



Open Research Online

Citation

Brignall, Robert (2012). Grid classes and partial well order. *Journal of Combinatorial Theory, Series A*, 119(1) pp. 99–116.

URL

<https://oro.open.ac.uk/29444/>

License

None Specified

Policy

This document has been downloaded from Open Research Online, The Open University's repository of research publications. This version is being made available in accordance with Open Research Online policies available from [Open Research Online \(ORO\) Policies](#)

Versions

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding

Grid Classes and Partial Well Order

Robert Brignall*

Department of Mathematics and Statistics
The Open University
Milton Keynes, UK
r.brignall@open.ac.uk
<http://users.mct.open.ac.uk/rb8599>

9th August 2011

Abstract

We prove necessary and sufficient conditions on a family of (generalised) gridding matrices to determine when the corresponding permutation classes are partially well-ordered. One direction requires an application of Higman’s Theorem and relies on there being only finitely many simple permutations in the only non-monotone cell of each component of the matrix. The other direction is proved by a more general result that allows the construction of infinite antichains in any grid class of a matrix whose graph has a component containing two or more non-monotone-griddable cells. The construction uses a generalisation of pin sequences to grid classes, together with a number of symmetry operations on the rows and columns of a gridding.

1 Introduction

A partial order is *partially well-ordered* if it contains neither an infinite *antichain* (a set of pairwise incomparable elements) nor an infinite descending chain. In the study of classes of combinatorial structures this latter condition is trivially satisfied, thus such a class is partially well-ordered if and only if it contains no infinite antichain. For many combinatorial structures we have only a quasi-ordering rather than a partial ordering, and in this case we call such a class *well quasi-ordered* when it contains no infinite antichain. Celebrated results affirming well quasi-ordering in different contexts range from Kruskal’s Tree Theorem [12] to the Robertson–Seymour Theorem [16] for minor-closed classes of graphs, but there are many known examples of quasi-orders that are not well quasi-ordered, such as hereditary properties of graphs. Higman’s Theorem (reproduced here in Section 3) is one of the few general tools available to prove that a given quasi-order is well quasi-ordered, but attention has been given more recently to develop a general theory of infinite antichains — see, for example, Gustedt [9] and Cherlin and Latka [7].

In this paper we are concerned with permutations, though there is no particular reason why parts of these results cannot be extended to other structures. A sequence a_1, \dots, a_n of length n of distinct real numbers is said to be *order isomorphic* to another sequence b_1, \dots, b_n if, for

*This research was conducted while the author was being supported by the Heilbronn Institute for Mathematical Research.

all $i, j \in [n] = \{1, 2, \dots, n\}$, $a_i < a_j$ if and only if $b_i < b_j$. In this way every sequence of real numbers of length n is order isomorphic to some permutation π of length n : $a_i < a_j$ if and only if $\pi(i) < \pi(j)$. This order isomorphism induces the *containment* ordering on permutations: we say that a permutation α is *contained* in π , $\alpha \leq \pi$, if there is some subsequence of π order isomorphic to α . Such a subsequence of π is called a *copy* of α in π . Conversely, if π does not contain the permutation β , then π is said to *avoid* β . For example, $\pi = 918572346$ contains 51342 because of the subsequence 91572 ($= \pi(1)\pi(2)\pi(4)\pi(5)\pi(6)$), but avoids 3142.

The containment ordering on permutations defines a partial order on the set of all permutations. A *permutation class* is a set of permutations closed downward in this partial order, i.e. if π is a permutation in the class \mathcal{C} and $\alpha \leq \pi$, then $\alpha \in \mathcal{C}$. These classes have received a lot of attention in recent years, and the question of partial well-order has played a central role: there is a vast library of infinite antichains (see, in particular Murphy’s thesis [14]), while Higman’s Theorem has been applied in the other direction by Atkinson, Murphy and Ruškuc [4] and Albert and Atkinson [1].

The traditional description of a class \mathcal{C} is by the unique antichain B that forms its *basis*: we write $\mathcal{C} = \text{Av}(B)$ to mean $\mathcal{C} = \{\pi : \beta \not\leq \pi \text{ for all } \beta \in B\}$. However, in recent years a new description of permutation classes has arisen, namely “grid classes” of matrices whose entries are themselves permutation classes — for formal definitions see Section 2. These have played a role in the development of the “Fibonacci” and “Vatter” dichotomies [11, 17], providing a complete answer to the possible growth rates¹ of permutation classes below $\kappa \approx 2.20557$, and in particular proving that there are only countably many classes below this growth rate. Grid classes are now being intensely studied in topics ranging from direct enumeration [3] to connections with geometry [2], but of particular relevance to this paper is Murphy and Vatter [15] where grid classes and partial well-order first met, and subsequent work in Waton’s thesis [19], later published in an article with Vatter [18]. In this paper, we will prove the following:

Theorem 1.1. *Let \mathcal{M} be a gridding matrix whose non-empty entries are monotone classes, or non-monotone-griddable classes containing only finitely many simple permutations. Then the permutation class $\text{Grid}(\mathcal{M})$ is partially well-ordered if and only if the graph of \mathcal{M} is a forest, and at most one cell in each component is not monotone.*

The bulk of the work in proving Theorem 1.1 is in showing:

Theorem 1.2. *Let \mathcal{M} be a gridding matrix where every non-empty cell is an infinite permutation class. Then $\text{Grid}(\mathcal{M})$ is not partially well-ordered if \mathcal{M} has a cycle, or a component containing two or more cells that are not monotone griddable.*

After introducing the necessary definitions in Section 2, Section 3 presents Higman’s theorem and completes the proof of the right-to-left direction of Theorem 1.1; the remainder of the paper is devoted to proving Theorem 1.2. In Section 4 we introduce a number of symmetries of griddings which reduces the number of classes that have to be considered. In Section 5 we introduce a family of grid matrices and show that they are the only ones we need to consider, and in Section 6 we show that these classes are not partially well-ordered by constructing antichains that lie in them which satisfy the additional properties required by the symmetry arguments.

¹All permutation classes have an *upper growth rate*, $\overline{\text{gr}}(\mathcal{C}) = \limsup_{n \rightarrow \infty} \sqrt[n]{|\mathcal{C}_n|}$ (see [13]) where \mathcal{C}_n is the set of permutations in \mathcal{C} of length n , but it is still not known in general whether the true *growth rate*, $\lim_{n \rightarrow \infty} \sqrt[n]{|\mathcal{C}_n|}$, exists for all permutation classes.

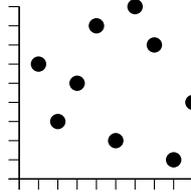


Figure 1: The plot of the permutation $\pi = 635829714$.

2 Definitions

As has become increasingly the case in the study of permutation patterns in recent years, it will prove very useful to view permutations and order isomorphism graphically. Two sets S and T of points in the plane are said to be order isomorphic if we can stretch and shrink the axes for the set S to map the points of S bijectively onto the points of T , i.e. if there are strictly increasing functions $f, g : \mathbb{R} \rightarrow \mathbb{R}$ such that $\{(f(s_1), g(s_2)) : (s_1, s_2) \in S\} = T$. Note that this forms an equivalence relation since the inverse of a strictly increasing function is also strictly increasing. The *plot* of the permutation π is the point set $\{(i, \pi(i))\}$, and every finite point set in the plane in which no two points share a coordinate (often called a *generic* or *noncorectilinear* set) is order isomorphic to the plot of a unique permutation (see Figure 1 for an example). Note that, with a slight abuse of terminology, we will say that a point set is order isomorphic to a permutation.

Inflations and Simple Permutations. An *interval* of a permutation π corresponds to a set of contiguous indices $I = [a, b] = \{a, a + 1, \dots, b\}$ such that the set of values $\pi(I) = \{\pi(i) : i \in I\}$ is also contiguous. For example, $645 = \pi(345)$ is an interval in $\pi = 72645813$.

We form an *inflation* of σ by the permutations τ_1, \dots, τ_k by replacing the entry $\sigma(i)$ with an interval order isomorphic to τ_i , and denote it by $\sigma[\tau_1, \dots, \tau_k]$. For example, $2413[21, 312, 1, 12] = 32867145$. Two special cases of inflations are the *direct sum* $\tau_1 \oplus \tau_2 = 12[\tau_1, \tau_2]$ and the *skew sum* $\tau_1 \ominus \tau_2 = 21[\tau_1, \tau_2]$. A *lenient inflation* is an inflation $\sigma[\tau_1, \dots, \tau_k]$ where we allow one or more of the τ_ℓ to be empty. A class \mathcal{C} is *substitution-closed* (or, in some texts, *wreath-closed*) if $\sigma[\tau_1, \dots, \tau_k] \in \mathcal{C}$ for all $\sigma, \tau_1, \dots, \tau_k \in \mathcal{C}$. The *substitution closure* of a set X is the smallest substitution-closed class containing X , and is denoted $\langle X \rangle$.

A *simple permutation* is a permutation which has no non-trivial intervals, or equivalently a permutation which cannot be expressed as an inflation of some smaller non-singleton permutation. Conversely:

Proposition 2.1 (Albert and Atkinson [1]). *Every permutation except 1 can be expressed as the inflation of a unique simple permutation of length at least 2.*

This proposition shows how simple permutations can be thought of as the “building blocks” of all other permutations, and consequently they play an important role in the study of permutation classes and have received much attention in recent years — see [5] for a survey. We will denote by $\text{Si}(\mathcal{C})$ the set of simple permutations in the class \mathcal{C} . Note that $\text{Si}(\mathcal{C}) = \text{Si}(\langle \mathcal{C} \rangle)$, and also that $\langle \mathcal{C} \rangle = \langle \text{Si}(\mathcal{C}) \rangle$.

Grid Classes. We will present here only a brief survey of the necessary results, and refer the reader to Vatter [17] for a more complete treatment of this topic. To draw a parallel with

the way we view permutations graphically, we will index matrices and grids starting from the bottom-left corner, and with the order of indices swapped. In other words, the ij th entry of a matrix (respectively, ij th cell of a grid) corresponds to the entry (cell) in column i and row j , and an $m \times n$ matrix has m columns and n rows.

An $m \times n$ -gridding of a permutation π is a collection of $m - 1$ distinct vertical and $n - 1$ distinct horizontal lines that divide the plot of π into mn cells. A permutation equipped with a particular $m \times n$ -gridding is called an $m \times n$ -gridded permutation, and for such a gridded permutation π , π^{st} denotes the set of points contained in the st th cell.

Let \mathcal{M} be an $m \times n$ matrix where each entry is a permutation class (noting that we permit the empty class \emptyset): \mathcal{M} is called a *gridding matrix*. (To avoid trivialities, we will always assume that \mathcal{M} does not have any rows or columns consisting entirely of empty cells.) An \mathcal{M} -gridding of a permutation π is an $m \times n$ gridding of π such that π^{st} lies in the class \mathcal{M}_{st} for all $s \in [m]$ and $t \in [n]$. If π possesses an \mathcal{M} -gridding, then π is said to be \mathcal{M} -griddable, and equipping π with such a gridding gives rise to an \mathcal{M} -gridded permutation. Similarly, a permutation class \mathcal{C} is said to be \mathcal{M} -griddable if every $\pi \in \mathcal{C}$ is \mathcal{M} -griddable. The largest permutation class that is \mathcal{M} -griddable (i.e. the class consisting of all \mathcal{M} -griddable permutations) is called the *grid class* of \mathcal{M} , and is denoted $\text{Grid}(\mathcal{M})$. One special case that has received particular attention has been that of *monotone grid classes*, where \mathcal{M} has only monotone (i.e. the classes $\text{Av}(21)$ and $\text{Av}(12)$) or empty entries.

Now let \mathcal{C} and \mathcal{D} be permutation classes. We say that \mathcal{C} is \mathcal{D} -griddable if there is some matrix \mathcal{M} whose entries are all subclasses of \mathcal{D} for which \mathcal{C} is \mathcal{M} -griddable. The following theorem gives a good characterisation of \mathcal{D} -griddability:

Theorem 2.2 (Vatter [17]). *A permutation class \mathcal{C} is \mathcal{D} -griddable if and only if it does not contain arbitrarily long direct sums or skew sums of basis elements of \mathcal{D} .*

A particular instance of this theorem is that a permutation class is monotone griddable if and only if it does not contain arbitrarily long direct sums of 21 or skew sums of 12. Define the *sum completion* of a permutation π to be the permutation class $\oplus\pi = \{\alpha_1 \oplus \alpha_2 \oplus \cdots \oplus \alpha_k : \alpha_i \leq \pi \text{ for all } i \leq k \in \mathbb{N}\}$, and the *skew completion* $\ominus\pi$ analogously. Thus:

Corollary 2.3. *A permutation class \mathcal{C} is monotone griddable if and only if it contains neither the class $\oplus 21$ nor the class $\ominus 12$.*

Grid classes and partial well-order. The *graph* of the gridding matrix \mathcal{M} is the graph $G_{\mathcal{M}}$ whose vertices are the non-empty cells of \mathcal{M} , with two vertices being adjacent if they share a row or a column of \mathcal{M} and all cells between them are empty. A *component* of \mathcal{M} is a submatrix \mathcal{M}' of \mathcal{M} for which $G_{\mathcal{M}'}$ is a connected component of $G_{\mathcal{M}}$. In determining whether grid classes are partially well-ordered, it is sufficient to look at these components individually:

Proposition 2.4 (Vatter [17]). *$\text{Grid}(\mathcal{M})$ is partially well-ordered if and only if $\text{Grid}(\mathcal{M}')$ is partially well-ordered for every connected component \mathcal{M}' of \mathcal{M} .*

In the case of monotone grid classes, the connection between $G_{\mathcal{M}}$ and partial well-order is well known:

Theorem 2.5 (Murphy and Vatter [15]). *The monotone grid class $\text{Grid}(\mathcal{M})$ is partially well-ordered if and only if $G_{\mathcal{M}}$ is a forest.*

One direction of this theorem is proved by constructing an antichain that “winds around” the cells corresponding to a cycle of $G_{\mathcal{M}}$, while the other requires Higman’s Theorem and has been reproved more efficiently by Waton [19]. Our proof of Theorems 1.1 and 1.2 will borrow a lot from the techniques in these two publications.

3 Partially Well Ordered Grid Classes

We complete one half of the proof of Theorem 1.1 by proving the following theorem.

Theorem 3.1. *Let \mathcal{M} be a gridding matrix whose entries are all permutation classes containing only finitely many simple permutations, and for which $G_{\mathcal{M}}$ is a forest and every component of \mathcal{M} contains at most one cell labelled by a class that is not monotone. Then $\text{Grid}(\mathcal{M})$ is partially well-ordered.*

We begin by giving a complete presentation of Higman’s Theorem, which will form the backbone of the proof of Theorem 3.1. We say that (A, M) is an *abstract algebra* if A is a set of elements and M a set of operations where each $\mu \in M$ has an *arity*: we say that μ is a k -ary operation if $\mu : A^k \rightarrow A$ for some positive integer k . Denote the set of k -ary operations in M by M_k , and suppose that M_k is empty for every $k > n$ for some n . (Note that we will allow 0-ary operations.) The abstract algebra (A, M) is said to be *minimal* if no subset B of A allows (B, M) to be an abstract algebra.

A partial order \leq_A on the set of elements A is a *divisibility order* on (A, M) if every operation $\mu \in M_k$, $k = 0, 1, \dots, n$, satisfies

- $a \leq_A b$ implies $\mu(\mathbf{x}, a, \mathbf{y}) \leq_A \mu(\mathbf{x}, b, \mathbf{y})$,
- $a \leq_A \mu(\mathbf{x}, a, \mathbf{y})$,

where \mathbf{x} and \mathbf{y} are arbitrary sequences comprising elements of A whose lengths sum to $k - 1$. Furthermore, given partial orders \leq_{M_k} on M_k , $k = 0, 1, \dots, n$, we say that \leq_A is *compatible* with these partial orders if, for $\lambda, \mu \in M_k$,

- $\lambda \leq_{M_k} \mu$ implies $\lambda(\mathbf{x}) \leq_A \mu(\mathbf{x})$ for all $\mathbf{x} \in A^k$.

Theorem 3.2 (Higman [10]). *Suppose that (A, M) is a minimal abstract algebra for which, for some n , the set M_k of k -ary operations in M is partially well-ordered for each $k = 0, 1, \dots, n$ and empty for $k > n$. Then (A, M) is partially well-ordered under any divisibility ordering compatible with the orders of M_k .*

Applying this to permutation classes, one type of operation that has been particularly amenable to this approach is the inflation of one permutation by others; inflating a permutation σ of length k by τ_1, \dots, τ_k may be thought of as a k -ary operation that acts on the permutations τ_1, \dots, τ_k . It is clear both that inflation is compatible with the permutation containment ordering and that permutation containment is a divisibility ordering with respect to inflations of this type. To satisfy the conditions of Higman’s Theorem, however, we cannot inflate arbitrarily large permutations. Roughly speaking, if a permutation class \mathcal{C} is a subclass of some substitution-closed class \mathcal{D} that can be expressed as the substitution closure of some finite set X , then Higman’s Theorem can be applied to prove that \mathcal{D} (and consequently \mathcal{C}) is partially well-ordered. Consequently, by Proposition 2.1:

Theorem 3.3 (Albert and Atkinson [1]). *Let \mathcal{C} be a class containing only finitely many simple permutations. Then \mathcal{C} is partially well-ordered.*

On the other hand, since any set X satisfying $\langle X \rangle = \langle \mathcal{C} \rangle$ must contain every permutation in $\text{Si}(\mathcal{C})$, we cannot arrange that X is finite when \mathcal{C} contains infinitely many simple permutations, and Higman's Theorem cannot be used in this way. This, however, does not mean that any class containing infinitely many simple permutations is not partially well-ordered: for example, $\text{Grid}(\text{Av}(21) \text{ Av}(21))$ is partially well-ordered by Theorem 2.5, but contains arbitrarily long simple permutations of the form $2\ 4\ 6 \cdots 2k\ 1\ 3\ 5 \cdots 2k - 1$.

Let us now extend this use of Higman's Theorem to gridding matrices. We first define an order on the set of $m \times n$ -gridded permutations. For $m \times n$ -gridded permutations α and π of lengths k and ℓ , respectively, we say that α is contained in π , $\alpha \leq_{mn} \pi$, if and only if there is a sequence of indices $1 \leq i_1 < \cdots < i_k \leq \ell$ such that $\pi(i_1) \cdots \pi(i_k)$ is order isomorphic to α as ungridded permutations, and for $j = 1, \dots, k$, $\pi(i_j)$ and $\alpha(j)$ lie in the same cell in the $m \times n$ -griddings. Similarly, for a specific $m \times n$ gridding matrix \mathcal{M} and \mathcal{M} -gridded permutations α and π , we write $\alpha \leq_{\mathcal{M}} \pi$ to mean $\alpha \leq_{mn} \pi$, but also recognising that both α and π are \mathcal{M} -gridded.

Suppose that \mathcal{M} is a gridding matrix whose graph is acyclic, and every non-empty cell of \mathcal{M} is labelled by a monotone class, except for the uv th cell which is labelled by some arbitrary class \mathcal{D} . Viewing $G_{\mathcal{M}}$ as a tree rooted on the uv th cell, each cell other than the uv th is the *child* of some *parent* cell, i.e. the cell lying directly above it in the rooted tree.

Now let τ_1, \dots, τ_k be \mathcal{M} -gridded permutations each with at least one point in cell uv , and let $\sigma \in \mathcal{D}$ be of length k . The \mathcal{M} -inflation of σ by τ_1, \dots, τ_k is the \mathcal{M} -gridded permutation $\pi = \sigma[\tau_1, \dots, \tau_k]_{\mathcal{M}}$, and is formed by first taking the inflation $\pi^{uv} = \sigma[\tau_1^{uv}, \dots, \tau_k^{uv}]$. For every other non-empty cell st of \mathcal{M} , if $\mathcal{M}_{st} = \text{Av}(21)$ then π^{st} is a (possibly lenient) inflation of $1 \cdots k$ by $\tau_1^{st}, \dots, \tau_k^{st}$ in some order, while if $\mathcal{M}_{st} = \text{Av}(12)$ then π^{st} is a (possibly lenient) inflation of $k \cdots 1$ by $\tau_1^{st}, \dots, \tau_k^{st}$ in some order. In either case, the order in which $\tau_1^{st}, \dots, \tau_k^{st}$ appears in the inflation is defined recursively in terms of its parent cell $\pi^{s't'}$ (where either $s = s'$ or $t = t'$) in the tree $G_{\mathcal{M}}$ rooted on cell uv : if cells st and $s't'$ share a column (i.e. if $s = s'$) then reading from left-to-right the τ_i^{st} appear in the same order as the $\tau_i^{s't'}$. Similarly, if $t = t'$ then reading from bottom-to-top the τ_i^{st} appear in the same order as the $\tau_i^{s't'}$. For each $i \in [k]$, the positions of the points in τ_i^{st} relative to the points in $\tau_i^{s't'}$ are exactly the same as the corresponding points in the \mathcal{M} -gridded permutation τ_i . Finally, for each $i \in [k]$, τ_i^{st} interacts with no other $\tau_j^{s't'}$, $j \neq i$: i.e. all the points in τ_i^{st} are above or below, and to the left or to the right of all points in each $\tau_j^{s't'}$. Note that we must remember the left-to-right and bottom-to-top order of $\tau_1^{st}, \dots, \tau_k^{st}$ in every non-empty cell st of \mathcal{M} even if one or more of the τ_i^{st} contains no points, so that we know the order of the cells for any subsequent descendants. See Figure 2 for an illustration.

A *lenient \mathcal{M} -inflation* of σ by τ_1, \dots, τ_k is defined in exactly the same way, except that we do not stipulate that each τ_i^{uv} be non-empty.

Proof of Theorem 3.1. By Proposition 2.4, we may assume that $G_{\mathcal{M}}$ consists of exactly one component. Thus \mathcal{M} is an $m \times n$ gridding matrix such that $G_{\mathcal{M}}$ is a tree and every non-empty cell of \mathcal{M} is labelled by a monotone class, except for the uv th cell which is labelled by some class \mathcal{D} containing only finitely many simple permutations. We will also assume that \mathcal{D} is substitution closed, as otherwise we may replace it with $\langle \mathcal{D} \rangle$ and prove the result for this larger class.

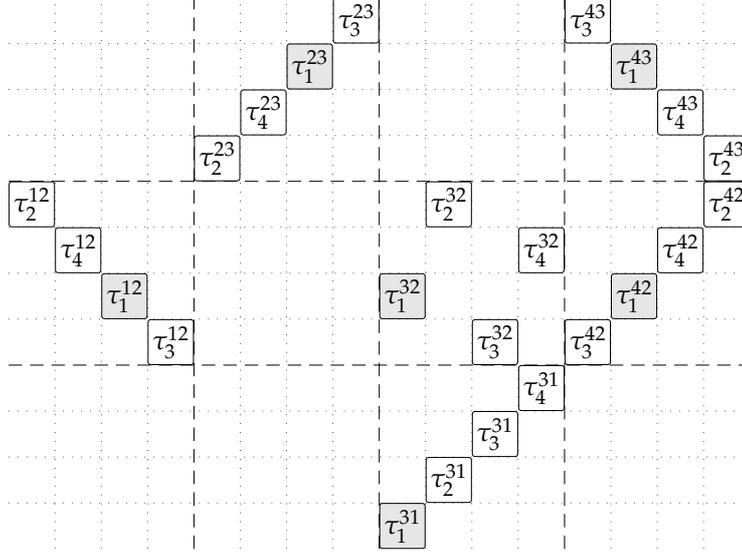


Figure 2: Forming the \mathcal{M} -inflation $2413[\tau_1, \tau_2, \tau_3, \tau_4]_{\mathcal{M}}$ for a 4×3 gridding matrix \mathcal{M} . The highlighted cells correspond to the \mathcal{M} -gridded permutation τ_1 .

For each $\sigma \in \text{Si}(\mathcal{D})$ of length k , we view an \mathcal{M} -inflation of σ as a k -ary operation. We claim that $\text{Grid}(\mathcal{M})$ is generated by this finite list of \mathcal{M} -inflations and all the \mathcal{M} -griddings of the singleton permutation 1. It will then follow by Higman's Theorem 3.2 that $\text{Grid}(\mathcal{M})$ is partially well-ordered. We will prove the result for \mathcal{M} -gridded permutations, and then the result for ungridded permutations in $\text{Grid}(\mathcal{M})$ will follow by applying the homomorphism that removes the gridlines.

We proceed by induction on the length of \mathcal{M} -gridded permutations. As we already have all the \mathcal{M} -gridded permutations of length 1, it is enough to show that any \mathcal{M} -gridded $\pi \in \text{Grid}(\mathcal{M})$ with $|\pi| \geq 2$ can be expressed as an \mathcal{M} -inflation of some $\sigma \in \text{Si}(\mathcal{D})$. Given one such π , suppose first that π^{uv} contains at least two points. By Proposition 2.1 there exists some $\sigma \in \text{Si}(\mathcal{D})$ such that π^{uv} is an inflation of σ , i.e. $\pi^{uv} = \sigma[\tau_1^{uv}, \dots, \tau_k^{uv}]$, for some permutations $\tau_1^{uv}, \dots, \tau_k^{uv}$. Label each point of π^{uv} with the symbol from $1, \dots, k$ corresponding to which of $\tau_1^{uv}, \dots, \tau_k^{uv}$ it belongs. We now label each cell recursively, working down the tree $G_{\mathcal{M}}$ rooted at the cell uv . Consider a cell st whose parent rw has been labelled. We will label each point p in π^{st} as follows:

- If the child shares a column with its parent (i.e. $r = s$), then p is assigned the same label as the rightmost point in π^{rw} that lies to its left. If there is no point in π^{rw} to the left of p , give p the label of the leftmost point of π^{rw} . If there are no points in π^{rw} , label every point of π^{st} with the label 1.
- If the child shares a row with its parent (i.e. $t = w$), then p is assigned the same label as the highest point of π^{rw} that lies below it. If there is no such point in π^{rw} , give p the label of the lowest point of π^{rw} . If there are no points in π^{rw} , label every point of π^{st} with the label 1.

For each $i \in [k]$, now create the \mathcal{M} -gridded permutation τ_i by taking all points of π with label

i. It is now clear to see that $\pi = \sigma[\tau_1, \dots, \tau_k]_{\mathcal{M}}$, as required.

This leaves the case where π^{uv} contains a singleton or is empty. Since $|\pi| \geq 2$, either there is a cell of π containing at least two points, or there are at least two non-empty cells. If there is a cell π^{st} containing at least two points, label the leftmost point with the label 1 and all other points in this cell with label 2. Then view $G_{\mathcal{M}}$ as a tree rooted at the cell st and label the points in the cells of π recursively as described above. Using these labels, now form the gridded permutations τ_1 and τ_2 as before, and observe that π is a lenient \mathcal{M} -inflation of 12 or 21 with τ_1 and τ_2 , in some order.

Finally, if all of the non-empty cells of π contain only one point, then label the point in any one non-empty cell of π with the symbol 1 and the point in any other non-empty cell with the symbol 2. Now assign every other point in every other cell either the label 1 or 2 in such a way that, forming the gridded permutations τ_1 and τ_2 from the labels, π can be expressed as a lenient inflation of 12 or 21 by τ_1 and τ_2 in some order. \square

4 Grid Classes by Symmetry

For the remainder of this paper we will be working towards proving Theorem 1.2 by showing that certain types of grid class are not partially well-ordered. Among these non-partially well-ordered grid classes will be those needed to prove the remaining direction of Theorem 1.1. We begin by showing how we may divide grid classes into families using “grid mappings”.

Let \mathcal{M} be an $m \times n$ gridding matrix, and let π be an \mathcal{M} -gridded permutation. Recall that the *inverse* of a permutation π is π^{-1} , defined by $\pi^{-1}(i) = j$ if and only if $\pi(j) = i$, and we extend this in two ways: first to an \mathcal{M} -gridded permutation π by mapping any vertical line between positions i and $i + 1$ ($i = 0, \dots, n$) to a horizontal line between values i and $i + 1$ and vice versa, and second to a permutation class \mathcal{C} by setting $\mathcal{C}^{-1} = \{\pi^{-1} : \pi \in \mathcal{C}\}$. We consider the effect of taking the inverse of π on the gridding of π , and consequently the effect on \mathcal{M} of taking the inverse of $\text{Grid}(\mathcal{M})$.

Lemma 4.1. *Let \mathcal{M} be an $m \times n$ gridding matrix. Then $\text{Grid}(\mathcal{M})^{-1} = \text{Grid}(\phi(\mathcal{M}))$ where $\phi(\mathcal{M})$ is defined by $(\phi(\mathcal{M}))_{ij} = \mathcal{M}_{ji}^{-1}$.*

We will call the map ϕ the *grid inverse* map.

Proof. First note that $\phi(\phi(\mathcal{M})) = \mathcal{M}$, so it suffices to show that $\text{Grid}(\mathcal{M})^{-1} \subseteq \text{Grid}(\phi(\mathcal{M}))$. Let π be any permutation in $\text{Grid}(\mathcal{M})^{-1}$, so $\pi^{-1} \in \text{Grid}(\mathcal{M})$ is \mathcal{M} -griddable. Pick any \mathcal{M} -gridding of π^{-1} , and apply the inverse operation to this gridded matrix to recover a gridding of π . By definition, all points of the ij th cell of the gridded version of π^{-1} are mapped under inverse to the ji th cell of the gridded π . Moreover, if the points in the ij th cell of π^{-1} form the permutation σ , then it is clear that σ^{-1} is the permutation formed by the points in the ji th cell of π , and so $\pi \in \text{Grid}(\phi(\mathcal{M}))$. \square

Given a permutation π of length k , the *reverse* of π , written $r(\pi)$, is the permutation obtained by reading the entries of π from right to left, i.e. for $i \in [k]$, we have $r(\pi)(i) = \pi(k + 1 - i)$. Similarly, the *complement* of π , denoted $c(\pi)$, is formed by reading the permutation from top to bottom, i.e. $c(\pi)(i) = k + 1 - \pi(i)$. Accordingly, the *reverse* of a set of permutations X is $r(X) = \{r(\pi) : \pi \in X\}$, and the *complement* is $c(X) = \{c(\pi) : \pi \in X\}$. Note in particular that if $\mathcal{C} = \text{Av}(B)$ is a permutation class with basis B then $r(\mathcal{C}) = \text{Av}(r(B))$ and $c(\mathcal{C}) = \text{Av}(c(B))$.

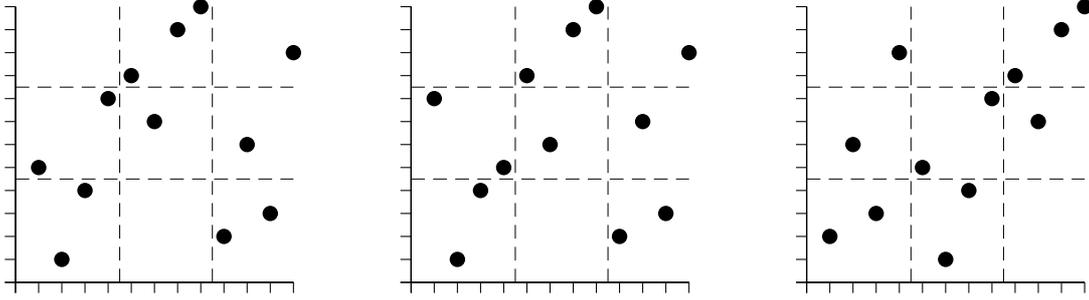


Figure 3: From left to right, the 3×3 gridded permutation $\pi = 5\ 1\ 4\ 8\ 9\ 7\ 11\ 12\ 2\ 6\ 3\ 10$, the 2nd row complement $c_2(\pi) = 8\ 1\ 4\ 5\ 9\ 6\ 11\ 12\ 2\ 7\ 3\ 10$, and the permutation $\mu(\pi) = 2\ 6\ 3\ 10\ 5\ 1\ 4\ 8\ 9\ 7\ 11\ 12$ where $\mu = 312$ is a permutation of the columns.

Now let \mathcal{M} be any $m \times n$ gridding matrix. For fixed $i \in [m]$, let $r_i(\mathcal{M})$ be the i th column reverse of \mathcal{M} , formed by applying the reverse map r to every cell in column i . Thus for all $j \in [n]$, for any $i' \neq i$ we have $(r_i(\mathcal{M}))_{i'j} = \mathcal{M}_{i'j}$, while $(r_i(\mathcal{M}))_{ij} = r(\mathcal{M}_{ij})$. We define the j th row complement analogously: $(c_j(\mathcal{M}))_{i'j} = \mathcal{M}_{i'j}$ whenever $j' \neq j$, and $(c_j(\mathcal{M}))_{ij} = c(\mathcal{M}_{ij})$ for all $i \in [m]$. Next, if μ is a permutation of length m , then let $\mu(\mathcal{M})$ be the gridding matrix formed by permuting the columns of \mathcal{M} as prescribed by μ , so that $(\mu(\mathcal{M}))_{ij} = \mathcal{M}_{\mu(i)j}$. We say that μ is a permutation of the columns of \mathcal{M} . Similarly, a permutation of the rows of \mathcal{M} is a permutation ν of length n satisfying $(\nu(\mathcal{M}))_{ij} = \mathcal{M}_{i\nu(j)}$.

We also extend the definitions of complements, reverses and permutations to gridded permutations in the obvious way. For example, if π is a gridded permutation for which the set of points in row j have values $a, a+1, \dots, b$, then the j th row complement of π is $c_j(\pi)$ defined by $c_j(\pi)(i) = b + a - \pi(i)$ if $(i, \pi(i))$ lies in row j , and $c_j(\pi)(i) = \pi(i)$ otherwise. See Figure 3.

A grid mapping is any composition of grid inverse, row complements, column reverses and row and column permutations, and we say that two matrices \mathcal{M} and \mathcal{N} are equivalent under the grid mapping f if $f(\mathcal{M}) = \mathcal{N}$. (Note that this extends in a natural way to an equivalence relation.) Grid mappings do not in general preserve the normal permutation containment ordering, but they do respect gridded containment (defined in Section 3).

Lemma 4.2. *Let \mathcal{M} be a gridding matrix, α and π two \mathcal{M} -gridded permutations and f any grid mapping of \mathcal{M} . Then $\alpha \leq_{\mathcal{M}} \pi$ if and only if $f(\alpha) \leq_{f(\mathcal{M})} f(\pi)$.*

Proof. It suffices to show that $\alpha \leq_{\mathcal{M}} \pi$ implies $f(\alpha) \leq_{f(\mathcal{M})} f(\pi)$ where f is a grid inverse, row complement, column reverse, or a row or column permutation. We will consider only the grid inverse and column permutation cases, the others following by similar arguments.

Suppose that α is of length k and π of length ℓ , and that the indices $1 \leq i_1 < \dots < i_k \leq \ell$ give rise to a subsequence $\pi(i_1) \dots \pi(i_k)$ that witnesses the gridded containment $\alpha \leq_{\mathcal{M}} \pi$. If $f = \phi$ is the grid inverse mapping then $\alpha \leq_{\mathcal{M}} \pi$ immediately implies $f(\alpha) = \alpha^{-1} \leq \pi^{-1} = f(\pi)$ as this is the normal inverse for permutations. Moreover, if $\pi(i_j)$ and $\alpha(j)$ ($j = 1, \dots, k$) lie in cell st of \mathcal{M} , then their images under f both lie in cell ts of $f(\mathcal{M})$, from which we conclude that $f(\alpha) \leq_{f(\mathcal{M})} f(\pi)$.

Now suppose that f is a column permutation, and note (by composing functions) that we can suppose that f swaps two columns, u and v , say. It is clear that the images of $\pi(i_j)$ and $\alpha(j)$ under f both lie in the same cell, so it remains to show that $f(\alpha) \leq f(\pi)$ as ungridded

permutations. This, however, is also straightforward: f simply swaps the segments of α that lie in columns u and v , and it does likewise in π . In particular, f swaps the two subsequences of $\pi(i_1) \cdots \pi(i_k)$ lying in columns u and v , and this image is a copy of $f(\alpha)$ in $f(\pi)$. \square

We next make a simple observation, which allows us to pass between grid containment and normal permutation containment.

Lemma 4.3. *Let α and π be \mathcal{M} -griddable permutations with $\alpha \leq \pi$ as ungridded permutations. Then for any \mathcal{M} -gridding of π , there exists an \mathcal{M} -gridding of α such that $\alpha \leq_{\mathcal{M}} \pi$.*

Proof. This follows trivially by considering any subsequence of π order isomorphic to α , and then adding the \mathcal{M} -gridding to π , and hence to α . \square

We will use Lemma 4.3 on permutations that have a unique gridding: if α and π are two permutations which have unique \mathcal{M} -griddings for some matrix \mathcal{M} , then $\alpha \not\leq_{\mathcal{M}} \pi$ implies $\alpha \not\leq \pi$. However, unique griddability is not in general preserved by grid mappings. For example, 135246 has a unique gridding in $\text{Grid}(\text{Av}(21) \text{ Av}(21))$, but applying a column reverse to the first column yields the permutation 531246, which can be gridded in two different ways in $\text{Grid}(\text{Av}(12) \text{ Av}(21))$. Thus, for a gridding matrix \mathcal{M} , we say that an \mathcal{M} -gridded permutation π is *strongly uniquely \mathcal{M} -griddable* if the given \mathcal{M} -gridding of π is unique and, for every grid mapping f of \mathcal{M} , $f(\pi)$ is also the unique $f(\mathcal{M})$ -gridding of $f(\pi)$. This extra condition gives us what we need:

Theorem 4.4. *Let \mathcal{M} be a gridding matrix, and let A be an infinite antichain for which infinitely many elements are strongly uniquely \mathcal{M} -griddable. Then the grid class of any gridding matrix \mathcal{N} that is equivalent to \mathcal{M} under some grid mapping is not partially well-ordered.*

Proof. First, we may assume that A consists only of strongly uniquely \mathcal{M} -griddable permutations, as we may discard any elements that are not. Note that $\text{Grid}(\mathcal{M})$ contains A and so is not partially well-ordered. Let f be any grid mapping of \mathcal{M} , and let $\mathcal{N} = f(\mathcal{M})$. Take any pair of distinct permutations $\alpha, \beta \in A$ (noting that $\alpha \not\leq \beta$), and equip each permutation with its unique \mathcal{M} -gridding. With these griddings $f(\alpha)$ and $f(\beta)$ are \mathcal{N} -gridded permutations, and since α and β are strongly uniquely \mathcal{M} -griddable these \mathcal{N} -griddings are the unique griddings of the underlying permutations of $f(\alpha)$ and $f(\beta)$. Now, since $\alpha \not\leq \beta$ we have $\alpha \not\leq_{\mathcal{M}} \beta$, and consequently $f(\alpha) \not\leq_{\mathcal{N}} f(\beta)$ by Lemma 4.2. Additionally, we have $f(\alpha) \not\leq f(\beta)$ as ungridded permutations by Lemma 4.3. Similarly, $\beta \not\leq \alpha$ implies $f(\beta) \not\leq f(\alpha)$, and so $f(\alpha)$ and $f(\beta)$ are incomparable permutations lying in $\text{Grid}(\mathcal{N})$, completing the proof. \square

5 A Family of Grid Matrices

In this section we will reduce the number of gridding matrices that we need to consider to prove Theorem 1.2 to a family in which we can easily build infinite antichains. In agreement with this Theorem, from now on we will now only consider gridding matrices where every non-empty cell is an infinite class, and for which the graph of the matrix either has a cycle, or a component with at least two non-monotone griddable classes. Let \mathcal{M} be such a matrix. First note that if $G_{\mathcal{M}}$ contains a cycle then it contains as a subclass a cyclic monotone grid class, which is not

partially well-ordered by Theorem 2.5. Thus we will assume from now on that $G_{\mathcal{M}}$ is acyclic and has a component with at least two non-monotone griddable entries.

Let \mathcal{F} be the family of all gridding matrices \mathcal{M} satisfying: (i) $G_{\mathcal{M}}$ is a path, (ii) the cells corresponding to the end points of $G_{\mathcal{M}}$ are either $\oplus 21$ or $\ominus 12$, and (iii) the cells corresponding to the internal vertices are monotone. First, we observe that it suffices to show that every grid class of a matrix from this family \mathcal{F} contains an infinite antichain:

Lemma 5.1. *Every grid class $\text{Grid}(\mathcal{M})$, for which $G_{\mathcal{M}}$ has a component with at least two non-monotone-griddable classes, contains $\text{Grid}(\mathcal{M}')$ for some $\mathcal{M}' \in \mathcal{F}$.*

Proof. Pick any two non-monotone-griddable cells lying in the same component of $G_{\mathcal{M}}$, and consider any path in $G_{\mathcal{M}}$ between these two cells. By Corollary 2.3, each of the two selected non-monotone-griddable cells contain $\oplus 21$ or $\ominus 12$, and all other cells on the path contain $\text{Av}(12)$ or $\text{Av}(21)$ (by Erdős and Szekeres [8], since we have assumed these cells are infinite). The Lemma now follows routinely by setting all cells of \mathcal{M} not on this path to be empty (and deleting any resulting empty rows and columns), and replacing the cells on this path with appropriate subclasses. \square

Now that we have the family of matrices \mathcal{F} , we consider the effect of grid mappings on this family. As observed earlier, grid mappings define an equivalence relation, and we wish to find a suitable representative from each class in which to build an antichain that satisfies Theorem 4.4.

Let $\mathcal{C} = \mathcal{D}^+ = \oplus 21$ and $\mathcal{D}^- = \ominus 12$. For $k \in \mathbb{N}$ define \mathcal{M}^k recursively as follows:

- $\mathcal{M}^1 = (\mathcal{C} \ \mathcal{D}^-)$.
- When $k = 4\ell + 1$, \mathcal{M}^k is a $(2\ell + 2) \times (2\ell + 1)$ matrix with $\mathcal{M}_{ij}^k = \mathcal{M}_{ij}^{k-1}$ for all $i \in [1, 2\ell + 1], j \in [2, 2\ell + 1]$; $\mathcal{M}_{(2\ell+2)1}^k = \mathcal{D}^-$; $\mathcal{M}_{11}^k = \text{Av}(21)$; and all other entries are \emptyset .
- When $k = 4\ell + 2$, \mathcal{M}^k is a $(2\ell + 2) \times (2\ell + 2)$ matrix with $\mathcal{M}_{ij}^k = \mathcal{M}_{ij}^{k-1}$ for all $i, j \in [1, 2\ell + 1]$; $\mathcal{M}_{(2\ell+2)(2\ell+2)}^k = \mathcal{D}^+$; $\mathcal{M}_{(2\ell+2)1}^k = \text{Av}(12)$; and all other entries are \emptyset .
- When $k = 4\ell + 3$, \mathcal{M}^k is a $(2\ell + 3) \times (2\ell + 2)$ matrix with $\mathcal{M}_{ij}^k = \mathcal{M}_{(i-1)j}^{k-1}$ for all $i \in [2, 2\ell + 3], j \in [1, 2\ell + 1]$; $\mathcal{M}_{1(2\ell+2)}^k = \mathcal{D}^-$; $\mathcal{M}_{(2\ell+3)(2\ell+2)}^k = \text{Av}(21)$; and all other entries are \emptyset .
- When $k = 4\ell + 4$, \mathcal{M}^k is a $(2\ell + 3) \times (2\ell + 3)$ matrix with $\mathcal{M}_{ij}^k = \mathcal{M}_{i(j-1)}^{k-1}$ for all $i \in [1, 2\ell + 2], j \in [2, 2\ell + 3]$; $\mathcal{M}_{11}^k = \mathcal{D}^+$; $\mathcal{M}_{1(2\ell+3)}^k = \text{Av}(12)$; and all other entries are \emptyset .

Suppressing the labels of empty cells, the first few such matrices are:

$$\begin{aligned}\mathcal{M}^1 &= (\mathcal{C} \ \mathcal{D}^-) \\ \mathcal{M}^2 &= \begin{pmatrix} & \mathcal{D}^+ \\ \mathcal{C} & \text{Av}(12) \end{pmatrix} \\ \mathcal{M}^3 &= \begin{pmatrix} \mathcal{D}^- & & \text{Av}(21) \\ & \mathcal{C} & \text{Av}(12) \end{pmatrix} \\ \mathcal{M}^4 &= \begin{pmatrix} \text{Av}(12) & & \text{Av}(21) \\ & \mathcal{C} & \text{Av}(12) \\ & \mathcal{D}^+ & \end{pmatrix} \\ \mathcal{M}^5 &= \begin{pmatrix} \text{Av}(12) & & \text{Av}(21) & & \\ & \mathcal{C} & \text{Av}(12) & & \\ \text{Av}(21) & & & & \mathcal{D}^- \end{pmatrix}.\end{aligned}$$

Note that $G_{\mathcal{M}^k}$ is a path of length k , one end of which is labelled by \mathcal{C} and the other by either \mathcal{D}^- or \mathcal{D}^+ , and whose internal vertices are labelled by $\text{Av}(21)$ or $\text{Av}(12)$.

Theorem 5.2. *Every gridding matrix $\mathcal{M} \in \mathcal{F}$ is equivalent under some grid mapping to some \mathcal{M}^k .*

Proof. We will form the grid mapping $f : \mathcal{M}^k \rightarrow \mathcal{M}$ in three stages. First, we select k and check whether we need to apply the grid inverse map to \mathcal{M}^k so that it has the same dimensions as \mathcal{M} . Next, we permute the rows and columns of \mathcal{M}^k (or $\phi(\mathcal{M}^k)$) to form an intermediate matrix \mathcal{N} that has empty cells in exactly the same positions as \mathcal{M} . Finally, we use row complements and column reversals on \mathcal{N} to match the non-empty cells to those of \mathcal{M} .

We pick k so that \mathcal{M}^k or its transpose has the same dimensions as \mathcal{M} . There are precisely two cases where we need to apply the grid inverse map ϕ : The first is if \mathcal{M} is not square and it has the same dimensions as the transpose of \mathcal{M} . The second case is where \mathcal{M} is square, but both the first and last edges in the path $G_{\mathcal{M}}$ correspond to pairs of vertices which share a column: when \mathcal{M}^k is square, both the first and last edges in the path $G_{\mathcal{M}^k}$ arise from pairs of vertices which share a row, hence the need for the grid inverse map.

Suppose without loss that we did not need to apply the grid inverse map. We now need to permute the rows and columns of \mathcal{M}^k so that the non-empty cells are moved to the same places as those of \mathcal{M} . This, however, is straightforward: the graphs of the matrices \mathcal{M} and \mathcal{M}^k are the same, so we simply have to apply row and column transpositions to move each non-empty cell of \mathcal{M}^k in turn to the required position. Call the resulting matrix \mathcal{N} .

All that remains is to fix the non-empty cells of \mathcal{N} to have the same labels as \mathcal{M} . This can be done by starting at one end of the path, and applying successive row complements and/or column reverses so that each non-empty cell in turn is labelled correctly. \square

6 Grid Pin Sequences and Antichains

By the results of the previous section, to complete the proof of Theorem 1.2, all we require by Theorems 4.4 and 5.2 is to find a strongly uniquely \mathcal{M}^k -griddable antichain for each $k \in \mathbb{N}$. The reason we chose the matrices \mathcal{M}^k is that they admit antichains that are easily described in terms of “grid pin sequences”. We now define these pin sequences, and prove some elementary

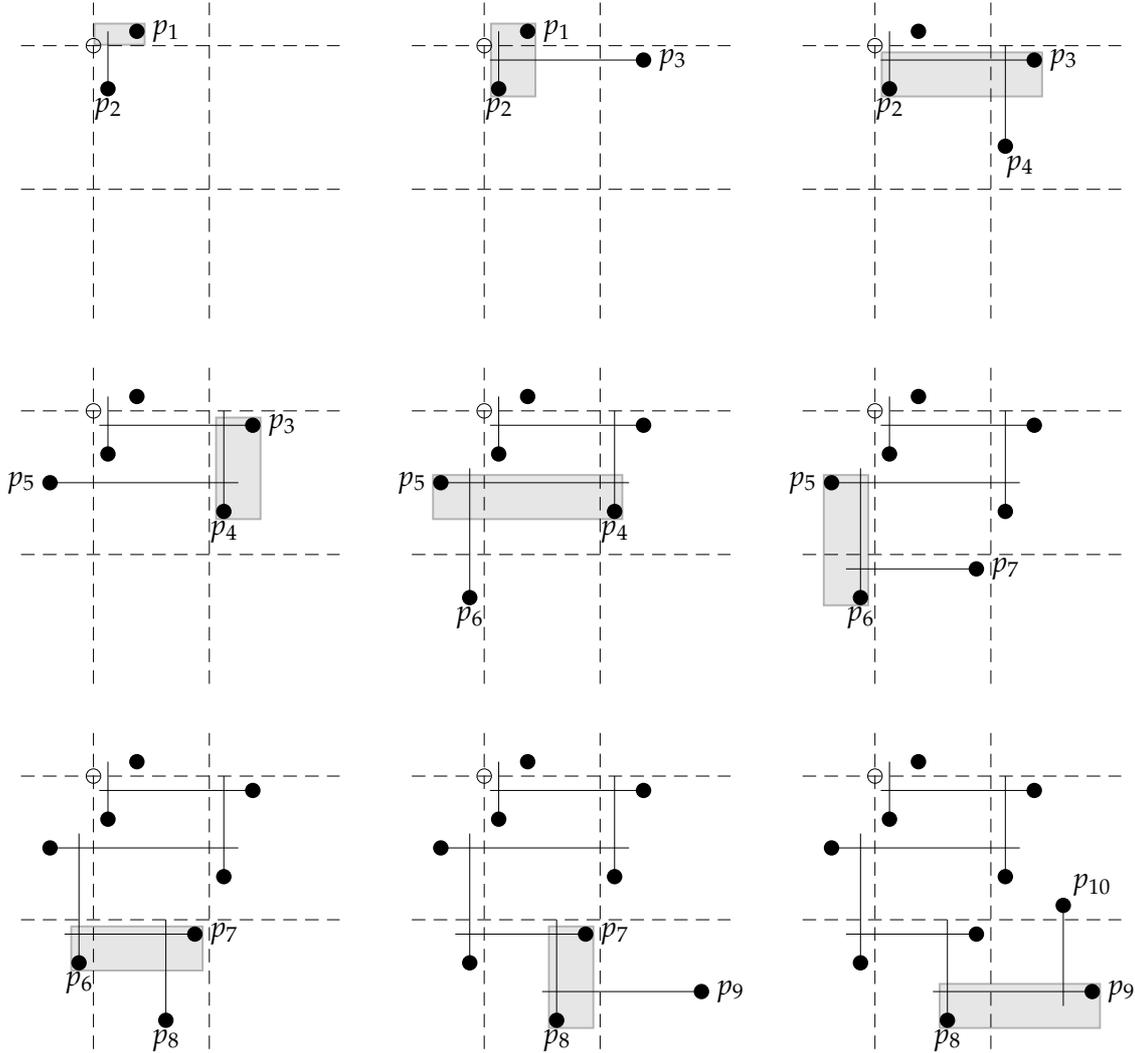


Figure 4: A grid pin sequence on the 3×3 grid.

results about them that should assist in our description of the antichains we wish to construct — it is not our aim here to produce a complete theory of these sequences.

Given points p_1, p_2, \dots in the plane, denote by $\text{rect}(p_1, p_2, \dots)$ the smallest axes-parallel rectangle containing them. A *grid pin sequence* is a sequence of points (called *pins*) p_1, p_2, \dots in an $m \times n$ gridded plane which for $i \geq 2$ must satisfy four conditions:

- *Local separation*: Each pin p_{i+1} separates p_i from p_{i-1} by position or by value.
- *Local externality*: Each pin p_{i+1} lies outside all of $\text{rect}(p_0, p_1), \text{rect}(p_1, p_2), \dots, \text{rect}(p_{i-1}, p_i)$. The *direction* of p_{i+1} denotes its placement relative to $\text{rect}(p_{i-1}, p_i)$: if p_{i+1} lies above (respectively, below, to the left, or to the right) $\text{rect}(p_{i-1}, p_i)$, then p_{i+1} is an up (respectively, down, left, right) pin.
- *Row-column agreement*: If p_{i+1} is an up or a down pin, it must lie in the same column as p_i , while if p_{i+1} is a left or a right pin, it must lie in the same row.

- *Non-interaction*: Each pin p_{i+1} , could not have been used as a grid pin earlier in the pin sequence. I.e. for every $2 \leq j < i$ the pin p_{i+1} must violate one of local separation or row-column agreement with respect to p_j and p_{j-1} (note that it cannot violate local externality).

It still remains to explain how to initiate a grid pin sequence. We begin by placing a fictional pin p_0 corresponding to an *origin* at the intersection of two chosen perpendicular grid lines. Our next pin, p_1 , is then placed in one of the four cells adjacent to this origin and has two directions given by its position relative to p_0 . For example, if p_1 lies below and to the left of p_0 , then p_1 is both a left pin and a down pin. The second pin is then placed to satisfy the four above conditions relative to p_0 and p_1 .

Grid pin sequences should be thought of as a generalisation of proper pin sequences, which were introduced in [6] in a Ramsey-type argument on simple permutations. Proper pin sequences can be recovered from our definition by restricting our view to a 2×2 grid: local separation and row-column agreement combine to form the separation condition of [6], and local externality and non-interaction combine to give the externality condition.

The directions of pins p_2, p_3, \dots must alternate:

Lemma 6.1. *In a grid pin sequence, if p_i ($i > 1$) is a left or a right pin, then p_{i+1} is an up or a down pin. Similarly, if p_i is an up or a down pin then p_{i+1} must be a left or a right pin.*

Proof. Suppose without loss that p_i is a left pin. By local externality and local separation, p_{i+1} must extend from $\text{rect}(p_{i-1}, p_i)$. However, if p_{i+1} is a left or a right pin, then p_{i+1} also either extends from $\text{rect}(p_{i-2}, p_{i-1})$ contradicting non-interaction, or it lies in $\text{rect}(p_{i-2}, p_{i-1})$ contradicting local externality. \square

Lemma 6.2. *If p_{i+1} is a left pin for the grid pin sequence p_1, \dots, p_i , then p_{i+1} lies further left than all previous left pins in its column, and to the right of all previous right pins in its column. Analogous statements hold if p_{i+1} is a right, up or down pin.*

Proof. We prove both statements of the first sentence simultaneously by induction on the total number of left and right pins in a given column. The base case, where there is just a single left or right pin, is trivial. So now suppose for a contradiction to the first statement that p_{i+1} is a left pin which lies to the right of some earlier left pin p_j ($j < i$) in the same column, and assume without loss that there are no other left pins between p_{i+1} and p_j . If p_{j-1} (the predecessor of p_j) lies to the right of p_{i+1} , then p_{i+1} separates p_{j-1} from p_j , is not contained in any of $\text{rect}(p_{j-1}, p_j), \dots, \text{rect}(p_0, p_1)$ and shares a column with p_j , and hence is a pin for p_1, \dots, p_j , contradicting non-interaction. Thus p_{j-1} lies between p_j and p_{i+1} and so lies in the same column. Now consider the pin p_{j-2} , which was a left or a right pin, or p_0 . (Note that p_1 is always either a left or a right pin.) It cannot be a left pin as it would necessarily have to lie by position between p_j and p_{i+1} but by our assumption there are no left pins between p_j and p_{i+1} ; it cannot be a right pin since by row-column agreement it must lie in the same column as p_{j-1} but to the right of p_j , contradicting the inductive hypothesis; finally, it cannot be p_0 as then, in order to lie on an adjacent grid line and ensure that $p_j = p_2$ extends from $\text{rect}(p_0, p_1)$, it must lie to the right of p_{i+1} , but then p_{i+1} either lies in $\text{rect}(p_0, p_1)$ contradicting local externality or it satisfies the conditions to be a pin for p_0, p_1 contradicting non-interaction.

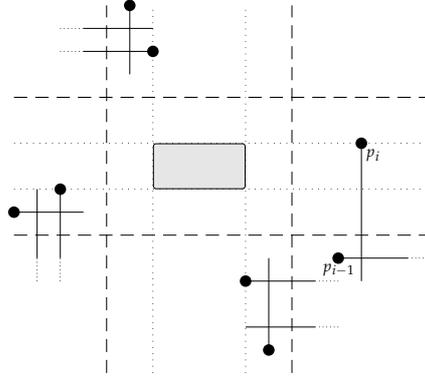


Figure 5: The shaded region denotes the area where pin p_{i+1} can be placed.

A similar argument may be applied to show that p_{i+1} lies to the right of all previous right pins in its column, and so by induction the first sentence of the lemma is true. Finally, symmetry proves the analogous statements in the other three directions. \square

Unlike the 2×2 case, for an arbitrary $m \times n$ grid the direction of the pin is not sufficient to describe the placement of the pin, so we need to be more specific. A *horizontal pin* is either a left or a right pin, while a *vertical pin* is either an up or a down pin.

Lemma 6.3. *Let p_1, p_2, \dots, p_i be a grid pin sequence of length $i \geq 2$ in an $m \times n$ grid. Then if p_{i+1} is a horizontal pin, its placement relative to p_1, \dots, p_i is uniquely determined (up to order isomorphism) by the column in which it lies. Similarly, if p_{i+1} is a vertical pin, its placement relative to p_1, \dots, p_i is uniquely determined by the row in which it lies.*

Proof. We prove only the case where p_{i+1} is a horizontal pin and p_i is an up pin. By row-column agreement, p_{i+1} must be made to lie in the same row as p_i , so coupling this information with the knowledge that p_{i+1} must lie in a specified column is enough to determine the cell into which p_i is placed. In particular, if the column that is to contain p_{i+1} is to the left (respectively, right) of the column containing p_i , then p_{i+1} is a left (resp. right) pin. If p_{i+1} is to lie in the same column as p_i , then the direction of p_{i+1} must match the direction of p_{i-1} to satisfy Lemma 6.3. (Note that if $p_{i-1} = p_1$, then the direction of p_{i+1} matches the horizontal direction of p_1 .)

By Lemma 6.2, p_i must be placed in the region of the cell to the left of all earlier left pins in its column, and to the right of all right pins in the column. This defines a vertical strip extending the length of the column that is devoid of pins. Similarly, p_{i+1} must lie below p_i and above p_{i-1} to satisfy separation, and additionally it must lie above all up pins other than p_i in its row to satisfy non-interaction. This defines a horizontal strip extending to the ends of the row which is devoid of pins.

The intersection of the horizontal strip and the vertical strip defines a rectangular region in the correct cell in which p_{i+1} can be placed — see Figure 5. By its construction, there are no points among p_1, \dots, p_i separating this region, and so all placements of p_{i+1} within this region produce the same permutation up to order isomorphism. \square

Note that the above lemma can be extended to include pin p_2 , but this requires a little further thought. It is not sufficient to state which cell it is to be placed in as there are two different placements of p_2 if it is to lie in the same cell as p_1 : one horizontal, one vertical. However, if the placement of p_2 is specified by a row, then we know p_2 is to be a vertical pin lying in the same column as p_1 , and if specified by a column then p_2 is a horizontal pin lying in the same row as p_1 .

Before we embark on constructing our antichain, we recall the definition of an inflation from Section 2 and extend this to grid pin sequences. (Note that this extension of the definition of an inflation is unrelated to the one used in Section 3.) Letting p_1, \dots, p_n be a grid pin sequence, the *grid pin sequence inflation of p_1, \dots, p_n by the permutations $\alpha_1, \dots, \alpha_n$* is the permutation formed by taking the permutation corresponding to the grid pin sequence p_1, \dots, p_n , and inflating each point p_i ($i = 1, \dots, n$) with the permutation α_i . This is denoted $p_1[\alpha_1], p_2[\alpha_2], \dots, p_n[\alpha_n]$, but whenever $\alpha_i = 1$ we denote the trivially inflated pin $p_i[1]$ simply by p_i . We call such a permutation an *inflated grid pin permutation*.

For each k , we now use inflated grid pin sequences to construct an infinite set of permutations A^k lying in $\text{Grid}(\mathcal{M}^k)$. This construction is accompanied by Figure 6. We begin by showing how to construct the infinite uninflated grid pin sequence p_1, p_2, \dots that will be used to construct all the permutations of A^k : First place the imaginary pin p_0 in the top-right corner of the cell labelled \mathcal{C} lying in the middle of \mathcal{M}^k , and the pin p_1 as a left and down pin (also in the cell labelled by \mathcal{C}). This cell is the only one in its column, but there is one other non-empty cell in the same row, into which we place a right pin p_2 . We then recursively place each pin p_{i+1} so that it does not lie in the same cell as p_i , but shares a row or column with p_i . (Note that by Lemma 6.3, this is a sufficient description, as we know whether p_i was a horizontal or a vertical pin.) Once we have placed our first pin p_j in the cell labelled by \mathcal{D}^+ or \mathcal{D}^- , we place the next pin p_{j+1} in the same cell, and then p_{j+2} is placed in the cell that contained p_{j-1} . Again we place one pin per cell back around until we reach the cell labelled by \mathcal{C} . Once in the cell with label \mathcal{C} , we place a second point in this cell to “turn around”, and repeat.

Finally, $A^k = \{\alpha_1, \alpha_2, \dots\}$, where $\alpha_i = p_1[21], p_2, \dots, p_{(2i-1)k-1}, p_{(2i-1)k}[\beta]$ is an inflated grid pin permutation of length $(2i-1)k+2$, with $\beta = 21$ if k is even and $\beta = 12$ otherwise.

Lemma 6.4. *Every permutation in A^k is \mathcal{M}^k -griddable.*

Proof. This is clear by considering the permitted region in which to place successive pins described in Lemma 6.3. In particular, cells labelled by monotone classes contain only monotone sequences of the right type, and the non-monotone cells contain permutations from $\oplus 21$ or $\ominus 12$ as required. \square

We need an infinite subset of A^k that is both an antichain and strongly uniquely \mathcal{M}^k -griddable. We will in fact find the latter first: knowing the uniqueness of the \mathcal{M}^k -gridding will assist us in proving that elements of A^k are incomparable. The methods of our proofs are similar in flavour to those used by Murphy and Vatter [15]. We begin by making the following straightforward observation, which we will use repeatedly.

Lemma 6.5. *Let $\alpha \in A^k$ be of length $n+2$, α' be the \mathcal{M}^k -gridded permutation of length n corresponding to the uninflated grid pin sequence p_1, \dots, p_n of α , and f be any grid mapping of \mathcal{M}^k . Then, if in $f(\alpha')$ the pins p_i and p_j ($1 < i < j < n$) lie in the same cell and are adjacent by position, they are separated precisely by the pins p_{i-1} and p_{j+1} by value. The same holds swapping “position” and “value”.*

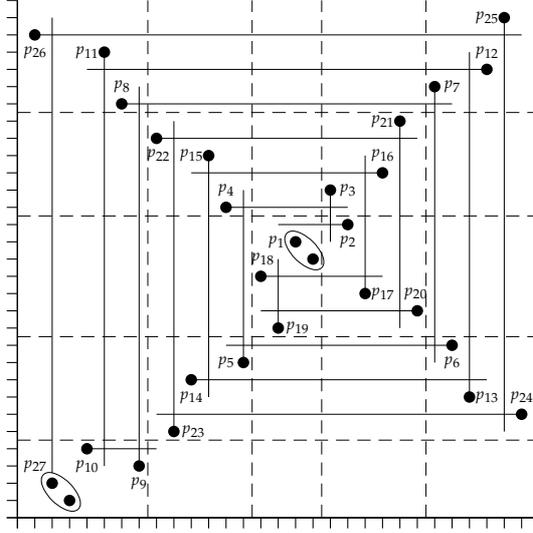


Figure 6: An element of A^8 in the grid class of \mathcal{M}^8 .

Proof. First, by composition of functions it is sufficient to prove the statement when f is the grid inverse map, a row or column permutation, a row complement, a column reversal or the identity map. The effect of the grid inverse mapping ϕ of \mathcal{M}^k on α' is merely to swap the terms “position” and “value” in the statement of the lemma, so we can discount this case. Moreover, every other grid mapping that we need to consider preserves the relative orderings of points by position and value in any given row or column, except possibly to reverse the order. Thus the lemma is true if we can show it is true when f is the identity grid mapping, and this is easily seen by considering the placement of successive pins as described in Lemma 6.3. \square

This lemma is all that is required to prove what we need.

Lemma 6.6. *Every permutation of length at least $2(k+1)^2 + 3$ in A^k is strongly uniquely \mathcal{M}^k -griddable.*

Proof. Given $\alpha \in A^k$ of length $n+2 \geq 2(k+1)^2 + 3$, consider the permutation α' corresponding to the uninflated grid pin sequence of length $n \geq 2(k+1)^2 + 1$ used to create α . We will prove the result for α' , from which the required result easily follows. Label the points of $f(\alpha')$ with p_1, p_2, \dots, p_n according to the grid pin sequence used to construct α' , i.e. the point with label p_i in $f(\alpha')$ is the image under f of the pin p_i . Our approach is to find pairs of points that are adjacent by position or value and lie together in a single cell in both griddings, then apply Lemma 6.5 to find another pair with this property, and repeat the process. Note that in what follows, it does not particularly matter precisely what the map f is: we will simply be exploiting particular features of $f(\mathcal{M}^k)$, for example $G_{f(\mathcal{M}^k)}$ is a path of length k whose internal vertices are labelled by monotone classes and whose end vertices are $\oplus 21$ or $\ominus 12$.

Suppose that $f(\alpha')$ has two $f(\mathcal{M}^k)$ -griddings, the first being the gridding inherited from the original gridding of α' under the grid mapping f , and the second some other gridding. Since $f(\alpha')$ contains at least $2(k+1)^2 + 1$ points, the second gridding has some cell that contains at least $2(k+1) + 1$ points. Of these points, at least three must lie together in some cell in the first gridding, and so we can identify three — p_h, p_i and p_j with $h < i < j$ — which in each gridding

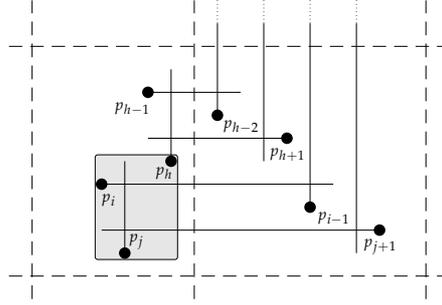


Figure 7: Finding points adjacent by position or by value in the end cell for the proof of Lemma 6.6.

lie in a common cell. (Note that these two common cells do not yet need to correspond to the same cell of $f(\mathcal{M}^k)$.) In each gridding, since $\text{rect}(p_h, p_i, p_j)$ is necessarily contained within the cell, so any other point inside $\text{rect}(p_h, p_i, p_j)$ must lie in the cell. Thus, by shrinking the rectangle and relabelling if necessary, we can assume that $\text{rect}(p_h, p_i, p_j)$ contains only p_h , p_i and p_j . Our first aim is to show that the griddings “line up”: i.e. p_h , p_i and p_j and any other points we find lie in the same cell in both griddings.

We begin by supposing that p_h , p_i and p_j lie in a non-monotone cell in the first gridding. Up to grid mappings, we have a situation such as the one depicted in Figure 7 (note that other arrangements are possible). As p_h , p_i and p_j cannot form a monotone sequence, they also lie in a non-monotone cell in the second gridding. For any such set of three points, observe that one of (p_h, p_i) or (p_i, p_j) are a pair of consecutive pins (i.e. $h = i - 1$ or $i = j - 1$) and hence adjacent by position or by value — without loss we will suppose $i = j - 1$ and that p_i and p_j are adjacent by position. Note that this also excludes the possibility that either of p_i or p_j is p_1 or p_n , and so we may apply Lemma 6.5. This yields a pair (p_{i-1}, p_{j+1}) , which in the first gridding lie in the only other non-empty cell in this row. In the second gridding, (p_{i-1}, p_{j+1}) must also lie together and in a different cell from the one containing p_h , p_i and p_j , to avoid creating a forbidden pattern. We now apply Lemma 6.5 again, this time to the pair (p_{i-1}, p_{j+1}) which are adjacent by value, yielding another pair (p_{i-2}, p_{j+2}) adjacent by position. In both griddings, (p_{i-2}, p_{j+2}) must occupy the only other non-empty cell in the same column as (p_{i-1}, p_{j+1}) . Repeating this process, we follow points around both griddings until we reach the other non-monotone cell. At this point, all the pins we have listed so far are forced to lie in the same cells in both griddings.

In the case where p_h , p_i and p_j lie in a monotone cell in the first gridding, we observe that p_h and p_i are adjacent by position or value, and p_i and p_j are, respectively, adjacent by value or position. We now repeatedly apply Lemma 6.5 to both of these sets of pairs as before. In the first gridding, one pair will produce a sequence that finishes at one of the non-monotone cells, and the other pair will generate a sequence to reach the other. In the second gridding, the positions of all of these points must be the same: first, the cell containing p_h , p_i and p_j is monotone, as otherwise it is the unique cell in its row or column, and so must also contain both points from one of the pairs (p_{h-1}, p_{i+1}) or (p_{i-1}, p_{j+1}) which would produce a forbidden pattern. Thus this cell is monotone, and so the pairs (p_{h-1}, p_{i+1}) and (p_{i-1}, p_{j+1}) must (in some order) lie in the only other non-empty cells in its row and column. A similar argument applies to the cells containing each pair in turn, and hence each of these pairs lies in the same cell in

both griddings.

In either of the above cases for the position of p_h , p_i and p_j , the cells of both griddings now line up, and we have a sequence of pairs of points connecting one non-monotone cell to the other. All that remains is to show how to “turn around” in the non-monotone cells, as we can then follow sequences of pairs of points between the two end cells until we have considered every point. Suppose, therefore, that we have reached a pair of non-consecutive points labelled p_k and p_ℓ with $k < \ell - 1$ in a non-monotone cell. Without loss we assume that p_k and p_ℓ are adjacent by position, so that in both griddings the non-monotone cell is the only non-empty cell in its row, and the previous pair $(p_{k+1}, p_{\ell-1})$ lie together in the only other nonempty cell in its column. Unless $k = 1$ or $\ell = n$ (which we will consider shortly), we can apply Lemma 6.5 again to find pins p_{k-1} and $p_{\ell+1}$ separating p_k and p_ℓ by value. Reading from left to right, they appear in the order $p_{k-1}p_kp_\ell p_{\ell+1}$ or its reverse, so all four points must lie in the same cell in both griddings. We now have two pairs of points which are adjacent by value, namely (p_{k-1}, p_k) and $(p_\ell, p_{\ell+1})$, to which we can again repeatedly apply Lemma 6.5. We follow both of these sequences until we reach the other non-monotone cell, at which point we can again “turn around”. Note, however, that we only need to follow pairs which give rise to pins that have not yet been seen.

Finally, when we encounter a pair (p_k, p_ℓ) in a non-monotone cell with $k = 1$ (or, by a similar argument, $\ell = n$), we find that they are separated by exactly one point, namely $p_{\ell+1}$. Now the pair $(p_\ell, p_{\ell+1})$ is adjacent by position or value, so we can use Lemma 6.5 on it as before. This then generates a sequence which we can follow back and forth until we reach the pin p_n , whence we have covered every point of the sequence.

Thus all points p_1, \dots, p_n must be placed in the same cells in both griddings, as required. The extension to $f(\alpha)$ is trivial: when we encountered pin p_1 or p_n in the above argument, we now encounter two points, both of which have their cell placements forced. \square

Lemma 6.7. *The set of permutations of length at least $2(k+1)^2 + 3$ in A^k is an antichain with respect to permutation containment.*

Proof. Let α, β be two permutations in A^k of lengths m and n respectively, both of length at least $2(k+1)^2 + 3$. Assuming $m < n$, suppose for a contradiction that $\alpha \leq \beta$, and fix one such embedding. Since both α and β have unique \mathcal{M}^k -griddings, this implies not only that $\alpha \leq_{\mathcal{M}^k} \beta$, but that our fixed embedding witnesses this gridded containment. By their construction, we can write α and β as inflated grid pin permutations, thus $\alpha = p_1[21], p_2, \dots, p_m[\gamma]$ and $\beta = q_1[21], q_2, \dots, q_n[\gamma]$, where $\gamma = 12$ or 21 depending on the parity of k . Moreover, since the griddings must match up, the fictive pin p_0 is placed in exactly the same position as q_0 , and so we will assume that p_0 is mapped to q_0 .

We claim that the inflated pin $p_1[21]$ must be mapped to $q_1[21]$. If not, then $p_1[21]$ must be mapped to two consecutive pins in the same cell, this being the only other way to form a 21 pattern in the cell labelled by \mathcal{C} . Thus suppose $p_1[21]$ is mapped to the pins q_{2ki} and q_{2ki+1} . Then the left pin p_2 must be mapped to some left pin with index at most $2kj + 2$, where $j < i$, since all later pins do not separate q_{2ki} from q_0 . Next, by a similar argument, p_3 can be mapped to a pin in β with index at most $2kj + 3$, and so on, until we find that pin p_{2k} (which is the next pin of α we encounter in the cell labelled by \mathcal{C}) must be mapped to a pin with index at most $2kj + 2k \leq 2ki$. This, however, is impossible because p_{2k} must lie below and to the left of $p_1[21]$, but yet in β its image cannot.

Now, since p_0 and q_0 , and $p_1[21]$ and $q_1[21]$ coincide, by the properties of grid pin sequences we conclude that p_2 must be mapped to q_2 , p_3 to q_3 , and so on. This, however, becomes impossible when we try to map $p_m[\gamma]$ into q_m : by non-interaction, there are no pins other than q_m in β that separate $\text{rect}(q_{m-2}, q_{m-1})$, but yet we need two in α to separate $\text{rect}(p_{m-2}, p_{m-1})$. \square

Thus we have:

Proof of Theorem 1.2. First note that if $G_{\mathcal{M}}$ contains a cycle then it contains a non-partially well-ordered class by Theorem 2.5, so we may assume that $G_{\mathcal{M}}$ is acyclic and has a component with at least two non-monotone griddable entries. By Lemma 5.1, it suffices to prove that $\text{Grid}(\mathcal{M})$ is not partially well-ordered for a gridding matrix \mathcal{M} belonging to the family of matrices \mathcal{F} , since every other grid class that needs to be considered contains such a class.

By Lemmas 6.6 and 6.7, A^k contains an infinite antichain of strongly uniquely \mathcal{M}^k -griddable permutations. By Theorem 5.2, \mathcal{M} can be obtained from one of the matrices \mathcal{M}^k for some k via a grid mapping, and so $\text{Grid}(\mathcal{M})$ is not partially well-ordered by Theorem 4.4. \square

The proof of Theorem 1.1 now follows by combining Theorems 1.2 and 3.1.

7 Concluding Remarks

Monotone griddable classes. Theorem 1.1 cannot immediately be extended by replacing “monotone classes” with “monotone gridable classes”. For example, if $\mathcal{M} = (\mathcal{C} \ \mathcal{D})$ where $\mathcal{C} = \begin{pmatrix} \text{Av}(12) \\ \text{Av}(21) \end{pmatrix}$ and $\mathcal{D} = \begin{pmatrix} \text{Av}(21) \\ \text{Av}(12) \end{pmatrix}$, then $\text{Grid}(\mathcal{M})$ contains $\text{Grid} \begin{pmatrix} \text{Av}(12) & \text{Av}(21) \\ \text{Av}(21) & \text{Av}(12) \end{pmatrix}$ which is not partially well-ordered by Theorem 2.5. (Note also that both \mathcal{C} and \mathcal{D} contain only finitely many simple permutations, so adding this restriction would not help.) However, Theorem 1.1 can be used indirectly for such gridding matrices by refining the gridding until all cells are monotone or non-monotone griddable — the details of such a refinement are beyond the scope of this paper, but see Vatter [17] for more details on griddability.

Grid pin sequences and antichains. Currently, every known infinite antichain in the permutation containment order can be built, via grid symmetries, from an infinite grid pin sequence. Note, however, that not every known infinite antichain is constructed simply by inflating the first and last points of a grid pin sequence: see Murphy’s thesis [14] for some other “anchoring” constructions. This naturally leads one to wonder whether there are infinite antichains that cannot be formed in this way. The *closure* of a set A of permutations is the class of permutations contained in the permutations of A , $\text{Cl}(A) = \{\pi : \pi \leq \alpha \text{ for some } \alpha \in A\}$.

Question 7.1. *Does there exist an infinite antichain A for which $\text{Cl}(A)$ does not contain arbitrarily long grid pin sequences?*

Partial well-order decidability. A first step in answering more general questions of decidability could be to consider the following question.

Question 7.2. *Is it decidable whether a given gridding matrix whose entries are partially well-ordered permutation classes defines a grid class that is partially well-ordered or not?*

Theorems 1.1 and 1.2 make some progress towards answering this, particularly in the extension to monotone griddable classes discussed earlier. However, a complete answer would also need to consider gridding matrices where each component is a tree with entries given by monotone classes except for one cell, which is labelled by a non-monotone-griddable class with arbitrarily long simple permutations. This situation is currently amenable neither to Higman’s Theorem nor grid pin sequences.

Acknowledgements. The author wishes to thank Vince Vatter and Michael Albert for helpful discussions, and to the two referees whose comments greatly improved the presentation of these results.

References

- [1] ALBERT, M. H., AND ATKINSON, M. D. Simple permutations and pattern restricted permutations. *Discrete Math.* 300, 1-3 (2005), 1–15.
- [2] ALBERT, M. H., ATKINSON, M. D., BOUVEL, M., RUŠKUC, N., AND VATTER, V. Geometric grid classes of permutations. In preparation.
- [3] ALBERT, M. H., ATKINSON, M. D., AND BRIGNALL, R. The enumeration of permutations avoiding 2143 and 4231. Accepted, *Pure Mathematics and its Applications*.
- [4] ATKINSON, M. D., MURPHY, M. M., AND RUŠKUC, N. Partially well-ordered closed sets of permutations. *Order* 19, 2 (2002), 101–113.
- [5] BRIGNALL, R. A survey of simple permutations. In *Permutation Patterns*, S. Linton, N. Ruškuc, and V. Vatter, Eds., vol. 376 of *London Math. Soc. Lecture Note Ser.* Cambridge Univ. Press, Cambridge, 2010, pp. 41–65.
- [6] BRIGNALL, R., HUCZYNSKA, S., AND VATTER, V. Decomposing simple permutations, with enumerative consequences. *Combinatorica* 28, 4 (2008), 385–400.
- [7] CHERLIN, G. L., AND LATKA, B. J. Minimal antichains in well-founded quasi-orders with an application to tournaments. *J. Combin. Theory Ser. B* 80, 2 (2000), 258–276.
- [8] ERDŐS, P., AND SZEKERES, G. A combinatorial problem in geometry. *Compos. Math.* 2 (1935), 463–470.
- [9] GUSTEDT, J. Finiteness theorems for graphs and posets obtained by compositions. *Order* 15 (1999), 203–220.
- [10] HIGMAN, G. Ordering by divisibility in abstract algebras. *Proc. London Math. Soc.* (3) 2 (1952), 326–336.
- [11] HUCZYNSKA, S., AND VATTER, V. Grid classes and the Fibonacci dichotomy for restricted permutations. *Electron. J. Combin.* 13 (2006), Research paper 54, 14 pp. (electronic).
- [12] KRUSKAL, J. B. Well-quasi-ordering, the Tree Theorem, and Vazsonyi’s conjecture. *Trans. Amer. Math. Soc.* 95 (1960), 210–225.

- [13] MARCUS, A., AND TARDOS, G. Excluded permutation matrices and the Stanley-Wilf conjecture. *J. Combin. Theory Ser. A* 107, 1 (2004), 153–160.
- [14] MURPHY, M. M. *Restricted permutations, antichains, atomic classes, and stack sorting*. PhD thesis, Univ. of St Andrews, 2002.
- [15] MURPHY, M. M., AND VATTER, V. Profile classes and partial well-order for permutations. *Electron. J. Combin.* 9, 2 (2003), Research paper 17, 30 pp. (electronic).
- [16] ROBERTSON, N., AND SEYMOUR, P. D. Graph minors. XX. Wagner’s conjecture. *J. Combin. Theory Ser. B* 92, 2 (2004), 325–357.
- [17] VATTER, V. Small permutation classes. To appear in *Proc. London Math. Soc.*
- [18] VATTER, V., AND WATON, S. On partial well-order for monotone grid classes of permutations. *Order* 28 (2011), 193–199.
- [19] WATON, S. *On Permutation Classes Generated by Token Passing Networks, Gridding Matrices and Pictures: Three Flavours of Involvement*. PhD thesis, Univ. of St Andrews, 2007.