \psi(i_s) \quad \text{or} \quad \psi(i_s) < \psi(i_{s-1}) < \dots < \psi(i_1).$$

This follows from Proposition 6.4 for $s \geq 3$, and is trivial for $s \leq 2$.

If $\psi(i_1) < \dots < \psi(i_s)$, then these entries are encoded from left to right in the word w , while if $\psi(i_s) < \dots < \psi(i_1)$, then they are encoded from right to left. We refer to this order as the *horizontal reading order of the letter*. If there is only one entry encoded by some letter, then we may choose the horizontal reading order arbitrarily to be either left-to-right or right-to-left.

Moreover, since the entries encoded by a form a monotone subsequence of π (because they lie together in a single cell of the monotonically gridded permutation $\pi^\#$), they are also encoded in w either from bottom to top or from top to bottom. We refer to this order as the *vertical reading order of the letter*. (Again, if there is only one such entry, then its vertical reading order may be chosen arbitrarily.) Note that the vertical (resp., horizontal) reading order of a letter is completely determined by its horizontal (resp., vertical) reading order together with the sign of the cell in which the entries encoded by that letter lie; for example, if the entries encoded by some letter lie in a decreasing cell and are encoded from bottom to top, then this means that they are also encoded from right to left.

We refer to the combination of the horizontal and vertical reading orders of a letter as simply the *reading order of the letter*. By our comments above, if we have the one cell per letter property, then the reading order of a letter is completely determined by its horizontal or vertical reading order and the sign of the cell in which the entries encoded by that letter lie.

6.4. Refining the gridding. Our next step is to refine the monotone gridding $\pi^\#$, by slicing it with additional vertical and horizontal lines, in a manner similar to that employed in [28, section 5]. The immediate aim of this additional slicing is to obtain the property in Proposition 6.5, which allows us to appeal to Proposition 6.2 in proving the subsequent results of this section.

The *rectangular hull* of a set of points in the plane is the smallest axis-parallel rectangle containing them. Suppose that we have our monotone gridded permutation $\pi^\#$ and that we have relettered in order to achieve the one cell per letter property. For each letter $(a, k, \ell) \in \Sigma^\#$, let $R_{(a,k,\ell)}$ denote the rectangular hull of the entries of $\pi^\#$ encoded by that letter. We now refine the gridding by adding, for each rectangle $R_{(a,k,\ell)}$, four lines that lie just outside its four sides. Thus for each letter $(a, k, \ell) \in \Sigma^\#$, we add two new horizontal lines and two new vertical lines. Consequently, because $\pi^\#$ is a $t \times u$ gridding, the maximum possible size of the new gridding is $(t + 2|\Sigma^\#|) \times (u + 2|\Sigma^\#|)$, which is at most $t(1 + 2u|\Sigma|) \times u(1 + 2t|\Sigma|)$ by our previous bound on $|\Sigma^\#|$.

As the original gridding is denoted by the sharp symbol $\#$, we use the double sharp symbol $\#$ to denote the *regridded permutation* $\pi^\#$. As stated above, the regridding gives us the following property.

PROPOSITION 6.5. *Let $\pi^\#$ and π^* be as above and suppose that for some indices $i_2 < i_3$, the entries $\pi(i_2)$ and $\pi(i_3)$ lie in the same column of π^* but are encoded by different letters of $\Sigma^\#$. Then there is an index i_1 satisfying $i_1 < i_2 < i_3$ so that $\pi(i_1)$ and $\pi(i_3)$ are encoded by the same letter of $\Sigma^\#$.*

Proof. Let $\pi(i_1)$ denote the leftmost entry of $\pi^\#$ that is encoded by the same letter of $\Sigma^\#$ as $\pi(i_3)$. In forming the regrided permutation π^* , we sliced by a line that lies just to the left of $\pi(i_1)$. Since $\pi(i_2)$ and $\pi(i_3)$ lie in the same column of π^* , this line must lie to the left of $\pi(i_2)$, and thus we must have $i_1 < i_2 < i_3$. \square

It is worth noting that the entries $\pi(i_1)$ and $\pi(i_3)$ of Proposition 6.5 must lie in the same cell of $\pi^\#$ since they are encoded by the same letter of $\Sigma^\#$, but that they need not lie in the same cell of π^* .

We now consider the reading orders of letters in the regrided permutation π^* . Note that π^* likely no longer has the one cell per letter property,⁴ but that property was only necessary to obtain reading orders for the letters of $\Sigma^\#$. In the regrided permutation π^* , we retain the reading orders of the letters of $\Sigma^\#$ from the relettering of section 6.3. Further note that if the reading order of a letter of $\Sigma^\#$ was chosen arbitrarily before, then that was because the letter encoded a single entry, and so in π^* , that entry is isolated in a cell by itself that shares neither a column nor a row with another nonempty cell. We now argue that these reading orders are compatible with each other.

PROPOSITION 6.6. *Suppose that π has only trivial monotone intervals and that $w^\#, \psi, \pi^\#,$ and π^* are as above. Then the letters of $\Sigma^\#$ encoding any pair of entries lying in the same column of π^* share the same horizontal reading order. Thus, by symmetry, the letters of $\Sigma^\#$ encoding any pair of entries lying in the same row of π^* also share the same vertical reading order.*

Proof. Suppose that for some indices $i_2 < i_3$, the entries $\pi(i_2)$ and $\pi(i_3)$ lie in the same column of π^* . We want to show that the letters encoding these two entries have the same horizontal reading order. If these entries are encoded by the same letter, then the result is trivial, so we assume that they are encoded by different letters.

First, suppose that $\pi(i_2)$ and $\pi(i_3)$ lie in different cells of $\pi^\#$. Proposition 6.5 shows directly that there is an index i_1 satisfying $i_1 < i_2 < i_3$ so that $\pi(i_1)$ and $\pi(i_3)$ are encoded by the same letter of $\Sigma^\#$. By symmetry, it also implies that there is an index i_4 satisfying $i_2 < i_3 < i_4$ so that $\pi(i_2)$ and $\pi(i_4)$ are encoded by the same letter of $\Sigma^\#$. As we are assuming that $\pi(i_2)$ and $\pi(i_3)$ do not share a cell in $\pi^\#$, the hypotheses of Proposition 6.2 are satisfied. Thus $\psi(i_1) < \psi(i_2) < \psi(i_3) < \psi(i_4)$ or the reverse, and this proves that the horizontal reading orders of the corresponding letters agree. By symmetry, we have the analogous statement about entries that lie in the same row of π^* but not the same cell of $\pi^\#$.

To complete the proof, suppose that the entries $\pi(i_2)$ and $\pi(i_3)$ lie in the same cell of $\pi^\#$. Because π has only trivial monotone intervals, these two entries must be separated, say by the entry $\pi(x)$. Thus $\pi(x)$ must lie in a different cell of $\pi^\#$ from $\pi(i_2)$ and $\pi(i_3)$, but in the same column or row of $\pi^\#$ as these entries.

If $\pi(x)$ lies in the same column as $\pi(i_2)$ and $\pi(i_3)$, then, by the previous case applied to the pair $\pi(x)$ and $\pi(i_2)$, we see that the horizontal reading order of the letter encoding $\pi(x)$ agrees with the horizontal reading order of the letter encoding $\pi(i_2)$. Similarly, the horizontal reading order of the letter encoding $\pi(x)$ agrees with the

⁴We could, of course, enlarge our alphabet again to regain the one cell per letter property, but there is no need.

horizontal reading order of the letter encoding $\pi(i_3)$. Thus, all three horizontal reading orders agree. Because $\pi(i_2)$ and $\pi(i_3)$ lie in the same cell of π^\sharp , this forces the vertical reading orders of the letters encoding these entries to agree as well. A symmetrical argument handles the case where $\pi(x)$ lies in the same row as $\pi(i_2)$ and $\pi(i_3)$. \square

With our next result we show that the regrided permutation π^* has a geometric realization, which implies that π itself lies in the geometric grid class of a matrix of bounded size.

PROPOSITION 6.7. *Suppose that the permutation π has only trivial monotone intervals, that $\pi \in \text{Grid}(M)$ for a $0/\pm 1$ matrix M of size $t \times u$, and that $G_\pi \cong \Gamma_D(w)$, where $|\Sigma| \leq r$, $D \subseteq \Sigma^2$, and $w \in \Sigma^*$. Then there is a $0/\pm 1$ matrix M^* of size at most $t(1 + 2ur) \times u(1 + 2tr)$ for which $\pi \in \text{Geom}(M^*)$.*

Proof. Let $\psi : V(G_\pi) \rightarrow V(\Gamma_D(w))$ be an isomorphism, and let π^\sharp be an M -gridding of π . Then reletter and regrid as described above to form a word $w^\sharp \in (\Sigma^\sharp)^*$ and a regrided permutation π^* . Under these hypotheses, our previous discussion implies that π^* has at most $t(1 + 2ur)$ columns and at most $u(1 + 2tr)$ rows. We may remove any empty columns or rows in π^* , and so we may assume that every column and every row of π^* contains at least one entry.

We now define column and row signs (c_k) and (r_ℓ) . Consider first the k th column of π^* . By removing empty columns, we have ensured that there is at least one letter that encodes an entry of this column. We set $c_k = 1$ if the horizontal reading order of this letter is left-to-right, and $c_k = -1$ if the horizontal reading order of this letter is right-to-left. Proposition 6.6 guarantees that this column sign is well defined—every letter encoding entries in this column has the same horizontal reading order. We assign row signs analogously, setting $r_\ell = 1$ if a letter encoding an entry of the ℓ th row has vertical reading order bottom-to-top, and $r_\ell = -1$ if it has vertical reading order top-to-bottom. Proposition 6.6 again shows that this row sign is well defined.

We now define a $0/\pm 1$ matrix M^* by $M^*(k, \ell) = c_k r_\ell$ for every k and ℓ . This is a partial multiplication matrix by definition, and it follows from the definition of M^* that π^* is an M^* -gridding of π . The proof will be complete if we can show that π^* has a geometric realization. To do this, we appeal to Proposition 5.2, which shows that π^* has a geometric realization if and only if the local orders $\{\triangleleft_k^{\text{col}}\}_k$ and $\{\triangleleft_\ell^{\text{row}}\}_\ell$ arising from π^* , (c_k) , and (r_ℓ) are consistent.

We prove that these local orders are consistent by establishing that ψ corresponds to a common linear extension of all of them. By symmetry, it suffices to consider a pair of entries in the same column, so suppose that for some indices $i_2 < i_3$, the entries $\pi(i_2)$ and $\pi(i_3)$ lie together in the k th column of π^* . We want to show that $\psi(i_2) < \psi(i_3)$ if and only if $c_k = 1$. This holds by the definition of c_k if $\pi(i_2)$ and $\pi(i_3)$ are encoded by the same letter of Σ^\sharp , so we may assume that these entries are encoded by different letters. In this case, we can apply Proposition 6.5 and its symmetry (as in the proof of Proposition 6.6) to see that there are indices i_1 and i_4 satisfying $i_1 < i_2 < i_3 < i_4$ so that $\pi(i_1)$ and $\pi(i_3)$ are encoded by the same letter of Σ^\sharp and $\pi(i_2)$ and $\pi(i_3)$ are encoded by the same letter of Σ^\sharp .

There are two cases, depending on whether $\pi(i_2)$ and $\pi(i_3)$ share a cell in π^\sharp . First, suppose that $\pi(i_2)$ and $\pi(i_3)$ lie in different cells of π^\sharp . Since $\pi(i_2)$ lies in a different cell of π^\sharp from $\pi(i_1)$ and $\pi(i_3)$, it follows that $\pi(i_2)$ separates these two entries. Thus, Proposition 6.1 shows that $\psi(i_1) < \psi(i_2) < \psi(i_3)$ or the reverse. Which of these inequalities we have is determined by the horizontal reading order of the letter that encodes $\pi(i_2)$ and $\pi(i_4)$. Since that reading order is encoded in c_k , it follows that $\psi(i_2) < \psi(i_3)$ if and only if $i_2 \triangleleft_k^{\text{col}} i_3$, which is what we desire in this case.

Now suppose that $\pi(i_2)$ and $\pi(i_3)$ lie in the same cell of π^\sharp , although note that this does not necessarily mean that they share the same cell of π^\times . Because π has only trivial monotone intervals, $\pi(i_2)$ and $\pi(i_3)$ must be separated by some entry $\pi(x)$ that lies outside their common cell in π^\sharp . Because all of the entries $\pi(i_1)$, $\pi(i_2)$, $\pi(i_3)$, and $\pi(i_4)$ lie in the same cell of π^\sharp , this means that $\pi(x)$ also separates $\pi(i_1)$ from $\pi(i_3)$ and $\pi(i_2)$ from $\pi(i_4)$. Therefore, we may apply Proposition 6.1 to see that if $c_k = 1$, then both $\psi(i_1) < \psi(x) < \psi(i_3)$ and $\psi(i_2) < \psi(x) < \psi(i_4)$, which implies that $\psi(i_2) < \psi(i_3)$, or if $c_k = -1$, then the reverse holds. In other words, as in the other case, $\psi(i_2) < \psi(i_3)$ if and only if $i_2 \leq_k^{\text{col}} i_3$.

This shows that the column orders are all consistent with ψ . By a symmetrical argument, it follows that the row orders are also consistent with ψ , and thus that the column and row orders are consistent with each other. This guarantees, via Proposition 5.2, that π^\times has a geometric realization, and thus completes the proof. \square

6.5. From one permutation to the entire class. We now conclude our proof of Theorem 1.1. Suppose we have a permutation class \mathcal{C} for which the corresponding graph class $G_{\mathcal{C}}$ has lettericity at most r . We want to show that \mathcal{C} is contained in some geometric grid class. By Corollary 3.3, we know that \mathcal{C} is monotonically griddable, so $\mathcal{C} \subseteq \text{Grid}(M)$ for some $0/\pm 1$ matrix M . Suppose that the matrix M has size $t \times u$.

Choose an arbitrary permutation $\pi \in \mathcal{C}$ of length n . As we remarked in section 6.2, it suffices to consider the case where π has only trivial monotone intervals. Because the lettericity of $G_{\mathcal{C}}$ is at most r , there are some alphabet Σ of size at most r , some decoder $D \subseteq \Sigma^2$, and some word $w \in \Sigma^n$ such that $G_\pi \cong \Gamma_D(w)$. We therefore have the hypotheses of Proposition 6.7, so there is some $0/\pm 1$ matrix M_π of size at most $t(1 + 2ur) \times u(1 + 2tr)$ for which $\pi \in \text{Geom}(M_\pi)$.

Therefore, for every permutation $\pi \in \mathcal{C}$ with only trivial monotone intervals, there is a $0/\pm 1$ matrix M_π of size at most $t(1 + 2ur) \times u(1 + 2tr)$ for which $\pi \in \text{Geom}(M_\pi)$. Thus, there is some finite $0/\pm 1$ matrix M that contains all such matrices M_π as submatrices,⁵ and thus we have $\mathcal{C} \subseteq \text{Geom}(M)$, completing the proof of Theorem 1.1.

7. Concluding remarks. We have proved that, when restricted to classes of inversion graphs, bounded lettericity is equivalent to geometric griddability of the corresponding permutation class. Several intriguing questions about both of these concepts remain. One area that is currently under active investigation is a characterization of the geometrically griddable permutation classes; Theorem 3.1 completely characterizes the *monotonically* griddable permutation classes, but as yet it has no geometric analogue. Similarly, there is no known characterization of graph classes of bounded lettericity, and the results established here suggest that it may be possible to answer both of these questions simultaneously.

Even if these characterizations remain out of reach, a decision procedure to determine if a given class (of permutations or graphs, specified by their minimal forbidden objects) is geometrically griddable or of bounded lettericity would be desirable.

There is an ongoing effort to develop robust tools for tackling these problems. In [4], Alecu and Lozin introduce the notion of *locally ordered hypergraphs*: hypergraphs enhanced with linear orders on their edges, subject to a local consistency

⁵In fact, this step of the proof does not require much expansion of these matrices. As in section 5, where we observed that $\text{Geom}(M) = \text{Geom}(M^{\times 2})$ for every $0/\pm 1$ matrix M , if we replace every entry by $\begin{pmatrix} -1 & \\ & -1 \end{pmatrix}$, it follows that for every $0/\pm 1$ matrix M of size $t \times u$, we have $\text{Geom}(M) \subseteq \text{Geom}(S_{t,u})$, where $S_{t,u}$ is the $2t \times 2u$ matrix defined by $S_{t,u}(k, \ell) = (-1)^{k+\ell-1}$. One therefore need only consider a large enough matrix of this form.

condition that the orders should agree where they overlap. This attempts to abstract the settings of monotone gridings and that of letter graph expressions. A proof of Theorem 1.1 with a similar outline as ours, but using the machinery of locally ordered hypergraphs, is presented in Alecu's thesis [3].

REFERENCES

- [1] M. ALBERT, M. ATKINSON, M. BOUVEL, N. RUŠKUC, AND V. VATTER, *Geometric grid classes of permutations*, Trans. Amer. Math. Soc., 365 (2013), pp. 5859–5881.
- [2] M. ALBERT, N. RUŠKUC, AND V. VATTER, *Inflations of geometric grid classes of permutations*, Israel J. Math., 205 (2015), pp. 73–108.
- [3] B. ALECU, *A Parametric Approach to Hereditary Classes*, Ph.D. thesis, University of Warwick, Warwick, UK, 2021.
- [4] B. ALECU AND V. LOZIN, *Understanding Lettericity I: A Structural Hierarchy*, preprint, <https://arxiv.org/abs/2106.03267>, 2021.
- [5] B. ALECU, V. LOZIN, AND D. DE WERRA, *The micro-world of cographs*, in Combinatorial Algorithms, L. Gąsieniec, R. Klasing, and T. Radzik, eds., Lecture Notes in Comput. Sci. 12126, Springer, Cham, 2020, pp. 30–42.
- [6] B. ALECU, V. LOZIN, D. DE WERRA, AND V. ZAMARAEV, *Letter graphs and geometric grid classes of permutations: Characterization and recognition*, Discrete Appl. Math., 283 (2020), pp. 482–494.
- [7] B. ALECU, V. LOZIN, V. ZAMARAEV, AND D. DE WERRA, *Letter graphs and geometric grid classes of permutations: Characterization and recognition*, in Combinatorial Algorithms, L. Brankovic, J. Ryan, and W. Smyth, eds., Lecture Notes in Comput. Sci. 10765, Springer, Cham, 2018, pp. 195–205.
- [8] M. ATKINSON AND T. STITT, *Restricted permutations and the wreath product*, Discrete Math., 259 (2002), pp. 19–36.
- [9] A. ATMINAS, *Classes of graphs without star forests and related graphs*, Discrete Math., 345 (2022), 113089, 22 pp.
- [10] A. ATMINAS AND V. LOZIN, *Labelled induced subgraphs and well-quasi-ordering*, Order, 32 (2015), pp. 313–328.
- [11] A. BRANDSTÄDT, V. B. LE, AND J. P. SPINRAD, *Graph Classes: A Survey*, SIAM Monogr. Discrete Math. Appl., SIAM, Philadelphia, 1999, <https://doi.org/10.1137/1.9780898719796>.
- [12] R. BRIGNALL, V. JELÍNEK, J. KYNČL, AND D. MARCHANT, *Zeros of the Möbius function of permutations*, Mathematika, 65 (2019), pp. 1074–1092.
- [13] R. BRIGNALL AND V. VATTER, *Labelled well-quasi-order for permutation classes*, Comb. Theory, to appear.
- [14] V. CHVÁTAL AND P. HAMMER, *Aggregation of inequalities in integer programming*, in Studies in Integer Programming, P. Hammer, E. Johnson, B. Korte, and G. Newhauser, eds., Ann. of Discrete Math. 1, North-Holland, Amsterdam, 1977, pp. 145–162.
- [15] V. DIEKERT, *Combinatorics on Traces*, Lecture Notes in Comput. Sci. 454, Springer-Verlag, Berlin, Germany, 1990.
- [16] R. FERGUSON AND V. VATTER, *Letter graphs and modular decomposition*, Discrete Appl. Math., 309 (2022), pp. 215–220.
- [17] S. FÖLDES AND P. HAMMER, *Split graphs*, Congr. Numer., 14 (1977), pp. 311–315.
- [18] M. GOLUBIC, *Threshold graphs and synchronizing parallel processes*, in Combinatorics (Proc. Fifth Hungarian Colloq.), Vol. I, A. Hajnal and V. Sós, eds., Colloq. Math. Soc. János Bolyai 18, North-Holland, Amsterdam, 1978, pp. 419–428.
- [19] M. GOLUBIC, *Algorithmic Graph Theory and Perfect Graphs*, 2nd ed., Ann. Discrete Math. 57, Elsevier, Amsterdam, 2004.
- [20] G. HIGMAN, *Ordering by divisibility in abstract algebras*, Proc. London Math. Soc. (3), 2 (1952), pp. 326–336.
- [21] C. HOMBERGER AND V. VATTER, *On the effective and automatic enumeration of polynomial permutation classes*, J. Symbolic Comput., 76 (2016), pp. 84–96.
- [22] S. HUCZYNSKA AND V. VATTER, *Grid classes and the Fibonacci dichotomy for restricted permutations*, Electron. J. Combin., 13 (2006), 54, 14 pp.
- [23] V. LOZIN, *From words to graphs, and back*, in Language and Automata Theory and Applications, C. Martín-Vide, A. Okhotin, and D. Shapira, eds., Lecture Notes in Comput. Sci. 11417, Springer, Cham, 2019, pp. 43–54.
- [24] N. MAHADEV AND U. PELED, *Threshold Graphs and Related Topics*, Ann. Discrete Math. 56, North-Holland, Amsterdam, 1995.

- [25] M. MURPHY AND V. VATTER, *Profile classes and partial well-order for permutations*, Electron. J. Combin., 9 (2003), 17, 30 pp.
- [26] M. PETKOVŠEK, *Letter graphs and well-quasi-order by induced subgraphs*, Discrete Math., 244 (2002), pp. 375–388.
- [27] Z. STANKOVA, *Forbidden subsequences*, Discrete Math., 132 (1994), pp. 291–316.
- [28] V. VATTER, *Small permutation classes*, Proc. Lond. Math. Soc. (3), 103 (2011), pp. 879–921.
- [29] V. VATTER, *Permutation classes*, in Handbook of Enumerative Combinatorics, M. Bóna, ed., CRC Press, Boca Raton, FL, 2015, pp. 754–833.
- [30] V. VATTER AND S. WATON, *On partial well-order for monotone grid classes of permutations*, Order, 28 (2011), pp. 193–199.
- [31] V. VATTER AND S. WATON, *On points drawn from a circle*, Electron. J. Combin., 18 (2011), 223, 10 pp.