GEOMETRIC GRID CLASSES OF PERMUTATIONS

MICHAEL H. ALBERT, M. D. ATKINSON, MATHILDE BOUVEL, NIK RUŠKUC,
 AND VINCENT VATTER

ABSTRACT. A geometric grid class consists of those permutations that can be drawn on a specified set of line segments of slope ± 1 arranged in a rectangular pattern governed by a matrix. Using a mixture of geometric and language theoretic methods, we prove that such classes are specified by finite sets of forbidden permutations, are partially well ordered, and have rational generating functions. Furthermore, we show that these properties are inherited by the subclasses (under permutation involvement) of such classes, and establish the basic lattice theoretic properties of the collection of all such subclasses.

1. INTRODUCTION

Subsequent to the resolution of the Stanley-Wilf Conjecture in 2004 by Marcus and Tardos [21], two major research programmes have emerged in the study of permutation classes:

- to characterise the possible growth rates of permutation classes and
- to provide necessary and sufficient conditions for permutation classes to have amenable generating functions.

With regard to the first programme, we point to the work of Kaiser and Klazar [20], who characterised the possible growth rates up to 2, and the work of Vatter [25], who extended this characterisation up to the algebraic number $\kappa \approx 2.20557$, the point at which infinite antichains begin to emerge, and where the transition from countably many to uncountably many permutation classes occurs. The second programme is illustrated by the work of Albert, Atkinson, and Vatter [3], who showed that all subclasses of the separable permutations not containing $\text{Av}(231)$ or a symmetry of this class have rational generating functions.

Both research programmes rely on structural descriptions of permutation classes, in particular, the notion of grid classes. Here we study the enumerative and order-theoretic properties of a certain type of grid class called a geometric grid class.

While we present formal definitions in the next two sections, geometric grid classes may be defined briefly as follows. Suppose that M is a $0/\pm 1$ matrix. The *standard figure* of M , which we typically denote by Λ , is the point set in $\mathbb{R}^2$ consisting of:

- the increasing open line segment from $(k-1, \ell-1)$ to (k, ℓ) if $M_{k,\ell} = 1$ or
- the decreasing open line segment from $(k-1, \ell)$ to $(k, \ell-1)$ if $M_{k,\ell} = -1$.

(Note that in order to simplify this correspondence, we index matrices first by column, counting left to right, and then by row, counting bottom to top throughout.)

FIGURE 1. The permutation 351624 on the left and the permutation 153426 on the right lie, respectively, in the geometric grid classes of

$$\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} \text{ and } \begin{pmatrix} -1 & 1 \\ 1 & -1 \end{pmatrix}.$$

The *geometric grid class* of M , denoted by $\text{Geom}(M)$, is then the set of all permutations that can be drawn on this figure in the following manner. Choose n points in the figure, with no two on a common horizontal or vertical line. Then label the points from 1 to n from bottom to top and record these labels reading from left to right.

Much of our perspective and inspiration comes from Steve Waton, who considered two particular geometric grid classes in his thesis [28]. These are the permutations which can be drawn from a circle, later studied by Vatter and Waton [27], and the permutations which can be drawn on an X , later studied by Elizalde [12]. Examples of these two grid classes are shown in Figure 1.¹

A permutation class is said to be *geometrically griddable* if it is contained in some geometric grid class. With that final piece of terminology, we can state the main results of this paper.

- **Theorem 6.1.** *Every geometrically griddable class is partially well ordered.* (Such classes do not contain infinite antichains.)
- **Theorem 6.2.** *Every geometrically griddable class is finitely based.* (These classes can be defined by only finitely many forbidden patterns.)
- **Theorem 8.1.** *Every geometrically griddable class is in bijection with a regular language and thus has a rational generating function.*
- **Theorem 9.1.** *The simple, sum indecomposable, and skew indecomposable permutations in every geometrically griddable class are each in bijection with a regular language and thus have rational generating functions.*
- **Theorem 10.3.** *The atomic geometrically griddable classes are precisely the geometric grid classes of $\bullet$ -isolated $0/\bullet/\pm 1$ matrices, and every geometrically griddable class can be expressed as a finite union of such classes.* (This type of geometric grid class is defined in Section 10.)

For the remainder of the introduction we present the (standard) definitions of permutation classes. In the next section we formalise the geometric notions of

¹The drawing on the left of Figure 1 has been drawn as a diamond to fit with the general approach of this paper, but by stretching and shrinking the x - and y -axes in the manner formalised in Section 2, it is clear that any permutation that can be “drawn on a diamond” can also be “drawn on a circle”.

this paper. Section 3 contains a brief discussion of grid classes of forests, while Section 4 introduces partial multiplication matrices. Sections 5–10 introduce a correspondence between permutations in a geometric grid class and words, and utilise this correspondence to establish the main results of the paper. Finally, in Section 11, we conclude with numerous open problems.

The permutation π of $\{1, 2, \dots, n\}$ *contains* or *involves* the permutation σ of $\{1, 2, \dots, k\}$ (written $\sigma \leq \pi$) if π has a subsequence of length k which is order isomorphic to σ . For example, $\pi = 391867452$ (written in list, or one-line notation) contains $\sigma = 51342$, as can be seen by considering the subsequence $\pi(2)\pi(3)\pi(5)\pi(6)\pi(9) = 91672$. A *permutation class* is a downset of permutations under this containment ordering; thus if $\mathcal{C}$ is a permutation class, $\pi \in \mathcal{C}$, and $\sigma \leq \pi$, then $\sigma \in \mathcal{C}$.

For any permutation class $\mathcal{C}$ there is a unique (and possibly infinite) antichain B such that

$$\mathcal{C} = \text{Av}(B) = \{\pi : \beta \not\leq \pi \text{ for all } \beta \in B\}.$$

This antichain B is called the *basis* of $\mathcal{C}$. We denote by $\mathcal{C}_n$ (for $n \in \mathbb{N}$) the set of permutations in $\mathcal{C}$ of length n , and we refer to

$$\sum_{n=0}^{\infty} |\mathcal{C}_n| x^n = \sum_{\pi \in \mathcal{C}} x^{|\pi|}$$

as the *generating function* of $\mathcal{C}$.

Finally, a permutation class, or indeed any partially ordered set, is said to be *partially well ordered* (*pwo*) if it contains neither an infinite strictly descending chain nor an infinite antichain. Of course, permutation classes cannot contain infinite strictly descending chains, so in this context being pwo is equivalent to having no infinite antichain.

2. THE GEOMETRIC PERSPECTIVE

Geometric ideas play a significant role in this paper, and to provide background we reintroduce permutations and the involvement relation in a somewhat nonstandard manner.

We call a subset of the plane a *figure*. We say that the figure $\mathcal{F} \subseteq \mathbb{R}^2$ is *involved* in the figure $\mathcal{G}$, denoted $\mathcal{F} \leq \mathcal{G}$, if there are subsets $A, B \subseteq \mathbb{R}$ and increasing injections $\phi_x : A \rightarrow \mathbb{R}$ and $\phi_y : B \rightarrow \mathbb{R}$ such that

$$\mathcal{F} \subseteq A \times B \text{ and } \phi(\mathcal{F}) = \{(\phi_x(a), \phi_y(b)) : (a, b) \in \mathcal{F}\} \subseteq \mathcal{G}.$$

The involvement relation is a *preorder* (it is reflexive and transitive but not necessarily antisymmetric) on the collection of all figures. If $\mathcal{F} \leq \mathcal{G}$ and $\mathcal{G} \leq \mathcal{F}$, then we say that $\mathcal{F}$ and $\mathcal{G}$ are (*figure*) *equivalent* and write $\mathcal{F} \approx \mathcal{G}$. Note that in the case of figures with only finitely many points, two figures are equivalent if and only if one can be transformed to the other by stretching and shrinking the axes. Three concrete examples are provided below.

- The figures $\mathcal{F} = \{(a, |a|) : a \in \mathbb{R}\}$ and $\mathcal{G} = \{(a, a^2) : a \in \mathbb{R}\}$ are equivalent. To see that $\mathcal{F} \leq \mathcal{G}$, take $A \times B = \mathbb{R} \times \mathbb{R}_{\geq 0}$, $\phi_x(a) = a$, and $\phi_y(b) = b^2$. To see that $\mathcal{G} \leq \mathcal{F}$, take $A \times B = \mathbb{R} \times \mathbb{R}_{\geq 0}$, $\phi_x(a) = a$, and $\phi_y(b) = \sqrt{b}$.
- The figures $\mathcal{F} = \{(a, a) : a \in \mathbb{R}\}$ and $\mathcal{G} = \{(a, a) : a \in [0, 1] \cup [2, 3]\}$ are equivalent. To see that $\mathcal{F} \leq \mathcal{G}$, take $A \times B = \mathbb{R}^2$, $\phi_x(a) = 1/(1 + e^{-a})$, and

$\phi_y(b) = 1/(1 + e^{-b})$. To see that $\mathcal{G} \leq \mathcal{F}$, consider the identity map with $A \times B = ([0, 1] \cup [2, 3])^2$.

- The “unit diamond” defined by $\mathcal{F} = \{(a, b) : |a| + |b| = 1\}$ (shown on the left of Figure 1) is equivalent to the unit circle $\mathcal{G} = \{(a, b) : a^2 + b^2 = 1\}$. To see that $\mathcal{F} \leq \mathcal{G}$, take $A \times B = [-1, 1]^2$, $\phi_x(a) = \sin(\pi a/2)$, and $\phi_y(b) = \cos(\pi(1 - b)/2)$. To see that $\mathcal{G} \leq \mathcal{F}$, simply consider the inverses of these maps.

To any permutation π of length n , we associate a figure which we call its *plot*, $\{(i, \pi(i))\}$. These figures have the important property that they are *independent*, by which we mean that no two points lie on a common horizontal or vertical line. It is clear that every finite independent figure is equivalent to the plot of a unique permutation, so we could *define* a permutation as an equivalence class of finite independent figures. Under this identification, the partial order on equivalence classes of finite independent figures is the same as the containment, or involvement, order defined in the introduction.

Every figure $\mathcal{F} \subseteq \mathbb{R}^2$ therefore naturally defines a permutation class,

$\text{Sub}(\mathcal{F}) = \{\text{permutations } \pi : \pi \text{ is equivalent to a finite independent subset of } \mathcal{F}\}$.

For example:

- Let $\mathcal{F} = \{(x, x) : x \in \mathbb{R}\}$. Then $\text{Sub}(\mathcal{F})$ contains a single permutation of each length, namely the identity, and its basis is $\{21\}$.
- Let $\mathcal{F} = \{(x, x) : x \in [0, 1]\} \cup \{(x+1, x) : x \in [0, 1]\}$. Then $\text{Sub}(\mathcal{F})$ consists of all permutations having at most one descent. Its basis is $\{321, 2143, 3142\}$, as can be seen by considering the ways in which two descents might occur.
- Let $\mathcal{F} = \{(x, \sin(x)) : x \in \mathbb{R}\}$. Then $\text{Sub}(\mathcal{F})$ is the set of all permutations. To establish this, note that every permutation can be broken into its increasing contiguous segments (“runs”) and points corresponding to each run can be chosen from increasing segments of $\mathcal{F}$. Its basis is, of course, the empty set.

A geometric grid class (as defined in the introduction) is precisely $\text{Sub}(\Lambda)$, where Λ denotes the standard figure of the defining matrix. Permutation classes of the form $\text{Sub}(\mathcal{F})$ in the special case where $\mathcal{F}$ is the plot of a bijection between two subsets of the real numbers have received some study before; we refer the reader to Atkinson, Murphy, and Ruškuc [7] and Huczynska and Ruškuc [18].

The lines $\{x = k : k = 0, \dots, t\}$ and $\{y = \ell : \ell = 0, \dots, u\}$ play a special role for standard figures, as they divide the figure into its cells. We extend this notion of *griddings* to all figures. First, though, we need to make a technical observation: because the real number line is order isomorphic to any open interval, it follows that every figure is equivalent to a bounded figure, and thus we may restrict our attention to bounded figures.

Let $R = [a, b] \times [c, d]$ be a rectangle in $\mathbb{R}^2$. A $t \times u$ -gridding of R is a tuple

$$G = (g_0, g_1, \dots, g_t; h_0, h_1, \dots, h_u)$$

of real numbers satisfying

$$\begin{aligned} a &= g_0 \leq g_1 \leq \dots \leq g_t = b, \\ c &= h_0 \leq h_1 \leq \dots \leq h_u = d. \end{aligned}$$

FIGURE 2. A 3×2 gridding of the rectangle $[a, b] \times [c, d]$ with the cells $C_{k,\ell}$ indicated.

These numbers are identified with the corresponding set of vertical and horizontal lines partitioning R into a rectangular collection of cells $C_{k,\ell}$ as shown in Figure 2. We often identify the gridding G and the collection of cells $C_{k,\ell}$.

A *gridded figure* is a pair $(\mathcal{F}, G)$, where $\mathcal{F}$ is a figure and G is a gridding containing $\mathcal{F}$ in its interior. To ensure that each point of $\mathcal{F}$ lies in a unique cell, we also require that the grid lines are disjoint from $\mathcal{F}$. By analogy with ungridded figures and permutations, we define the preorder $\leq$ and the equivalence relation $\approx$ for $t \times u$ gridded figures. The additional requirement is that the mapping $\phi = (\phi_x, \phi_y)$ appearing in the original definition maps the (k, ℓ) cell of one gridded figure to the (k, ℓ) cell of the other gridded figure. A *gridded permutation* is the equivalence class of a finite independent gridded figure.

The connection between finite figures, gridded finite figures, permutations, and gridded permutations can be formalised as follows. Let Φ and $\Phi^\sharp$, respectively, denote the set of all finite independent and finite gridded independent figures in $\mathbb{R}^2$, and let $\mathcal{S}$ and $\mathcal{S}^\sharp$, respectively, denote the set of all permutations and all gridded permutations. The (obvious) mappings connecting these sets are:

- $\Delta : \Phi^\sharp \rightarrow \Phi$, given by removing the grid lines, i.e., $(\mathcal{F}, G) \mapsto \mathcal{F}$;
- $\delta : \mathcal{S}^\sharp \rightarrow \mathcal{S}$, also given by removing the grid lines, i.e., in this context, mapping the equivalence class of $(\mathcal{F}, G)$ under $\approx$ to the equivalence class of $\mathcal{F}$ under $\approx$;
- $\pi^\sharp : \Phi^\sharp \rightarrow \mathcal{S}^\sharp$, which sends every element of $\Phi^\sharp$ to its equivalence class under $\approx$;
- $\pi : \Phi \rightarrow \mathcal{S}$, which sends each element of Φ to its equivalence class under $\approx$.

It is a routine matter to verify that the diagram

$$\begin{array}{ccc} \Phi^\sharp & \xrightarrow{\Delta} & \Phi \\ \pi^\sharp \downarrow & & \downarrow \pi \\ \mathcal{S}^\sharp & \xrightarrow{\delta} & \mathcal{S} \end{array}$$

commutes.

The gridings we are interested in restrict the content of cells in a way specified by a matrix. Given a $0/\pm 1$ matrix M , we say that the gridded figure $(\mathcal{F}, G)$ is

compatible with M if the following holds for all relevant k and ℓ :

$$\mathcal{F} \cap C_{k,\ell} \text{ is } \begin{cases} \text{increasing,} & \text{if } M_{k,\ell} = 1, \\ \text{decreasing,} & \text{if } M_{k,\ell} = -1, \\ \text{empty,} & \text{if } M_{k,\ell} = 0. \end{cases}$$

We are concerned with the images, under the maps $\pi^\sharp$ and $\pi\Delta = \delta\pi^\sharp$, of the set of all M -compatible, finite, independent, gridded figures. We denote these two sets by $\text{Grid}^\sharp(M)$ and $\text{Grid}(M)$, respectively. We refer to $\text{Grid}(M)$ as the (*monotone*) *grid class* of M .

In this context, the standard figure Λ (as defined in the introduction) of a $0/\pm 1$ matrix M of size $t \times u$ has *standard grid lines* G given by $\{x = k : k = 0, \dots, t\}$ and $\{y = \ell : \ell = 0, \dots, u\}$. We refer to the pair (Λ, G) as the *standard gridded figure* of M and denote it by $\Lambda^\sharp$; note that $\Lambda^\sharp$ is compatible (in the sense above) with M . The set of all gridded permutations contained in $\Lambda^\sharp$ is denoted by $\text{Geom}^\sharp(M)$, and its image under δ is $\text{Geom}(M)$, as defined in the introduction. Because the standard gridded figure of a $0/\pm 1$ matrix M is compatible with M , we always have

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

In the next section we characterise the matrices for which equality is achieved.

3. GRID CLASSES OF FORESTS

The first appearance of grid classes in the literature, in the special case where M is a permutation matrix and under the name “profile classes”, was in Atkinson [5]. Atkinson observed that such grid classes have polynomial enumeration. Huczynska and Vatter [19], who introduced the name “grid classes”, generalised Atkinson’s result to show that grid classes of signed permutation matrices have eventually polynomial enumeration. This result allowed them to give a more structural proof of Kaiser and Klazar’s “Fibonacci dichotomy” [20], which states that for every permutation class $\mathcal{C}$, either $|\mathcal{C}_n|$ is greater than the n th Fibonacci number for all n or $|\mathcal{C}_n|$ is eventually polynomial.

Atkinson, Murphy, and Ruškuc [6] and Albert, Atkinson, and Ruškuc [2] studied the special case of grid classes of $0/\pm 1$ vectors, under the name “ W -classes”. The former paper proves that such grid classes are pwo and finitely based, while the latter shows that they have rational generating functions.

The first appearance of grid classes of $0/\pm 1$ matrices in full generality was in Murphy and Vatter [22] (again under the name “profile classes”). They were interested in the two properties of such classes, and to state their result we need a definition.

The *cell graph* of M is the graph on the vertices $\{(k, \ell) : M_{k,\ell} \neq 0\}$ in which (k, ℓ) and (i, j) are adjacent if the corresponding cells of M share a row or a column and there are no nonzero entries between them in this row or column. We say that the matrix M is a *forest* if its cell graph is a forest. Viewing the absolute value of M as the adjacency matrix of a bipartite graph, we obtain a different graph, its *row-column graph*.² It is not difficult to show that the cell graph of a matrix is a forest if and only if its row-column graph is also a forest. An example of each

²In other words, the row-column graph of a $t \times u$ matrix M is the bipartite graph on the vertices $x_1, \dots, x_t, y_1, \dots, y_u$ where there is an edge between x_k and y_ℓ if and only if $M_{k,\ell} \neq 0$.

FIGURE 3. A matrix together with its cell graph (centre) and row-column graph (right).

of these graphs is shown in Figure 3. These graphs completely determine the two properties of grid classes:

Theorem 3.1 (Murphy and Vatter [22], later generalised by Brignall [10]). *The class $\text{Grid}(M)$ is two if and only if M is a forest.*

It has long been conjectured that $\text{Grid}(M)$ has a rational generating function if M is a forest;³ for example, by Huczynska and Vatter [19, Conjecture 2.8]. Indeed, this conjecture was the original impetus for the present work. However, as the work progressed, it became apparent that the geometric paradigm provided a viewpoint which was at once more insightful and more general, and thus our perspective shifted. The link with the original motivation is provided by the following result.

Theorem 3.2. *If M is a forest, then $\text{Grid}^\sharp(M) = \text{Geom}^\sharp(M)$ and thus $\text{Grid}(M) = \text{Geom}(M)$.*

Proof. The proof is by induction on the number of nonzero entries of M . For the case of a single nonzero cell, note that one can place any increasing (resp., decreasing) set of points on a line of slope 1 (resp., -1) by applying a horizontal transformation.

Now suppose that M has two or more nonzero entries, denote its standard gridded figure by $\Lambda_M^\sharp = (\Lambda_M, G)$, and let (k, ℓ) denote a leaf in the cell graph of M . By considering the transpose of M if necessary, we may assume that there are no other nonzero entries in column k of M . Let $\pi^\sharp$ be an arbitrary gridded permutation in $\text{Grid}^\sharp(M)$. We aim to show that $\pi^\sharp \in \text{Geom}^\sharp(M)$.

Denote by N the matrix obtained from M by setting the (k, ℓ) entry equal to 0, and denote its standard gridded figure by $\Lambda_N^\sharp = (\Lambda_N, G)$. Note that the grid lines of $\Lambda_N^\sharp$ and $\Lambda_M^\sharp$ are identical because the corresponding matrices are the same size. Let $\sigma^\sharp$ denote the gridded permutation obtained from $\pi^\sharp$ by removing all entries in the (k, ℓ) cell. Because $\sigma^\sharp$ lies in $\text{Grid}^\sharp(N)$, which by induction is equal to $\text{Geom}^\sharp(N)$, there is a finite independent point set $S \subseteq \Lambda_N \subseteq \Lambda_M$ such that $(S, G) \approx \sigma^\sharp$. As long as we do not demand that the new points belong to Λ_M , it is clear that we can extend S by adding points in the (k, ℓ) cell to arrive at a point set $P \supseteq S$ such that $(P, G) \approx \pi^\sharp$. Then we can apply a horizontal transformation to column k to move these new points onto the diagonal line segment of Λ_M in this cell. This horizontal transformation does not affect the points of S , because none of those points lie in column k , so we see that $P \subseteq \Lambda_M$, and thus $\pi^\sharp \in \text{Geom}^\sharp(M)$, as desired. $\square$

³Note that $\text{Grid}(M)$ can have a nonrational generating function when M is not a forest. An example of this is given in the conclusion.

FIGURE 4. The permutation 2413 lies in $\text{Grid} \left(\begin{smallmatrix} -1 & 1 \\ 1 & -1 \end{smallmatrix} \right)$ but not $\text{Geom} \left(\begin{smallmatrix} -1 & 1 \\ 1 & -1 \end{smallmatrix} \right)$. If 2413 did lie in this geometric grid class, then beginning with the 2 and moving in a clockwise direction, we see that the 4 lies to the right of the 2 and thus closer to the centre, the 3 lies closer than the 4 to the centre, the 1 lies closer than the 3 to the centre, and finally, to reach a contradiction, the 2 must lie even closer to the centre than the 1.

Because of Theorem 3.2, all of our results about geometric grid classes yield immediate corollaries to grid classes of forests, which we shall generally not mention. For example, our upcoming Theorem 6.1 generalises one direction of Theorem 3.1.

The fact that $\text{Grid}(M)$ is not pwo when M is not a forest (the other direction of Theorem 3.1), combined with Theorem 6.1 which shows that all geometric grid classes are pwo, implies that the converse to Theorem 3.2 also holds: if $\text{Grid}(M) = \text{Geom}(M)$, then M is a forest. This fact can also be established by arguments generalising those accompanying Figure 4.

4. PARTIAL MULTIPLICATION MATRICES

In this section we consider a particular “refinement” operation on matrices, which is central to our later arguments. Let M be a $0/\pm 1$ matrix of size $t \times u$ and q be a positive integer. The *refinement* $M^{\times q}$ of M is the $0/\pm 1$ matrix of size $qt \times qu$ obtained from M by replacing each 1 by a $q \times q$ identity matrix (which, by our conventions, has ones along its southwest-northeast diagonal), each -1 by a negative $q \times q$ anti-identity matrix, and each 0 by a $q \times q$ zero matrix. It is easy to see that the standard figure $M^{\times q}$ is equivalent to the standard figure of M , so $\text{Geom}(M^{\times q}) = \text{Geom}(M)$ for all q (although, of course, the corresponding gridded classes differ).

The refinements $M^{\times 2}$ play a special role throughout this paper. To explain this we first need a definition. We say that a $0/\pm 1$ 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 entry $M_{k,\ell}$ is equal to either 0 or the product $c_k r_\ell$.

As our next result shows, we are never far from a partial multiplication matrix.

Proposition 4.1. *For every $0/\pm 1$ matrix M , its refinement $M^{\times 2}$ is a partial multiplication matrix.*

Proof. By construction, $M^{\times 2}$ is made up of 2×2 blocks equal to

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

From this it follows that $(M^{\times 2})_{k,\ell} \in \{0, (-1)^{k+\ell}\}$. Therefore we may take $c_k = (-1)^k$ and $r_\ell = (-1)^\ell$ as our column and row signs. $\square$

Because $\text{Geom}(M) = \text{Geom}(M^{\times 2})$, we may always assume that the matrices we work with are partial multiplication matrices. We record this useful fact below.

Proposition 4.2. *Every geometric grid class is the geometric grid class of a partial multiplication matrix.*

5. WORDS AND ENCODINGS

From the point of view of our goals in this paper, subword-closed languages over a finite alphabet display model behaviour: all such languages are defined by finite sets of forbidden subwords, are pwo under the subword order, and have rational generating functions. The brunt of our subsequent effort is focused on transferring these favourable properties from words to permutations.

Let Σ be a finite alphabet and Σ^* be the set of all finite words (i.e., sequences) over Σ . This set is partially ordered by means of the *subword* or *subsequence* order: $v \leq w$ if one can obtain v from w by deleting letters.

Subsets of Σ^* are called *languages*. We say that a language is *subword-closed* if it is a downward closed set in the subword order (such languages are also called *piecewise testable* by some; for example, Simon [23]). To borrow terminology from permutation classes, we say that the *basis* of a subword-closed language L is the set of minimal words which do not lie in L . It follows that L consists of precisely those words which do not contain any element of its basis. Moreover, a special case of a result of Higman [15] implies that subword-closed languages have finite bases:

Higman's Theorem. *The set of words over any finite alphabet is pwo under the subword order.*

The following characterisation of subword-closed languages is folkloric and follows directly from the fact that, for any finite set of forbidden subwords, there exists a finite state automaton accepting words not containing these subwords, which is acyclic except for loops at individual states.

Proposition 5.1. *Let Σ be a finite alphabet. Every nonempty subword-closed language over Σ can be expressed as a finite union of languages of the form*

$$\Sigma_1^* \{\varepsilon, a_2\} \Sigma_3^* \{\varepsilon, a_4\} \dots \Sigma_{2q-1}^* \{\varepsilon, a_{2q}\} \Sigma_{2q+1}^*,$$

where $q \geq 0$, $a_2, \dots, a_{2q} \in \Sigma$ and $\Sigma_1, \dots, \Sigma_{2q+1} \subseteq \Sigma$.

This fact shows that subword-closed languages are *regular languages*. To recall their definition briefly, given a finite alphabet Σ , the empty language $\emptyset$, the empty word language $\{\varepsilon\}$, and the singleton languages $\{a\}$ for each $a \in \Sigma$ are regular. Moreover, given two regular languages K and L over Σ , their union $K \cup L$, their concatenation $KL = \{vw : v \in K \text{ and } w \in L\}$, and the star $L^* = \{v^{(1)} \dots v^{(m)} : m \geq 0 \text{ and } v^{(1)}, \dots, v^{(m)} \in L\}$ are also regular. Alternatively, one may define regular languages as those accepted by a deterministic finite state automaton. From this second viewpoint, it is easy to see that given two regular languages K and L ,

FIGURE 5. An example of the map φ for a matrix with row signs $r_1 = -1$ and $r_2 = 1$ and column signs $c_1 = -1$, $c_2 = c_3 = 1$. Here we see that $\varphi(a_{31}a_{31}a_{22}a_{21}a_{11}a_{32}a_{22}) = 1527436$.

their complement $K \setminus L$ is also regular, a property we use many times. We say that the *generating function* of the language L is $\sum x^{|w|}$, where the sum is taken over all $w \in L$ and $|w|$ denotes the number of letters in w , i.e., its *length*. We use only the most basic properties of regular languages, for which we refer the reader to Flajolet and Sedgewick [13, Section I.4 and Appendix A.7]. In particular, the following fact is of central importance throughout the paper.

Theorem 5.2. *Every regular language has a rational generating function.*

We now describe the correspondence between permutations in a geometric grid class, $\text{Geom}(M)$, and words over an appropriate finite alphabet. This encoding, essentially introduced by Vatter and Waton [26], is fundamental to all of our proofs.

By Proposition 4.2, we may assume that M is a $t \times u$ partial multiplication matrix with column and row signs $c_1, \dots, c_t$ and $r_1, \dots, r_u$. Let $\Lambda^\sharp$ denote the standard gridded figure of M , and define the *cell alphabet* of M to be

$$\Sigma = \{a_{k\ell} : M_{k,\ell} \neq 0\}.$$

Intuitively, the letter $a_{k\ell}$ represents an instruction to place a point in an appropriate position on the line in the (k, ℓ) cell of $\Lambda^\sharp$. This appropriate position is determined as follows, and the whole process is depicted in Figure 5.

We say that the *base line* of a column of $\Lambda^\sharp$ is the grid line to the left (resp., right) of that column if the corresponding column sign is 1 (resp., -1). Similarly, the base line of a row of $\Lambda^\sharp$ is the grid line below (resp., above) that row if the corresponding row sign is 1 (resp., -1). We designate the intersection of the two base lines of a cell as its *base point*. Note that the base point is an endpoint of the line segment of Λ lying in this cell. As this definition indicates, we interpret the column and row signs as specifying the direction in which the columns and rows are “read”. Owing to this interpretation, we represent the column and row signs in our figures by arrows, as shown in Figure 5.

To every word $w = w_1 \cdots w_n \in \Sigma^*$ we associate a permutation $\varphi(w)$ as follows. First we choose arbitrary distances

$$0 < d_1 < \cdots < d_n < 1.$$

Next, for each i , we let p_i be the point on the line segment in cell $C_{k,\ell}$, where $w_i = a_{k,\ell}$, at infinity-norm distance d_i from the base point of $C_{k,\ell}$. It follows from

our choice of distances $d_1, \dots, d_n$ that $p_1, \dots, p_n$ are independent, and we define $\varphi(w)$ to be the permutation which is equivalent to the set $\{p_1, \dots, p_n\}$ of points.

It is a routine exercise to show that $\varphi(w)$ does not depend on the particular choice of distances $d_1, \dots, d_n$, showing that the mapping $\varphi : \Sigma^* \rightarrow \text{Geom}(M)$ is well defined. Of course there is a gridded counterpart $\varphi^\# : \Sigma^* \rightarrow \text{Geom}^\#(M)$, whereby we retain the grid lines coming from the figure $\Lambda^\#$.

The basic properties of φ and $\varphi^\#$ are described by the following result.

Proposition 5.3. *The mappings φ and $\varphi^\#$ are length-preserving, finite-to-one, onto, and order-preserving.*

Proof. That φ is length-preserving is obvious, as it maps letters in a word to entries in a permutation; that it is finite-to-one follows immediately from this.

In order to prove that φ is onto, let $\pi \in \text{Geom}(M)$ and choose a finite set $P = \{p_1, \dots, p_n\} \subseteq \Lambda$ of points which represent π (where as usual Λ denotes the standard figure of M). Suppose that the point p_i belongs to the cell (k_i, ℓ_i) of $\Lambda^\#$, and let d_i denote the infinity-norm distance from p_i to the base point of this cell. The points in P are independent because they are equivalent to a permutation. Therefore, we may move the points of P independently by small amounts without affecting its (figure) equivalence class, and thus may assume that the distances d_i are distinct. By reordering the points if necessary, we may also assume that $d_1 < \dots < d_n$. It is then clear that $\varphi(a_{k_1 \ell_1} \dots a_{k_n \ell_n}) = \pi$, so φ is indeed onto.

It remains to show that φ is order-preserving. Suppose that $v = v_1 \dots v_k, w = w_1 \dots w_n \in \Sigma^*$ satisfy $v \leq w$. Thus there are indices $1 \leq i_1 < \dots < i_k \leq n$ such that $v = w_{i_1} \dots w_{i_k}$. Note that if $\varphi(w)$ is represented by the point set $\{p_1, \dots, p_n\}$ via the sequence of distances $d_1 < \dots < d_n$, then $\varphi(v)$ is represented by the point set $\{p_{i_1}, \dots, p_{i_k}\}$ via the sequence of distances $d_{i_1} < \dots < d_{i_k}$, so $\varphi(v) \leq \varphi(w)$.

The proofs for the gridded version $\varphi^\#$ are analogous. $\square$

Using this correspondence between words and permutations, one may give an alternative proof of Theorem 3.2, showing that $\text{Grid}^\#(M) = \text{Geom}^\#(M)$ and thus that $\text{Grid}(M) = \text{Geom}(M)$ when M is a forest. As Proposition 5.3 shows that $\varphi^\#$ maps onto $\text{Geom}^\#(M)$, one only needs to show that it also maps onto $\text{Grid}^\#(M)$ when M is a forest. This is proved directly in Vatter and Waton [26].

6. PARTIAL WELL ORDER AND FINITE BASES

We now use the encoding $\varphi : \Sigma^* \rightarrow \text{Geom}(M)$ from the previous section to establish structural properties of geometrically griddable classes, i.e., subclasses of geometric grid classes. We begin with partial well order. By Theorem 3.2, this generalises one direction of Theorem 3.1.

Theorem 6.1. *Every geometrically griddable class is partially well ordered.*

Proof. Let $\mathcal{C}$ be a geometrically griddable class. By Proposition 4.2, $\mathcal{C} \subseteq \text{Geom}(M)$ for some partial multiplication matrix M . As partial well order is inherited by subclasses, it suffices to prove that $\text{Geom}(M)$ is pwo. Let $\varphi : \Sigma^* \rightarrow \text{Geom}(M)$, where Σ is the cell alphabet of M , be the encoding introduced in Section 5. Take $A \subseteq \text{Geom}(M)$ to be an antichain. For every $\alpha \in A$ there is some $w_\alpha \in \Sigma^*$ such that $\varphi(w_\alpha) = \alpha$ because φ is onto (Proposition 5.3). The set $\{w_\alpha : \alpha \in A\}$ must be an antichain in Σ^* because φ is order-preserving (Proposition 5.3 again). Higman's Theorem therefore shows that this set, and thus also A , is finite, as desired. $\square$

Our next goal is the following result.

Theorem 6.2. *Every geometrically griddable class is finitely based.*

We first make an elementary observation.

Proposition 6.3. *The union of a finite number of geometrically griddable classes is geometrically griddable.*

Proof. Let $\mathcal{C}$ and $\mathcal{D}$ be geometrically griddable classes. Thus there are matrices M and N such that $\mathcal{C} \subseteq \text{Geom}(M)$ and $\mathcal{D} \subseteq \text{Geom}(N)$. It follows that $\mathcal{C} \cup \mathcal{D} \subseteq \text{Geom}(P)$ for any matrix P which contains copies of both M and N . The result for arbitrary finite unions follows by iteration. $\square$

Given any permutation class $\mathcal{C}$, we let $\mathcal{C}^{+1}$ denote the class of *one-point extensions* of elements of $\mathcal{C}$, that is, $\mathcal{C}^{+1}$ is the set of all permutations π which contain an entry whose removal yields a permutation in $\mathcal{C}$. Every basis element of a class $\mathcal{C}$ is necessarily a one-point extension of $\mathcal{C}$ because the removal of *any* entry of a basis element of $\mathcal{C}$ yields a permutation in $\mathcal{C}$. Since bases of permutation classes are necessarily antichains, Theorem 6.2 will follow from the following result and Theorem 6.1.

Theorem 6.4. *If the class $\mathcal{C}$ is geometrically griddable, then the class $\mathcal{C}^{+1}$ is also geometrically griddable.*

Proof. It suffices to prove the result for geometric grid classes themselves, and by Proposition 4.2 we may further restrict our attention to the case of $\text{Geom}(M)$ where M is a partial multiplication matrix. Take $\tau \in \text{Geom}(M)^{+1}$ to be a one-point extension of $\pi \in \text{Geom}(M)$. Letting $\Lambda^\sharp = (\Lambda, G)$ denote the standard gridded figure of M , there is a finite independent set $P \subseteq \Lambda$ such that (P, G) is equivalent to some gridding of π . Now we may add a point, say x , to P to obtain an independent point set which is equivalent to τ . By moving x without affecting the equivalence class of $P \cup \{x\}$, we may further assume that x lies in the interior of a cell.

Let h and v denote, respectively, the horizontal and vertical lines passing through x , as shown on the left of Figure 6. Our goal is to create a new, refined gridding of Λ which contains v and h as grid lines.

The *offset* of a horizontal (resp., vertical) line is the distance between that line and the base line of the row (resp., column) of $\Lambda^\sharp$ that it slices through (recall from Section 5 that the standard gridded figure of a partial multiplication matrix has designated base lines determined by its column and row signs). From the definition of a partial multiplication matrix, it follows that if a vertical line of offset d slices a nonempty cell of Λ , it intersects the line segment in this cell precisely where the horizontal line of offset d slices the line segment.

Because the cells of a standard gridded figure have unit width, v and h have offsets strictly between 0 and 1, say $0 < d_1 \leq d_2 < 1$. By possibly moving x slightly without affecting the equivalence class of $P \cup \{x\}$ (which we may do because $P \cup \{x\}$ is independent), we may assume that $0 < d_1 < d_2 < 1$. We now define our refined gridding H , an example of which is shown on the right of Figure 6. We take H to consist of the grid lines of G together with the $2t$ vertical lines which pass through each column of Λ at the offsets d_1 and d_2 , and the $2u$ horizontal lines which pass through each row of Λ at the same two offsets. It follows from our observations above that (Λ, H) is equivalent to the standard gridded figure of $M^{\times 3}$

FIGURE 6. On the left is the standard gridded figure of a partial multiplication matrix, together with two additional (dashed) lines intersecting at the point x . On the right is the refinement defined in the proof of Theorem 6.4.

and in particular consists of a grid containing line segments of slope ± 1 , each of which runs from corner to corner in its cell. By possibly moving the points slightly, we may assume that no point of P lies on a grid line in H (although x lies on two such lines).

By cutting the figure Λ at the lines v and h and moving the four resulting pieces, we can create a new column and row whose cell of intersection contains x . (One can also view this as expanding the lines v and h into a column and a row, respectively.) We can then fill this cell of intersection with a line segment of slope ± 1 (the choice is immaterial) running from corner to corner; clearly, x can be shifted onto this line segment without affecting the equivalence class of $P \cup \{x\}$. The resulting gridded figure consists of line segments that run from corner to corner in their cells and is equivalent to the standard gridded figure of some partial multiplication matrix N which is of size $(3t + 1) \times (3u + 1)$, from which it follows that $\tau \in \text{Geom}(N)$. As there are only finitely many such matrices, we see that $\text{Geom}^{+1}(M)$ is contained in a finite union of geometric grid classes and so is geometrically griddable by Proposition 6.3. $\square$

7. GRIDDED PERMUTATIONS AND TRACE MONOIDS

In the next two sections we show that geometrically griddable classes have rational generating functions. We divide this task into two stages. First, using trace monoids, we show that every *gridded* class $\text{Geom}^\sharp(M)$ is in bijection with a regular language. Then, because a permutation may have many valid griddings, the second part of our task, undertaken in the next section, is to remove this multiplicity, allowing us to enumerate the geometric grid class $\text{Geom}(M)$ and all its subclasses.

To illustrate the issue we seek to understand in this section, we refer the reader to Figure 7. This example shows two different words which map to the same gridded permutation under $\varphi^\sharp$. This happens because the order in which points are consecutively inserted into *independent* cells — i.e., cells which share neither a column nor a row — is immaterial. On the language level, this means that letters which correspond to such cells “commute”. For example, in Figure 7, the

FIGURE 7. Two drawings of a particular gridding of the permutation 2465371. The drawing on the left is encoded as $a_{31}a_{32}a_{11}a_{22}a_{31}a_{32}a_{11}$, while the drawing on the right is encoded as $a_{31}a_{11}a_{32}a_{22}a_{31}a_{32}a_{11}$.

second and third letters of the words can be interchanged without changing the gridded permutation, which corresponds to placing the points p_2 and p_3 at different distances from their base points.

More precisely, suppose that M is a partial multiplication matrix with cell alphabet Σ . We say that the two words $v, w \in \Sigma^*$ are *equivalent* if one can be obtained from the other by successively interchanging adjacent letters which represent independent cells. The equivalence classes of this relation therefore form a *trace monoid*, which can be defined by the presentation

$$\langle \Sigma \mid a_{ij}a_{k\ell} = a_{k\ell}a_{ij} \text{ whenever } i \neq k \text{ and } j \neq \ell \rangle.$$

An element of this monoid (an equivalence class of words in Σ^*) is called a *trace*. It is an elementary and foundational fact about presentations that two words lie in the same trace if and only if one can be obtained from the other by a finite sequence of applications of the defining relations $a_{ij}a_{k\ell} = a_{k\ell}a_{ij}$ whenever $i \neq k$ and $j \neq \ell$ (Howie [17, Proposition 1.5.9 and Section 1.6]). The theory of trace monoids is understood to considerable depth; see for example Diekert [11].

Proposition 7.1. *Let M be a partial multiplication matrix with cell alphabet Σ . For words $v, w \in \Sigma^*$, we have $\varphi^\#(v) = \varphi^\#(w)$ if and only if v and w have the same trace in the trace monoid of M .*

Proof. It is clear from the preceding discussion that $\varphi^\#(v) = \varphi^\#(w)$ whenever v and w have the same trace.

Now suppose that $\varphi^\#(v) = \varphi^\#(w)$. We aim to prove that v and w have the same trace in the trace monoid of M . Clearly v and w must have the same length, say n . We prove our assertion by induction on n . In the case where $n = 1$, note that $\varphi^\#(a_{k,\ell})$ consists of a single point in cell $C_{k,\ell}$, so if $|v| = |w| = 1$, then $\varphi^\#(v) = \varphi^\#(w)$ implies that $v = w$, and thus v and w are in the same trace trivially. Thus we assume that $n \geq 2$ and that the assertion is true for all words of length $n - 1$.

Write $v = v_1 \cdots v_n$ and $w = w_1 \cdots w_n$. If $v_1 = w_1$, then from the definition of $\varphi^\#$ we have $\varphi^\#(v_2 \cdots v_n) = \varphi^\#(w_2 \cdots w_n)$, and the assertion follows by induction. Therefore we may assume that $v_1 \neq w_1$.

Suppose that the column and row signs of M are $c_1, \dots, c_t$ and $r_1, \dots, r_u$ and that $v_1 = a_{k\ell}$. While there are four cases depending on the column and row signs c_k and r_ℓ , they are essentially identical, so we assume that $c_k = r_\ell = 1$. This means that

among the points created by following the definition of $\varphi^\sharp$, the point corresponding to v_1 in $\varphi^\sharp(v)$ is the leftmost point in its column and bottommost point in its row. Since $\varphi^\sharp(v) = \varphi^\sharp(w)$ as *gridded* permutations, $\varphi^\sharp(w)$ must also contain a point in cell $M_{k,\ell}$ which is the leftmost point in its column and bottommost point in its row; suppose this point corresponds to w_i . Because $c_k = 1$, the entries in column k are placed from left to right, so w_i must be the first letter of w which corresponds to a cell in column k . Similarly, w_i must be the first letter of w which corresponds to a cell in row ℓ . Therefore all of the letters $w_1, \dots, w_{i-1}$ correspond to cells which are in different rows and columns from $M_{k,\ell}$. This shows that w lies in the same trace as the word

$$w' = w_i w_1 \cdots w_{i-1} w_{i+1} \cdots w_n = v_1 w_1 \cdots w_{i-1} w_{i+1} \cdots w_n.$$

Observe that $\varphi^\sharp(w') = \varphi^\sharp(w)$ and thus $\varphi^\sharp(w') = \varphi^\sharp(v)$. Finally, because v and w' have the same first letter, we see that v and w' are in the same trace by the first case of this argument, which implies that v and w are in the same trace, proving the proposition. $\square$

This result established, we are reduced to the task of choosing from each trace a unique representative, which is a well-understood problem.

Proposition 7.2 (Anisimov and Knuth; see Diekert [11, Corollary 1.2.3]). *In any trace monoid, it is possible to choose a unique representative from each trace in such a way that the resulting set of representatives forms a regular language.*

We immediately obtain the following.

Corollary 7.3. *For every partial multiplication matrix M , the (gridded) class $\text{Geom}^\sharp(M)$ is in bijection with a regular language.*

8. REGULAR LANGUAGES AND RATIONAL GENERATING FUNCTIONS

We now shift our attention to the *ungridded* permutations in a geometrically griddable class. This is the most technical argument of the paper, and we outline the general approach before delving into the formalisation.

The crux of the issue relates to the ungridded version of the encoding map φ and arises because it is possible that $\varphi(v) = \varphi(w)$ for two words $v, w \in \Sigma^*$ even though $\varphi^\sharp(v) \neq \varphi^\sharp(w)$. This happens precisely when $\pi = \varphi(v) = \varphi(w)$ admits two different griddings. To discuss these different griddings, we say that the gridded permutation $(\pi, G) \in \text{Geom}^\sharp(M)$ is a $\text{Geom}^\sharp(M)$ *gridding* of the permutation $\pi \in \text{Geom}(M)$. Since our goal is to establish a bijection between any geometrically griddable class $\mathcal{C}$ and a regular language, if a permutation in $\mathcal{C}$ has multiple $\text{Geom}^\sharp(M)$ griddings we must choose only one. To do so, we introduce a total order on the set of all such griddings of a fixed permutation π and aim to keep only those which are minimal in this order.

The problem is thus translated to that of recognising minimal $\text{Geom}^\sharp(M)$ griddings: Given a word $w \in \Sigma^*$, how can we determine if the gridded permutation $\varphi^\sharp(w)$ represents the minimal $\text{Geom}^\sharp(M)$ gridding of $\pi = \varphi(w)$? If not, then there is a lesser gridding of π , given by some $\varphi^\sharp(v)$. The fact that $\varphi^\sharp(v)$ is less than $\varphi^\sharp(w)$ is witnessed by the position of one or more particular entries which lie in different cells in $\varphi^\sharp(v)$ and $\varphi^\sharp(w)$.

In order to discuss these entries, we use the terminology of marked permutations and marked words. A *marked permutation* is a permutation in which the entries are

allowed to be marked, which we designate with an overline. (Thus our marked permutations look like some other authors' signed permutations, although the marking is meant to convey a completely different concept.) The containment order extends naturally, whereby we make sure that the markings line up. Formally we say that the marked permutation π of length n contains the marked permutation σ of length k if π has a subsequence $\pi(i_1), \pi(i_2), \dots, \pi(i_k)$ such that:

- $\pi(i_1)\pi(i_2)\cdots\pi(i_k)$ is order isomorphic to σ (this is the standard containment order) and
- for each $1 \leq j \leq k$, $\pi(i_j)$ is marked if and only if $\sigma(j)$ is marked.

For example, $\pi = \overline{39}1\overline{86}7\overline{45}2$ contains $\sigma = \overline{5}1\overline{3}42$, as can be seen by considering the subsequence $\overline{9}1\overline{6}72$. A *marked permutation class* is then a set of marked permutations which is closed downward under this containment order.

The mappings φ and $\varphi^\sharp$ can be extended in the obvious manner to order-preserving mappings $\overline{\varphi}$ and $\overline{\varphi}^\sharp$ from $(\Sigma \cup \overline{\Sigma})^*$ to the marked version of $\text{Geom}(M)$, without and with grid lines, respectively, where here $\overline{\Sigma}$ is the *marked cell alphabet* $\{\overline{a} : a \in \Sigma\}$ and both mappings send marked letters to marked entries.

Theorem 8.1. *Every geometrically griddable class is in bijection with a regular language, and thus has a rational generating function.*

Proof. Let $\mathcal{C}$ be a geometrically griddable class. By Proposition 4.2, $\mathcal{C} \subseteq \text{Geom}(M)$ for a partial multiplication matrix M . By Corollary 7.3, there is a regular language, say $L^\sharp$, such that $\varphi^\sharp : L^\sharp \rightarrow \text{Geom}^\sharp(M)$ is a bijection.

We begin by defining a total order on the various griddings of each permutation $\pi \in \text{Geom}(M)$ and thus also on the $\text{Geom}^\sharp(M)$ griddings of permutations in $\mathcal{C}$. Roughly, this order prefers griddings in which the entries of π lie in cells as far to the left and bottom as possible or, in terms of grid lines, the order prefers griddings in which the grid lines lie as far to the right and top as possible. Suppose we have two different $\text{Geom}^\sharp(M)$ griddings of π given by the grid lines G and H . Because these griddings are different, there must be a leftmost vertical grid line or, failing that, a bottommost horizontal grid line, which moved. Note that in the former case, G and H will contain a different number of entries in the corresponding column, while in the latter case they will contain a different number of entries in the corresponding row.

Formally, if (π, G) contains the same number of entries as (π, H) in each of the leftmost $k-1$ columns but contains more entries than (π, H) in column k , then we write $(\pi, G) \sqsubset (\pi, H)$, and we say that column k *witnesses* this fact. Otherwise, if (π, G) and (π, H) contain the same number of entries in each column and in each of the bottom $\ell-1$ rows but (π, G) contains more entries than (π, H) in row ℓ , then $(\pi, G) \sqsubset (\pi, H)$, and we say that row ℓ *witnesses* this fact. Given a permutation $\pi \in \mathcal{C}$, our aim is to choose, from all $\text{Geom}^\sharp(M)$ griddings of π , the minimal one.

Given a triple (π, G, H) where (π, G) and (π, H) are $\text{Geom}^\sharp(M)$ griddings of some $\pi \in \mathcal{C}$ with $(\pi, G) \sqsubseteq (\pi, H)$, we mark the entries of the gridded permutation (π, H) in the following manner. If $(\pi, G) = (\pi, H)$, then no markings are applied. If $(\pi, G) \sqsubset (\pi, H)$ is witnessed by column k , then we mark those entries of π which lie in column k in (π, G) but not in (π, H) . Otherwise, if $(\pi, G) \sqsubset (\pi, H)$ is witnessed by row ℓ , then we similarly mark those entries of π which lie in row ℓ in (π, G) but not in (π, H) .

Let $\overline{\mathcal{C}}^\#$ denote the set of all marked gridded permutations obtained from triples (π, G, H) , $\pi \in \mathcal{C}$, in this manner. Because $\text{Geom}^\#(M)$ griddings are inherited by subpermutations and $\overline{\varphi}^\#$ is order-preserving, the language

$$\overline{\mathcal{J}} = (\overline{\varphi}^\#)^{-1}(\overline{\mathcal{C}}^\#)$$

is subword-closed in $(\Sigma \cup \overline{\Sigma})^*$. In particular, $\overline{\mathcal{J}}$ is a regular language. Loosely speaking, the words in $\overline{\mathcal{J}}$ containing marked letters carry information about all the nonminimal griddings. Our goal is therefore to recognise these nonminimal griddings in Σ^* and remove them from the language $L^\#$.

Consider any word $w \in L^\#$, which encodes the $\text{Geom}^\#(M)$ -gridded permutation $\varphi^\#(w) = (\pi, H)$. The gridding given by H is *not* the minimal $\text{Geom}^\#(M)$ gridding of π precisely if $\overline{\mathcal{J}}$ contains a copy of w with one or more marked letters. Let $\Gamma : (\Sigma \cup \overline{\Sigma})^* \rightarrow \Sigma^*$ denote the homomorphism which removes markings, i.e., the homomorphism given by $a, \overline{a} \mapsto a$. The words which represent nonminimal griddings (precisely the words we wish to remove from $L^\#$) are therefore the set

$$K = \Gamma\left(\overline{\mathcal{J}} \cap \left((\Sigma \cup \overline{\Sigma})^* \setminus \Sigma^*\right)\right).$$

By the basic properties of regular languages, it can be seen that K and hence $L = L^\# \setminus K$ are regular. The proof is then complete, as $\varphi : L \rightarrow \mathcal{C}$ is a bijection. $\square$

9. INDECOMPOSABLE AND SIMPLE PERMUTATIONS

Here we adapt the techniques of the previous section to establish bijections between regular languages and three structurally important subsets of geometrically griddable classes.

An *interval* in the permutation π is a set of contiguous indices $I = [a, b]$ such that the set of values $\pi(I) = \{\pi(i) : i \in I\}$ is also contiguous. Every permutation of length n has trivial intervals of lengths 0, 1, and n ; the permutation π of length at least 2 is said to be *simple* if it has no other intervals. The importance of simple permutations in the study of permutation classes has been recognised since Albert and Atkinson [1], whose terminology we follow. We refer to Brignall [9] for a recent survey.

Given a permutation σ of length m and nonempty permutations $\alpha_1, \dots, \alpha_m$, the *inflation* of σ by $\alpha_1, \dots, \alpha_m$ — denoted $\sigma[\alpha_1, \dots, \alpha_m]$ — is the permutation obtained by replacing each entry $\sigma(i)$ by an interval that is order isomorphic to α_i in such a way that the intervals are order isomorphic to σ . For example,

$$2413[1, 132, 321, 12] = 4\ 798\ 321\ 56.$$

Two particular types of inflations have their own names. These are the *direct sum*, or simply *sum*, $\pi \oplus \sigma = 12[\pi, \sigma]$ and the *skew sum* $\pi \ominus \sigma = 21[\pi, \sigma]$. We say that a permutation is *sum indecomposable* if it is not the sum of two shorter permutations and is *skew indecomposable* if it is not the skew sum of two shorter permutations.

These notions defined, we are ready to construct regular languages which encode the simple, sum indecomposable, and skew indecomposable permutations in a geometrically griddable class. Our approach mirrors the proof of Theorem 8.1.

Theorem 9.1. *The simple, sum indecomposable, and skew indecomposable permutations in every geometrically griddable class are each in bijection with a regular language and thus have rational generating functions.*

Proof. We give complete details only for the case of simple permutations. The minor modifications needed to handle sum indecomposable and skew indecomposable permutations are indicated at the end of the proof.

Let $\mathcal{C}$ be a geometrically griddable class. As usual, Proposition 4.2 implies that $\mathcal{C} \subseteq \text{Geom}(M)$ for a partial multiplication matrix M . By Theorem 8.1, there is a regular language L such that $\varphi : L \rightarrow \mathcal{C}$ is a bijection.

Let $\overline{\mathcal{C}}$ denote the set of all permutations in $\mathcal{C}$ with all possible markings. We say that the markings of a marked permutation are *interval consistent* if the marked entries of the permutation form a (possibly trivial) interval. Let $\overline{\mathcal{I}}$ consist of all marked permutations in $\overline{\mathcal{C}}$ with interval consistent markings. Because $\overline{\varphi}$ is order-preserving, the preimage

$$\overline{\mathcal{J}} = \overline{\varphi}^{-1}(\overline{\mathcal{I}})$$

is subword-closed in $(\Sigma \cup \overline{\Sigma})^*$ and thus is a regular language.

Now consider any permutation $\pi \in \mathcal{C}$. This permutation is simple if and only if it does not have a nontrivial interval. In terms of our markings, therefore, π is simple if and only if there is no interval consistent marking of π which contains at least two marked entries and at least one unmarked entry. On the language level, a word over $(\Sigma \cup \overline{\Sigma})^*$ has at least two marked entries and at least one unmarked entry precisely if it lies in

$$(\Sigma \cup \overline{\Sigma})^* \setminus (\Sigma^* \cup \overline{\Sigma}^* \cup \Sigma^* \overline{\Sigma} \Sigma^*).$$

Therefore the words in Σ^* which represent nonsimple permutations in $\mathcal{C}$ are precisely those in the set

$$K = \Gamma \left(\overline{\mathcal{J}} \cap \left((\Sigma \cup \overline{\Sigma})^* \setminus (\Sigma^* \cup \overline{\Sigma}^* \cup \Sigma^* \overline{\Sigma} \Sigma^*) \right) \right),$$

where, as in the proof of Theorem 8.1, $\Gamma : (\Sigma \cup \overline{\Sigma})^* \rightarrow \Sigma^*$ denotes the homomorphism which removes markings. The simple permutations in $\mathcal{C}$ are therefore encoded by the regular language $L \setminus K$, completing the proof of that case.

This proof can easily be adapted to the case of sum (resp., skew) indecomposable permutations by defining markings to be *sum consistent* (resp., *skew consistent*) if the underlying permutation is the sum (resp., skew sum) of its marked entries and its unmarked entries (in either order). $\square$

10. ATOMIC DECOMPOSITIONS

The intersection of two geometrically griddable classes is trivially geometrically griddable, and as we observed in Proposition 6.3, their union is geometrically griddable as well. Therefore, within the lattice of permutation classes, the collection of geometrically griddable classes forms a sublattice. In this section we consider geometrically griddable classes from a lattice-theoretic viewpoint.

The permutation class $\mathcal{C}$ is *join-irreducible* (in the usual lattice-theoretic sense) if $\mathcal{C} \neq \mathcal{D} \cup \mathcal{E}$ for two proper subclasses $\mathcal{D}, \mathcal{E} \subsetneq \mathcal{C}$. In deference to existing literature on permutation classes, we refer to join-irreducible classes as *atomic*. It is not difficult to show that the *joint embedding property* is a necessary and sufficient condition for

the permutation class $\mathcal{C}$ to be atomic. This condition states that for all $\pi, \sigma \in \mathcal{C}$, there is a $\tau \in \mathcal{C}$ containing both π and σ .

Fraïssé [14] studied atomic classes in the more general context of relational structures and established another necessary and sufficient condition. Specialised to the context of permutations, given two linearly ordered sets A and B and a bijection $f : A \rightarrow B$, every finite subset $\{a_1 < \dots < a_n\} \subseteq A$ maps to a finite sequence $f(a_1), \dots, f(a_n) \in B$ that is order isomorphic to a unique permutation. We call the set of permutations that arise in this manner the *age of f* , denoted $\text{Age}(f : A \rightarrow B)$. A proof of the following result in the language of permutations can also be found in Atkinson, Murphy and Ruškuc [7].

Theorem 10.1 (Fraïssé [14]; see also Hodges [16, Section 7.1]). *The following three conditions are equivalent for a permutation class $\mathcal{C}$:*

- (1) $\mathcal{C}$ is atomic,
- (2) $\mathcal{C}$ satisfies the joint embedding property, and
- (3) $\mathcal{C} = \text{Age}(f : A \rightarrow B)$ for a bijection f between two countable linear orders A and B .

The next proposition is a specialisation of standard lattice-theoretic facts which may be found in more general terms in many sources, such as Birkhoff [8].

Proposition 10.2. *Every pwo permutation class can be expressed as a finite union of atomic classes.*

In order to describe the atomic geometrically griddable classes as $\text{Geom}(M)$ for certain matrices M , we allow our matrices to contain entries equal to $\bullet$ in order to signify cells in which a permutation may contain at most one point. We have to be a bit careful here, as it is unclear how one should interpret $(\bullet \bullet)$. We simply forbid such configurations, in the sense formalised by the following definitions.

Suppose that M is a $0/\bullet/\pm 1$ matrix, meaning that each entry of M lies in $\{0, \bullet, 1, -1\}$. We say that M is $\bullet$ -isolated if every $\bullet$ entry is the only nonzero entry in its column and row. Given a $\bullet$ -isolated $0/\bullet/\pm 1$ matrix M , its *standard figure* is the point set in $\mathbb{R}^2$ consisting of:

- a single point at $(k - 1/2, \ell - 1/2)$ if $M_{k,\ell} = \bullet$,
- the line segment from $(k - 1, \ell - 1)$ to (k, ℓ) if $M_{k,\ell} = 1$, or
- the line segment from $(k - 1, \ell)$ to $(k, \ell - 1)$ if $M_{k,\ell} = -1$.

We can then extend the notion of geometric grid classes to $0/\bullet/\pm 1$ matrices in the obvious manner and obtain the following result.

Theorem 10.3. *The atomic geometrically griddable classes are precisely the geometric grid classes of $\bullet$ -isolated $0/\bullet/\pm 1$ matrices, and every geometrically griddable class can be expressed as a finite union of such classes.*

Proof. First suppose that M is a $\bullet$ -isolated $0/\bullet/\pm 1$ matrix. It is clear from the geometric description of $\text{Geom}(M)$ that given any two permutations $\pi, \sigma \in \text{Geom}(M)$, there is a permutation $\tau \in \text{Geom}(M)$ such that $\tau \geq \pi, \sigma$, so such classes satisfy the joint embedding property and are thus atomic by Theorem 10.1.

Next we show that every geometrically griddable class can be expressed as a finite union of classes of the form $\text{Geom}(M)$ where M is a $\bullet$ -isolated $0/\bullet/\pm 1$ matrix. Note that this will imply that the only atomic geometrically griddable classes are of the latter type.

FIGURE 8. The standard gridded figure of the matrix $\begin{pmatrix} 1 & -1 & 1 \\ -1 & 0 & 0 \end{pmatrix}$ is shown on the left, while the figure on the right displays the subfigure for the subclass encoded by the language $\{a_{11}, a_{12}\}^* \{\varepsilon, a_{22}\} \{a_{11}, a_{12}, a_{32}\}^* \{\varepsilon, a_{22}\} \{a_{12}, a_{32}\}^*$.

Let $\mathcal{C}$ be a geometrically griddable class. By Proposition 4.2, $\mathcal{C} \subseteq \text{Geom}(M)$ for some partial multiplication matrix M (note that M is a $0/\pm 1$ matrix) with cell alphabet Σ . Since the encoding map $\varphi : \Sigma^* \rightarrow \text{Geom}(M)$ is order-preserving (Proposition 5.3), the preimage $\varphi^{-1}(\mathcal{C})$ is a subword-closed language. By Proposition 5.1, we know that $\varphi^{-1}(\mathcal{C})$ is a finite union of languages of the form

$$(\dagger) \quad \Sigma_1^* \{\varepsilon, a_2\} \Sigma_3^* \{\varepsilon, a_4\} \dots \Sigma_{2q-1}^* \{\varepsilon, a_{2q}\} \Sigma_{2q+1}^*,$$

where $q \geq 0$, $\Sigma_1, \dots, \Sigma_{2q+1} \subseteq \Sigma$, and $a_2, \dots, a_{2q} \in \Sigma$.

Let L denote an arbitrary language of the form $(\dagger)$. We will show that $\varphi(L) = \text{Geom}(M_L)$ for some $\bullet$ -isolated $0/\bullet/\pm 1$ matrix M_L , from which the result will follow. We start with the standard gridded figure $\Lambda^\# = (\Lambda, G)$ of $M^{\times(2q+1)}$. Recall that each cell of the standard gridded figure of M becomes $(2q+1)^2$ cells in $\Lambda^\#$ of which $(2q+1)$ are nonempty. We use this to label the nonempty cells of $\Lambda^\#$ by $C_{k,\ell}^{(s)}$ for $s \in [2q+1]$, in order of increasing distance from the base point as it would be in the standard gridded figure of M .

The permutations in $\varphi(L)$ are then equivalent to finite independent sets $P \subseteq \Lambda$ of the following form:

- For odd $s \in [2q+1]$, P may contain any point of Λ belonging to cells $C_{k,\ell}^{(s)}$ for any k, ℓ such that $a_{k,\ell} \in \Sigma_s$, and no points from other cells.
- For even $s \in [2q+1]$, P may contain at most one point from the cell $C_{k,\ell}^{(s)}$, where $a_{k,\ell} = a_s$, and no points from any other cells.

Thus $\varphi(L) = \text{Sub}(\Lambda_L)$, where $\Lambda_L \subseteq \Lambda$ consists of:

- (1) all line segments of Λ in the cells $C_{k,\ell}^{(s)}$ where $s \in [2q+1]$ is odd and $a_{k,\ell} \in \Sigma_s$ and
- (2) the centre point of the subcell $C_{k,\ell}^{(s)}$ where $s \in [2q+1]$ is even and $a_s = a_{k,\ell}$.

Figure 8 shows an example of this construction. The subfigure Λ_L is clearly the standard gridded figure of some $0/\bullet/\pm 1$ matrix M_L . Moreover, as $\bullet$ entries can only arise from case (2) above, it follows that M_L is $\bullet$ -isolated, completing the proof. $\square$

11. CONCLUDING REMARKS

We have provided a comprehensive toolbox of results applicable to geometrically griddable classes, so perhaps the most immediate question is: how can one tell if a permutation class is geometrically griddable? Huczynska and Vatter [19] have shown that a class is contained in a monotone grid class (i.e., it is *griddable*) if and only if it does not contain arbitrarily long sums of 21 or skew sums of 12. However, $\text{Grid}(M) \neq \text{Geom}(M)$ when M is not a forest, so it remains to determine the precise border between griddability and *geometric* griddability.

None of the major proofs in the preceding sections are effective, in as much as they all appeal to the finiteness of certain antichains of words, which follows nonconstructively from Higman's Theorem. Therefore these proofs do not provide algorithms to accomplish any of the following:

- Given a $0/\pm 1$ matrix M , compute the basis of $\text{Geom}(M)$. In particular, any bound on the length of the basis elements would provide such an algorithm.
- Given a $0/\pm 1$ matrix M , compute the generating function for $\text{Geom}(M)$, the simple permutations in $\text{Geom}(M)$, etc.
- Given a $0/\pm 1$ matrix M and a finite set of permutations B , determine the atomic decomposition of $\text{Geom}(M) \cap \text{Av}(B)$ and/or its enumeration.

An intriguing, and somewhat different, question is the membership problem. Given a $0/\pm 1$ matrix M , how efficiently (as a function of n) can one determine if a permutation of length n lies in $\text{Geom}(M)$? Because geometric grid classes are finitely based, this problem is guaranteed to be polynomial-time, but it could conceivably be linear-time. Such a result would extend the parallel between geometric grid classes and subword-closed languages, because the latter (and indeed all regular languages) have linear-time membership problems.

While we believe that geometric grid classes play a special role in the structural theory of permutation classes, their nongeometric counterparts also present many natural questions. Perhaps the most natural is the finite basis question. Does the class $\text{Grid}(M)$ have a finite basis for every $0/\pm 1$ matrix M ? We feel that the answer should be “yes”, but have scant evidence. In his thesis, Waton [28] proves that the grid class

$$\text{Grid} \left(\begin{array}{cc} 1 & 1 \\ 1 & 1 \end{array} \right)$$

is finitely based.

Another example of a finitely based nongeometric grid class appears in one of the earliest papers on permutation patterns, where Stankova [24] proves that the class of permutations which can be expressed as the union of an increasing and a decreasing subsequence, called the *skew-merged permutations*, has the basis $\{2143, 3412\}$. In our notation, this class is

$$\text{Grid} \left(\begin{array}{cc} -1 & 1 \\ 1 & -1 \end{array} \right).$$

The class of skew-merged permutations is also notable because it is the only nongeometric grid class with a known generating function,

$$\frac{1 - 3x}{(1 - 2x)\sqrt{1 - 4x}},$$

due to Atkinson [4]. Could it be the case that all (monotone) grid classes have algebraic generating functions? A first step in this direction might be a more structural derivation of the generating function for skew-merged permutations.

REFERENCES

1. Michael H. Albert and M. D. Atkinson, *Simple permutations and pattern restricted permutations*, Discrete Math. **300** (2005), no. 1-3, 1–15. MR2170110 (2006d:05007)
2. Michael H. Albert, M. D. Atkinson, and Nik Ruškuc, *Regular closed sets of permutations*, Theoret. Comput. Sci. **306** (2003), no. 1-3, 85–100. MR2000167 (2004d:68106)
3. Michael H. Albert, M. D. Atkinson, and Vincent Vatter, *Subclasses of the separable permutations*, Bull. Lond. Math. Soc. **43** (2011), 859–870. MR2854557
4. M. D. Atkinson, *Permutations which are the union of an increasing and a decreasing subsequence*, Electron. J. Combin. **5** (1998), Research paper 6, 13 pp.
5. ———, *Restricted permutations*, Discrete Math. **195** (1999), no. 1-3, 27–38. MR1663866 (99i:05004)
6. M. D. Atkinson, M. M. Murphy, and Nik Ruškuc, *Partially well-ordered closed sets of permutations*, Order **19** (2002), no. 2, 101–113. MR1922032 (2003g:06002)
7. ———, *Pattern avoidance classes and subpermutations*, Electron. J. Combin. **12** (2005), no. 1, Research paper 60, 18 pp. MR2180797 (2006g:05007)
8. Garrett Birkhoff, *Lattice theory*, Third edition. American Mathematical Society Colloquium Publications, Vol. XXV, American Mathematical Society, Providence, R.I., 1967. MR0227053 (37:2638)
9. Robert Brignall, *A survey of simple permutations*, Permutation Patterns (Steve Linton, Nik Ruškuc, and Vincent Vatter, eds.), London Mathematical Society Lecture Note Series, vol. 376, Cambridge University Press, 2010, pp. 41–65. MR2732823
10. ———, *Grid classes and partial well order*, J. Combin. Theory Ser. A **119** (2012), 99–116. MR2844085 (2012k:06004)
11. Volker Diekert, *Combinatorics on traces*, Lecture Notes in Computer Science, vol. 454, Springer-Verlag, Berlin, 1990. MR1075995 (92d:68070)
12. Sergi Elizalde, *The $\mathcal{X}$ -class and almost-increasing permutations*, Ann. Comb. **15** (2011), 51–68. MR2785755
13. Philippe Flajolet and Robert Sedgewick, *Analytic combinatorics*, Cambridge University Press, Cambridge, 2009. MR2483235 (2010h:05005)
14. Roland Fraïssé, *Sur l'extension aux relations de quelques propriétés des ordres*, Ann. Sci. Ecole Norm. Sup. (3) **71** (1954), 363–388. MR0069239 (16,1006b)
15. Graham Higman, *Ordering by divisibility in abstract algebras*, Proc. London Math. Soc. (3) **2** (1952), 326–336. MR0049867 (14:238e)
16. Wilfrid Hodges, *Model theory*, Encyclopedia of Mathematics and its Applications, vol. 42, Cambridge University Press, Cambridge, 1993. MR1221741 (94e:03002)
17. John M. Howie, *Fundamentals of semigroup theory*, London Mathematical Society Monographs. New Series, vol. 12, The Clarendon Press Oxford University Press, New York, 1995, Oxford Science Publications. MR1455373 (98e:20059)
18. Sophie Huczynska and Nik Ruškuc, *Pattern classes of permutations via bijections between linearly ordered sets*, European J. Combin. **29** (2008), no. 1, 118–139. MR2368620 (2008j:05011)
19. Sophie Huczynska and Vincent Vatter, *Grid classes and the Fibonacci dichotomy for restricted permutations*, Electron. J. Combin. **13** (2006), Research paper 54, 14 pp. MR2240760 (2007c:05004)
20. Tomáš Kaiser and Martin Klazar, *On growth rates of closed permutation classes*, Electron. J. Combin. **9** (2003), no. 2, Research paper 10, 20 pp. MR2028280 (2004m:05026)
21. Adam Marcus and Gábor Tardos, *Excluded permutation matrices and the Stanley-Wilf conjecture*, J. Combin. Theory Ser. A **107** (2004), no. 1, 153–160. MR2063960 (2005b:05009)
22. Maximilian M. Murphy and Vincent Vatter, *Profile classes and partial well-order for permutations*, Electron. J. Combin. **9** (2003), no. 2, Research paper 17, 30 pp. MR2028286 (2004i:06004)
23. Imre Simon, *Piecewise testable events*, Automata theory and formal languages (Second GI Conf., Kaiserslautern, 1975) (Berlin), Lecture Notes in Comput. Sci., vol. 33, Springer, 1975, pp. 214–222. MR0427498 (55:530)

- 24. Zvezdelina E. Stankova, *Forbidden subsequences*, Discrete Math. **132** (1994), no. 1-3, 291–316. MR1297387 (96e:05006)
- 25. Vincent Vatter, *Small permutation classes*, Proc. Lond. Math. Soc. (3) **103** (2011), 879–921. MR2852292 (2012j:05022)
- 26. Vincent Vatter and Steve Waton, *On partial well-order for monotone grid classes of permutations*, Order **28** (2011), 193–199. MR2819219
- 27. ———, *On points drawn from a circle*, Electron. J. Combin. **18** (1) (2011), Paper 223, 10 pp. MR2861402 (2012m:05046)
- 28. Steve Waton, *On permutation classes defined by token passing networks, gridding matrices and pictures: Three flavours of involvement*, Ph.D. thesis, Univ. of St Andrews, 2007.

DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF OTAGO, DUNEDIN, NEW ZEALAND
E-mail address: `malbert@cs.otago.ac.nz`

DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF OTAGO, DUNEDIN, NEW ZEALAND
E-mail address: `mike@cs.otago.ac.nz`

CNRS, LABRI, UNIVERSITÉ BORDEAUX 1, BORDEAUX, FRANCE
E-mail address: `mathilde.bouvel@labri.fr`

SCHOOL OF MATHEMATICS AND STATISTICS, UNIVERSITY OF ST. ANDREWS, ST. ANDREWS,
SCOTLAND
E-mail address: `nik@mcs.st-and.ac.uk`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF FLORIDA, GAINESVILLE, FLORIDA 32611-8105
E-mail address: `vatter@ufl.edu`