

Profile classes and partial well-order for permutations

Maximillian M. Murphy

School of Mathematics and Statistics
University of St. Andrews
Scotland
`max@mcs.st-and.ac.uk`

Vincent R. Vatter*

Department of Mathematics
Rutgers University
USA
`vatter@math.rutgers.edu`

Submitted: May 3, 2003; Accepted: Oct 13, 2003; Published: Oct 23, 2003

MR Subject Classifications: 06A06, 06A07, 68R15

Keywords: Restricted permutation, forbidden subsequence, partial well-order, well-quasi-order

Abstract

It is known that the set of permutations, under the pattern containment ordering, is not a partial well-order. Characterizing the partially well-ordered closed sets (equivalently: down sets or ideals) in this poset remains a wide-open problem. Given a $0/\pm 1$ matrix M , we define a closed set of permutations called the profile class of M . These sets are generalizations of sets considered by Atkinson, Murphy, and Ruškuc. We show that the profile class of M is partially well-ordered if and only if a related graph is a forest. Related to the antichains we construct to prove one of the directions of this result, we construct exotic fundamental antichains, which lack the periodicity exhibited by all previously known fundamental antichains of permutations.

1 Introduction

It is an old and oft rediscovered fact that there are infinite antichains of permutations with respect to the pattern containment ordering, so the set of all finite permutations is not partially well-ordered. Numerous examples exist including Laver [10], Pratt [13], Tarjan [17], and Spielman and Bóna [16]. In order to show that certain subsets of permutations are partially well-ordered, Atkinson, Murphy, and Ruškuc [3] introduced profile

*Partially supported by an NSF VIGRE grant to the Rutgers University Department of Mathematics.

classes of $0/\pm 1$ vectors (although they gave these classes a different name). We extend their definition to $0/\pm 1$ matrices, give a simple method of determining whether such a profile class is partially well-ordered, and add to the growing library of infinite antichains by producing antichains for those profile classes that are not partially well-ordered. Finally, in Section 5 we generalize our antichain construction to produce exotic fundamental antichains.

The *reduction* of the length k word w of distinct integers is the k -permutation $\text{red}(w)$ obtained by replacing the smallest element of w by 1, the second smallest element by 2, and so on. If $q \in S_k$, we write $|q|$ for the length k of q and we say that the permutation $p \in S_n$ contains a q pattern, written $q \leq p$, if and only if there is a subsequence $1 \leq i_1 < \dots < i_k \leq n$ so that $p(i_1) \dots p(i_k)$ reduces to q . Otherwise we say that p is q -avoiding and write $q \not\leq p$. The problem of enumerating q -avoiding n -permutations has received much attention recently, see Wilf [18] for references.

The relation \leq is a partial order on permutations. Recall that the partially ordered set (X, \leq) is said to be partially well-ordered if it contains neither an infinite properly decreasing sequence nor an infinite antichain (a set of pairwise incomparable elements). Since $|q| < |p|$ whenever $q \leq p$ with $q \neq p$, no set of permutations may contain an infinite properly decreasing sequence, so a set of permutations is partially well-ordered if and only if it does not contain an infinite antichain.

If X is any set of permutations, we let $\mathcal{A}(X)$ denote the set of finite permutations that avoid every member of X . We also let $\text{cl}(X)$ denote the *closure* of X , that is, the set of all permutations p such that there is a $q \in X$ that contains p . We say that the set X is *closed* (or that it is an *order ideal* or a *down-set*) if $\text{cl}(X) = X$. Now that we have the notation, we state another result from Atkinson et al. [3].

Theorem 1.1. [3] *Let p be a permutation. Then $\mathcal{A}(p)$ is partially well-ordered if and only if $p \in \{1, 12, 21, 132, 213, 231, 312\}$.*

We will rely heavily on the result of Higman [8] that the set of finite words over a partially well-ordered set is partially well-ordered under the subsequence ordering. More precisely, if (X, \leq) is a poset, we let X^* denote the set of all finite words with letters from X . Then we say that $a = a_1 \dots a_k$ is a subsequence of $b = b_1 \dots b_n$ (and write $a \leq b$) if there is a subsequence $1 \leq i_1 < \dots < i_k \leq n$ such that $a_j \leq b_{i_j}$ for all $j \in [k]$.

Higman's Theorem. [8] *If (X, \leq) is partially well-ordered then so is (X^*, \leq) .*

Actually, the theorem above is a special case of Higman's result, but it is all that we will need.

If $p \in S_m$ and $p' \in S_n$, we define the direct sum of p and p' , $p \oplus p'$, to be the $(m+n)$ -permutation given by

$$(p \oplus p')(i) = \begin{cases} p(i) & \text{if } 1 \leq i \leq m, \\ p'(i-m) + m & \text{if } m+1 \leq i \leq m+n. \end{cases}$$

The skew sum of p and p' , $p \ominus p'$, is defined by

$$(p \ominus p')(i) = \begin{cases} p(i) + n & \text{if } 1 \leq i \leq m, \\ p'(i-m) & \text{if } m+1 \leq i \leq m+n. \end{cases}$$

Given a set X of permutations, the *sum completion* of X is the set of all permutations of the form $p_1 \oplus p_2 \oplus \dots \oplus p_k$ for some $p_1, p_2, \dots, p_k \in X$, and the *strong completion* of X is set of all permutations that can be obtained from X by a finite number of \oplus and \ominus operations. The following result is given in [3].

Proposition 1.2. [3] *If X is a partially well-ordered set of permutations, then so is the strong completion of X .*

For example, this proposition shows that the set of layered permutations is partially well-ordered, as they are precisely the sum completion of the chain $\{1, 21, 321, \dots\}$. Similarly, the set of separable permutations, the strong completion of the single permutation 1, is partially well-ordered.

2 Profile classes of $0/\pm 1$ matrices

This section is devoted to introducing the central object of our consideration: profile classes. We begin with notation. If M is an $m \times n$ matrix and $(i, j) \in [m] \times [n]$, we denote by $M_{i,j}$ the entry of M in row i and column j . For $I \subseteq [m]$ and $J \subseteq [n]$, we let $M_{I \times J}$ stand for the submatrix $(M_{i,j})_{i \in I, j \in J}$. We write M^t for the transpose of M .

Given a matrix of size $m \times n$, we define the its *support*, $\text{supp}(M)$, to be the set of pairs (i, j) such that $M_{i,j} \neq 0$. The permutation matrix corresponding to $p \in S_n$, M_p , is then the $n \times n$ 0/1 matrix with $\text{supp}(M_p) = \{(i, p(i)) : i \in [n]\}$.

If P and Q are matrices of size $m \times n$ and $r \times s$ respectively, we say that P contains a Q pattern if there is a submatrix P' of P of the same size as Q such that for all $(i, j) \in [r] \times [s]$,

$$Q_{i,j} \neq 0 \text{ implies } P'_{i,j} = Q_{i,j}.$$

(Note that we have implicitly re-indexed the support of P' here.) We write $Q \leq P$ when P contains a Q pattern and $Q \not\leq P$ otherwise. If q and p are permutations then $q \leq p$ if and only if $M_q \leq M_p$. Füredi and Hajnal studied this ordering for 0/1 matrices in [6].

We define the *reduction* of a matrix M to be the matrix $\text{red}(M)$ obtained from M by removing the all-zero columns and rows. Given a set of ordered pairs X let $\Delta(X)$ denote the smallest 0/1 matrix with $\text{supp}(\Delta(X)) = X$. If we are also given a matrix P , let $\Delta^{(P)}(X)$ denote the matrix of the same size as P with $\text{supp}(\Delta^{(P)}(X)) = X$, if such a matrix exists. If Q is a 0/1 matrix satisfying $\text{red}(Q) = Q$ (for instance if Q is a permutation matrix) then Q is contained in a 0/1 matrix P if and only if there is a set $X \subseteq \text{supp}(P)$ with $\text{red}(\Delta(X)) = Q$.

We say that M is a *quasi-permutation matrix* if there is a permutation matrix M' that contains an M pattern or, equivalently, if $\text{red}(M)$ is a permutation matrix. If M is a quasi-permutation matrix and $\text{supp}(M) = \{(i_1, j_1), \dots, (i_\ell, j_\ell)\}$ with $1 \leq i_1 < \dots < i_\ell$, we say that M is *increasing* if $1 \leq j_1 < \dots < j_\ell$ and *decreasing* if $j_1 > \dots > j_\ell \geq 1$. Hence increasing quasi-permutation matrices reduce to permutation matrices of increasing

We have previously defined $\text{Prof}(M)$ to be the set of permutations matrices admitting an M -partition. Now let $\text{Part}(M)$ consist of the triples (P, I, J) where $P \in \text{Prof}(M)$ and (I, J) is an M -partition of P . We let the other set, $\text{SubPart}(M)$, contain all triples (P, I, J) where P is a quasi-permutation matrix and (I, J) is an M -partition of P . Hence $\text{Part}(M) \subseteq \text{SubPart}(M)$.

Suppose that M is an $r \times s$ $0/\pm 1$ matrix with $(P, I, J), (P', I', J') \in \text{SubPart}(M)$ where

$$\begin{aligned} I &= \{i_1 \leq \dots \leq i_{r+1}\}, \\ J &= \{j_1 \leq \dots \leq j_{s+1}\}, \\ I' &= \{i'_1 \leq \dots \leq i'_{r+1}\}, \\ J' &= \{j'_1 \leq \dots \leq j'_{s+1}\}. \end{aligned}$$

We write $(P', I', J') \preceq (P, I, J)$ if there is a set $X \subseteq \text{supp}(P)$ such that $\text{red}(\Delta(X)) = \text{red}(P')$ and for all $k \in [r]$ and $\ell \in [s]$,

$$|X \cap ([i_k, i_{k+1}] \times [j_\ell, j_{\ell+1}])| = |\text{supp}(P') \cap ([i'_k, i'_{k+1}] \times [j'_\ell, j'_{\ell+1}])|.$$

Because $\text{Part}(M) \subseteq \text{SubPart}(M)$, we have also defined \preceq on $\text{Part}(M)$. It is routine to verify that \preceq is a partial order on both of these sets.

The poset we are really interested in, $(\text{Prof}(M), \leq)$, is a homomorphic image of $(\text{Part}(M), \preceq)$. Consequently, if for some M we can show that $(\text{Part}(M), \preceq)$ is partially well-ordered, then we may conclude that $(\text{Prof}(M), \leq)$ is partially well-ordered. This is similar to the approach Atkinson, Murphy, and Ruškuc [3] used to prove Theorem 2.1. First we examine two symmetries of partition classes.

Proposition 3.1. *If M is a $0/\pm 1$ matrix then $(\text{Part}(M^t), \preceq) \cong (\text{Part}(M), \preceq)$.*

Proof: The isomorphism is given by $(P, I, J) \mapsto (P^t, J, I)$. \diamond

Proposition 3.1 says almost nothing more than that for permutations p and q , $q \leq p$ if and only if $\text{inv}(q) \leq \text{inv}(p)$, where here inv denotes the group-theoretic inverse. Similarly, we could define the reverse of a matrix and see that $(\text{Part}(M), \preceq) \cong (\text{Part}(M'), \preceq)$ whenever M and M' lie in the same orbit under the dihedral group of order 4 generated by these two operations. In fact, we have the following more powerful symmetry.

Proposition 3.2. *If M and M' are $0/\pm 1$ matrices and M' can be obtained by permuting the rows and columns of M then $(\text{Part}(M), \preceq) \cong (\text{Part}(M'), \preceq)$.*

Proof: By Proposition 3.1, it suffices to prove this in the case where M' can be obtained by permuting just the rows of M . Furthermore, it suffices to show this claim in the case where M' can be obtained from M by interchanging two adjacent rows k and $k+1$. Let $(P, I = \{i_1 \leq \dots \leq i_{r+1}\}, J = \{j_1 \leq \dots \leq j_{s+1}\}) \in \text{Part}(M)$. Define P' by

$$\begin{aligned} P'_{[1, i_k] \times [n]} &= P_{[1, i_k] \times [n]}, \\ P'_{[i_k, i_k + i_{k+2} - i_{k+1}] \times [n]} &= P_{[i_{k+1}, i_{k+2}] \times [n]}, \\ P'_{[i_k + i_{k+2} - i_{k+1}, i_{k+2}] \times [n]} &= P_{[i_k, i_{k+1}] \times [n]}, \\ P'_{[i_{k+2}, n] \times [n]} &= P_{[i_{k+2}, n] \times [n]}, \end{aligned}$$

and set

$$I' = \{i_1 \leq \dots \leq i_k \leq i_k + i_{k+2} - i_{k+1} \leq i_{k+2} \leq \dots \leq i_{r+1}\}.$$

It is easy to check that $(P, I, J) \mapsto (P', I', J)$ is an isomorphism. \diamond

The analogue of Proposition 3.1 for the poset $(\text{Prof}(M), \leq)$ is true. However, the analogue of Proposition 3.2 fails in general. For example, $\text{Prof} \begin{pmatrix} 1 & 1 & -1 \end{pmatrix}^t$ contains 21 permutations of length four, excluding only 3214, 4213, and 4312, whereas $\text{Prof} \begin{pmatrix} 1 & -1 & 1 \end{pmatrix}^t$ is without 2143, 3142, 3241, 4132, and 4231. Propositions 3.1 and 3.2 suggest (although they fall short of proving) that whether or not $(\text{Part}(M), \preceq)$ is partially well-ordered depends only on the isomorphism class of $G(M)$, this hint was the original motivation for our main result, Theorem 2.2. We are now ready to prove one direction of this theorem.

Theorem 3.3. *Let M be a $0/\pm 1$ matrix. If $G(M)$ is a forest then $(\text{Part}(M), \preceq)$ is partially well-ordered.*

Proof: Let M be an $r \times s$ $0/\pm 1$ matrix satisfying the hypotheses of the theorem. By induction on $|\text{supp}(M)|$ we will construct two maps, μ and ν , such that if $(P, I, J) \in \text{SubPart}(M)$ then

$$\nu(M; P, I, J) = \nu_1(M; P, I, J) \dots \nu_{|\text{supp}(P)|}(M; P, I, J) \in ([r] \times [s])^{|\text{supp}(P)|},$$

and

$$\mu(M; P, I, J) = \mu_1(M; P, I, J) \dots \mu_{|\text{supp}(P)|}(M; P, I, J)$$

is a word containing each element of $\text{supp}(P)$ precisely once, thus specifying an order for us to read through the nonzero entries of P . The other map, ν , will then record which section of P each of these entries lie in. This is formalized in the first of three claims we make about these maps below.

(i) If $\nu_t(M; P, I, J) = (a, b)$ then $\mu_t(M; P, I, J) \in [i_a, i_{a+1}) \times [j_b, j_{b+1})$.

(ii) If $1 \leq a_1 < \dots < a_b \leq |\text{supp}(P)|$ then

$$\begin{aligned} \mu(M; \Delta^{(P)}(\{\mu_{a_1}(M; P, I, J), \dots, \mu_{a_b}(M; P, I, J)\}), I, J) \\ = \mu_{a_1}(M; P, I, J) \dots \mu_{a_b}(M; P, I, J). \end{aligned}$$

(iii) If $(P', I', J') \in \text{SubPart}(M)$ with $\nu(M; P', I', J') = \nu(M; P, I, J)$ then

$$\text{red}(P') = \text{red}(P).$$

First we show that this is enough to prove the theorem. Higman's Theorem tells us that in any infinite set of words from $([r] \times [s])^*$ there are two that are comparable. Hence in every infinite subset of $\text{Part}(M)$, there are elements (P', I', J') and (P, I, J) such that

$\nu(M; P', I', J') \leq \nu(M; P, I, J)$. Hence there are indices $1 \leq a_1 < \dots < a_b \leq |\text{supp}(P)|$ so that

$$\nu(M; P', I', J') = \nu_{a_1}(M; P, I, J) \dots \nu_{a_b}(M; P, I, J).$$

Now let $X = \{\mu_{a_1}(M; P, I, J), \dots, \mu_{a_b}(M; P, I, J)\}$. Claim (ii) implies that

$$\mu(M; \Delta^{(P)}(X), I, J) = \mu_{a_1}(M; P, I, J) \dots \mu_{a_b}(M; P, I, J),$$

and thus by claim (i) we have

$$\begin{aligned} \nu(M; \Delta^{(P)}(X), I, J) &= \nu_{a_1}(M; P, I, J) \dots \nu_{a_b}(M; P, I, J), \\ &= \nu(M; P', I', J'). \end{aligned}$$

Hence claim (iii) shows that

$$\text{red}(\Delta^{(P)}(X)) = \text{red}(P').$$

This implies that $P' \leq P$. The other part of what we need to conclude that $(P', I', J') \preceq (P, I, J)$ comes directly from claim (i). Therefore $\text{Part}(M)$ does not contain an infinite antichain, as desired.

We also need to say a few words about the symmetries of these matrices. Suppose that we have constructed $\mu(M; P, I, J)$, and thus $\nu(M; P, I, J)$, for every $(P, I, J) \in \text{SubPart}(M)$. We would like to claim that this shows how to construct $\mu(M^t; P, I, J)$ for every $(P, I, J) \in \text{SubPart}(M^t)$.

Let $(P, I, J) \in \text{SubPart}(M^t)$, so $(P^t, J, I) \in \text{SubPart}(M)$. We define $\mu(M^t; P, I, J)$ in the natural way by

$$\mu_t(M^t; P, I, J) = (b, a) \text{ if and only if } \mu_t(M; P^t, J, I) = (a, b).$$

Claim (i) then shows us how to define $\nu(M^t; P, I, J)$. Now suppose that $1 \leq a_1 < \dots < a_b \leq |\text{supp}(P)|$ and let $X = \{\mu_{a_1}(M^t; P, I, J), \dots, \mu_{a_b}(M^t; P, I, J)\}$. By definition,

$$\mu_t(M; \Delta^{(P)}(X), I, J) = (b, a)$$

for $t \in [b]$, where

$$(a, b) = \mu_t(M^t; (\Delta^{(P)}(X))^t, J, I),$$

and $(a, b) = \mu_{a_t}(M; P^t, J, I)$ by claim (ii) for M . This shows that $\mu_t(M; \Delta^{(P)}(X), I, J) = \mu_{a_t}(M; P, I, J)$, proving claim (ii). Claim (iii) is easier to prove: if $(P', I', J'), (P, I, J) \in \text{SubPart}(M^t)$ have $\nu(M^t; P', I', J') = \nu(M^t; P, I, J)$ then $\nu(M; (P')^t, J', I') = \nu(M; P^t, J, I)$ so $\text{red}((P')^t) = \text{red}(P^t)$ and thus $\text{red}(P') = \text{red}(P)$.

We would also like to know how to construct $\mu(\overline{M}; P, I, J)$ if \overline{M} is obtained by permuting the rows and columns of M . By our work above, it suffices to show this when \overline{M} can be obtained from M by interchanging rows k and $k+1$. Let $(P, I = \{i_1 \leq \dots \leq i_{r+1}\}, J = \{j_1 \leq \dots \leq j_{s+1}\}) \in \text{SubPart}(\overline{M})$ and define \overline{P} by

$$\begin{aligned} \overline{P}_{[1, i_k] \times [n]} &= P_{[1, i_k] \times [n]}, \\ \overline{P}_{[i_k, i_k + i_{k+2} - i_{k+1}] \times [n]} &= P_{[i_{k+1}, i_{k+2}] \times [n]}, \\ \overline{P}_{[i_k + i_{k+2} - i_{k+1}, i_{k+2}] \times [n]} &= P_{[i_k, i_{k+1}] \times [n]}, \\ \overline{P}_{[i_{k+2}, n] \times [n]} &= P_{[i_{k+2}, n] \times [n]}, \end{aligned}$$

and set

$$\bar{I} = \{i_1 \leq \dots \leq i_k \leq i_k + i_{k+2} - i_{k+1} \leq i_{k+2} \leq \dots \leq i_{r+1}\}.$$

Note that $(\bar{P}, \bar{I}, J) \in \text{SubPart}(M)$, so we can construct $\mu(M; \bar{P}, \bar{I}, J)$. Suppose that $\mu_t(M; \bar{P}, \bar{I}, J) = (a, b)$. We construct $\mu(\bar{M}; P, I, J)$ by

$$\mu_t(\bar{M}; P, I, J) = \begin{cases} (a, b) & \text{if } (a, b) \notin [i_k, i_{k+2}] \times [n], \\ (a + i_{k+2} - i_{k+1}, b) & \text{if } (a, b) \in [i_k, i_{k+1}] \times [n], \\ (a - (i_{k+1} - i_k), b) & \text{if } (a, b) \in [i_{k+1}, i_{k+2}] \times [n]. \end{cases}$$

As usual, claim (i) shows us how to construct $\nu(\bar{M}; P, I, J)$. Checking claims (ii) and (iii) is similar to what we did for the transpose, so we omit it.

We are now ready to begin constructing μ and ν . If $M = 0$, then the only members of $\text{SubPart}(M)$ are triples of the form (P, I, J) where $P = 0$. In this event we set $\nu(M; P, I, J)$ and $\mu(M; P, I, J)$ to the empty word, and claims (i)–(iii) hold quite trivially.

Otherwise $G(M)$ has at least one edge, so it contains a leaf. By our previous work, we may assume that $(r, s) \in \text{supp}(M)$ and $(r, \ell) \notin \text{supp}(M)$ for all $\ell < s$. In other words, the last row of M is identically 0 except in the bottom-right corner, where it contains either a 1 or -1 . Our construction of μ and ν will depend on the operations used to put M into this form but this is of no consequence to us since we have shown that any definition of μ and ν that satisfies (i)–(iii) suffices to prove the theorem.

Let $\bar{M} = M_{[r-1] \times [s]}$. Also, for any $(P, I = \{i_1 \leq \dots \leq i_{r+1}\}, J = \{j_1 \leq \dots \leq j_{s+1}\}) \in \text{Part}(M)$, let $\bar{P} = P_{[1, i_r] \times [1, j_{s+1}]}$ and $\bar{I} = \{i_1 \leq \dots \leq i_r\}$. We have that $(\bar{P}, \bar{I}, J) \in \text{SubPart}(\bar{M})$, and thus by induction we have maps

$$\begin{aligned} \nu(\bar{M}; \bar{P}, \bar{I}, J) &= \nu_1(\bar{M}; \bar{P}, \bar{I}, J) \dots \nu_{|\text{supp}(\bar{P})|}(\bar{M}; \bar{P}, \bar{I}, J) \in ([r-1] \times [s])^{|\text{supp}(\bar{P})|}, \\ \mu(\bar{M}; \bar{P}, \bar{I}, J) &= \mu_1(\bar{M}; \bar{P}, \bar{I}, J) \dots \mu_{|\text{supp}(\bar{P})|}(\bar{M}; \bar{P}, \bar{I}, J), \end{aligned}$$

that satisfy (i), (ii), and (iii).

Now let us build another map, $\mu^{(0)}(M; P, I, J)$ by reading P from left to right. In other words,

$$\mu^{(0)}(M; P, I, J) = \mu_1^{(0)}(M; P, I, J) \dots \mu_{|\text{supp}(P)|}^{(0)}(M; P, I, J),$$

where $\mu_a^{(0)}(M; P, I, J)$ is the element of $\text{supp}(P) - \{\mu_1^{(0)}(M; P, I, J), \dots, \mu_{a-1}^{(0)}(M; P, I, J)\}$ with least second coordinate.

Clearly $\mu^{(0)}(M; P, I, J)$ contains each entry of $\text{supp}(P)$ precisely once. We will now form $\mu(M; P, I, J)$ by rearranging the entries of $\mu^{(0)}(M; P, I, J)$ that also lie in $\text{supp}(\bar{P})$ according to $\mu(\bar{M}; \bar{P}, \bar{I}, J)$. More precisely, suppose that the elements of $\text{supp}(\bar{P})$ appear in positions $1 \leq a_1 < \dots < a_{|\text{supp}(\bar{P})|} \leq \text{supp}(P)$ of $\mu^{(0)}(M; P, I, J)$. Then let

$$\mu(M; P, I, J) = \mu_1(M; P, I, J) \dots \mu_{|\text{supp}(P)|}(M; P, I, J),$$

where

$$\mu_b(M; P, I, J) = \begin{cases} \mu_c(\bar{M}; \bar{P}, \bar{I}, J) & \text{if } b = a_c, \\ \mu_b^{(0)}(M; P, I, J) & \text{otherwise (i.e. when } \mu_b^{(0)}(M; P, I, J) \notin \text{supp}(\bar{P}) \text{)}. \end{cases}$$

By claim (i), this also defines $\nu(M; P, I, J)$. It remains to check that these maps have the desired properties. If $z \in \text{supp}(P)$, we will briefly use the notation $P - z$ to denote the matrix obtained from P by changing the entry at z to 0. To show (ii), it suffices to show that

$$\begin{aligned} \nu(M; P - \mu_a(M; P, I, J), I, J) = \\ \nu_1(M; P, I, J) \dots \nu_{a-1}(M; P, I, J) \nu_{a+1}(M; P, I, J) \dots \nu_{|\text{supp}(P)|}(M; P, I, J). \end{aligned}$$

There are two cases to consider. If $\mu_a(M; P, I, J) \in \text{supp}(\overline{P})$, then let b be such that

$$\mu_a(M; P, I, J) = \mu_b(\overline{M}; \overline{P}, \overline{I}, J),$$

and let c be such that

$$\mu_a(M; P, I, J) = \mu_c^{(0)}(M; P, I, J).$$

Clearly we have

$$\begin{aligned} \mu^{(0)}(M; P - \mu_a(M; P, I, J), I, J) = \\ \mu_1^{(0)}(M; P, I, J) \dots \mu_{c-1}^{(0)}(M; P, I, J) \mu_{c+1}^{(0)}(M; P, I, J) \dots \mu_{|\text{supp}(P)|}^{(0)}(M; P, I, J), \end{aligned}$$

and by induction,

$$\begin{aligned} \mu(\overline{M}; \overline{P} - \overline{\mu_a(M; P, I, J)}, \overline{I}, J) = \\ \mu_1(\overline{M}; \overline{P}, \overline{I}, J) \dots \mu_{b-1}(\overline{M}; \overline{P}, \overline{I}, J) \mu_{b+1}(\overline{M}; \overline{P}, \overline{I}, J) \dots \mu_{|\text{supp}(\overline{P})|}(\overline{M}; \overline{P}, \overline{I}, J). \end{aligned}$$

This implies the claim. The other case, where $\mu_a(M; P, I, J) \notin \text{supp}(\overline{P})$, is easier.

We now have only claim (iii) to show. Suppose to the contrary that (P', I', J') , $(P, I, J) \in \text{SubPart}(M)$ satisfy $\nu(M; P', I', J') = \nu(M; P, I, J)$ but $\text{red}(P') \neq \text{red}(P)$, and choose P' and P with $|\text{supp}(P)| = |\text{supp}(P')|$ minimal subject to this constraint. If (r, s) occurs in neither of these words then we are done because $\overline{P'} = P'$, $\overline{P} = P$, and

$$\nu(\overline{M}; \overline{P'}, \overline{I}', J') = \nu(M; P', I', J') = \nu(M; P, I, J) = \nu(\overline{M}; \overline{P}, \overline{I}, J),$$

so by induction on $|\text{supp}(M)|$, $\text{red}(P') = \text{red}(P)$, a contradiction.

Otherwise (r, s) occurs in both $\nu(M; P', I', J')$ and $\nu(M; P, I, J)$. This is the only part of our proof that depends on the sign of $M_{r,s}$. Since both cases are similar, we will show only the case where $M_{r,s} = 1$. Let a be the position of the last occurrence of (r, s) in $\nu(M; P', I', J')$ and $\nu(M; P, I, J)$, so for all $b > a$, $\nu_b(M; P', I', J') = \nu_b(M; P, I, J) \neq (r, s)$. By our assumptions on M and the construction of μ and ν , we know that of all elements in $\text{supp}(P')$, $\mu_a(M; P', I', J')$ has the greatest first coordinate. We also know the analogous fact for $\mu_a(M; P, I, J)$.

Furthermore, by claims (i) and (ii), we get

$$\nu(M; P' - \mu_a(M; P', I', J'), I', J') = \nu(M; P - \mu_a(M; P, I, J), I, J),$$

so by our choice of P and P' , we have

$$\text{red}(P' - \mu_a(M; P', I', J')) = \text{red}(P - \mu_a(M; P, I, J)).$$

Due to our construction of μ and ν and our choice of a , the entries $\mu_1(M; P', I', J')$, \dots , $\mu_{a-1}(M; P', I', J')$ lie to the upper-left of $\mu_a(M; P', I', J')$, that is, they have lesser first and second coordinates. In addition all of the other entries, $\mu_{a+1}(M; P', I', J')$, \dots , $\mu_{|\text{supp}(P')|}(M; P', I', J')$, lie to the upper-right of $\mu_a(M; P', I', J')$. Completely analogously, we have the same facts for (P, I, J) . This is enough to conclude that $\text{red}(P') = \text{red}(P)$, a contradiction, proving the theorem. \diamond

Theorem 3.3 and Proposition 1.2 together imply the following corollary.

Corollary 3.4. *If M is a finite $0/\pm 1$ matrix and $G(M)$ is a forest then the strong completion of $(\text{Prof}(M), \leq)$ is partially well-ordered.*

4 When profile classes are not partially well-ordered

We have half of Theorem 2.2 left to prove, and its proof will occupy this section. We would like to show that if M is a $0/\pm 1$ matrix for which $G(M)$ is not forest, i.e., it contains a cycle, then $(\text{Prof}(M), \leq)$ contains an infinite antichain. Our construction will generalize the Widderschin antichain introduced in the second section.

First an overview. We will begin by constructing a chain

$$(1) = \overline{P}_1 \leq \overline{P}_2 \leq \dots$$

of permutation matrices, each formed by inserting a new 1 into the previous matrix in a specified manner. Then from \overline{P}_n we will form the $(n+2) \times (n+2)$ permutation matrix P_n by expanding the “first” and “last” entries of \overline{P}_n into appropriate 2×2 matrices. Finally, we will show that there is some constant K depending only on M for which each P_n with $n \geq K$ has a unique M -partition, and from this it will follow that $\{P_n : n \geq K\}$ forms an antichain.

Before we begin, we need to make a technical observation. If $M' \leq M$ then $\text{Prof}(M') \subseteq \text{Prof}(M)$, so we will assume throughout this section that $G(M)$ is precisely a cycle. This requirement is not strictly necessary, but it will simplify the proofs greatly.

Now we are ready to construct \overline{P}_n , which will be an $n \times n$ permutation matrix containing \overline{P}_{n-1} . To the nonzero entries of \overline{P}_n we attach three pieces of information:

- (i) a number; the entry we insert into \overline{P}_{n-1} in order to form \overline{P}_n will receive number n ,
- (ii) a yearn, which must be one of top-left, top-right, bottom-left, or bottom-right, and
- (iii) a nonzero entry of M .

We call the resulting object a *batch*, which will help us keep it separate from the entries of M .

When thinking about these three pieces of information, it might be best to think about starting with an empty matrix partitioned into blocks corresponding to the cells of M . We will insert the batches in the order given by their number. Each batch will be inserted into the block corresponding to the entry of M given by (iii). Within this block, each entry will be placed — with some restrictions — in the corner given by its yearn, so we might say colorfully that each batch yearns toward a corner of its block. This implies that if the entry of M corresponding to a batch is a 1, then the yearn of that batch must be either top-left or bottom-right. Otherwise the yearn must be top-right or bottom-left. Finally, the entries that successive batches correspond to by (iii) will trace out the cycle in $G(M)$.

We have already stated that $\overline{P}_1 = (1)$, but we have not specified properties (ii) and (iii) of batch number 1. We can choose to correspond with the first batch any nonzero entry of M , but for the purpose of being as concise as possible, let us always take it to correspond to the left-most nonzero entry on the first row of M . Such an entry exists because $G(M)$ has been assumed not to have isolated vertices. Upon fixing this entry of M , we have two choices for the yearn of the first batch (although, up to symmetry, the two choices result in the same antichain, see Figure 2 for an example of this). Let us always assume that the first batch yearns right-ward (either bottom-right if the corresponding entry of M is a 1 or top-right if it is a -1).

Having completed the definition of the first batch, we move on to the second. Since $G(M)$ is precisely a cycle, there is a unique nonzero entry of M on the same row as the entry that the first batch corresponds to. We will choose this entry to correspond to the second batch. (We have a choice to take the entry in the same row or the entry in the same column, but again it turns out that these two result in symmetric antichains.) Finally, we specify that the second batch be top-yearning if the first batch was top-yearning and bottom-yearning if the first batch was bottom-yearning. This, together with the sign of the corresponding entry of M , determines the yearn.

Before describing where to insert the second batch into \overline{P}_1 to form \overline{P}_2 , let us define the other batches. The n th batch will correspond to a nonzero cell of M that shares either a row or column with the $n - 1$ st batch, but is not the same cell that either the $n - 1$ st batch or the $n - 2$ nd batch correspond to. Such an entry exists because $G(M)$ is an even cycle (since $G(M)$ is bipartite for any M). If the n th batch shares a row with the $n - 1$ st batch then the n th batch will have the same vertical yearning as the $n - 1$ st batch, that is, it will be top-yearning if the $n - 1$ st batch is top-yearning, and it will be bottom-yearning if the $n - 1$ st batch is bottom-yearning. If the n th batch shares a column with the $n - 1$ st batch, then the two must share the same horizontal yearning. Together with the sign of the corresponding entry of M , this determines the yearn of the n th batch.

Now suppose that we have \overline{P}_{n-1} and want to insert the n th batch. Suppose that this batch corresponds to the cell $(i, j) \in \text{supp}(M)$. Then our first requirements are that the batch must be inserted

- (1) below all batches corresponding to matrix entries (x, y) with $x < i$,

- (2) above all batches corresponding to matrix entries (x, y) with $x > i$,
- (3) to the right of all batches corresponding to matrix entries (x, y) with $y < j$, and
- (4) to the left of all batches corresponding to matrix entries (x, y) with $y > j$.

These four restrictions are enough to insure that the n th batch ends up in the desired “block” of \overline{P}_n . Now we need to insure that it ends up in the correct position within this block. To this end we place the n th batch as far towards its yearning as possible subject to (1)-(4) and one additional condition. The n th and $n - 1$ st batches share either a column or a row, and due to this they must also share either their horizontal or vertical yearning, respectively. The additional condition is simply that the n th batch must not overtake the $n - 1$ st batch in this yearning.

For example, suppose that the n th batch has top-left yearn and that the n th and $n - 1$ st batches share a row, so the $n - 1$ st batch also yearns to be high. Then the n th batch must be placed below $n - 1$ st batch, but otherwise, subject to (1)-(4), as high and far to the left as possible.

Once we have constructed \overline{P}_n , we form P_n by replacing the first and last batches by $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ if that batch corresponds to an 1 in M and by $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ if that batch corresponds to a -1 in M .

Before beginning the proof that the P matrices form an antichain we do a small example, constructing $\overline{P}_1, \overline{P}_2, \dots, \overline{P}_6$ for the matrix

$$M = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}.$$

Let us take the first batch to correspond to entry $(1, 1)$ of M and to have bottom-right yearn. As for any M , we have

$$\overline{P}_1 = (\mathbf{1}).$$

The second batch then corresponds to entry $(1, 2)$ of M . Since the first and second batches share a row, the second batch must be bottom-yearning, and thus its yearn must be bottom-left because $M_{1,2} = -1$. To place the second batch into \overline{P}_2 , we note that conditions (1)-(4) simply state that the second batch must be placed to the right of the first batch. The other requirements insist that the second batch be placed above the first batch, so we end up with

$$\overline{P}_2 = \begin{pmatrix} & \mathbf{1} \\ 1 & \end{pmatrix}.$$

(Here we have made the second batch bold and, as usual, suppressed the 0s.)

The third batch then corresponds to entry $(2, 2)$ of M . It must yearn leftward because it shares a column with the second batch, and since $M_{2,2} = 1$, this means that its yearn must be top-left. Conditions (1)-(4) imply that the third batch must be below the first

and second batches and to the right of the first batch. The other requirements imply that the third batch must be to the right of the second batch. Hence we have

$$\overline{P}_3 = \begin{pmatrix} & 1 & \\ 1 & & \\ & & \mathbf{1} \end{pmatrix}.$$

The fourth batch then has top-right yearn, and must lie below all the previous batches, to the right of the first batch, and to the left of the second and third batches, so

$$\overline{P}_4 = \begin{pmatrix} & & 1 & \\ 1 & & & \\ & & & 1 \\ & \mathbf{1} & & \end{pmatrix}.$$

The fifth batch, like the first batch, corresponds to entry $(1, 1)$ of M . It has the same yearn as the first batch, bottom-right, and must be to the left of batches 2, 3, and 4, above batches 3 and 4, but otherwise as far down and to the right as possible. We then have

$$\overline{P}_5 = \begin{pmatrix} & & & 1 & \\ 1 & & & & \\ & \mathbf{1} & & & \\ & & & & 1 \\ & & 1 & & \end{pmatrix}.$$

Finally, the sixth batch corresponds to entry $(1, 2)$ of M and has bottom-left yearn, like the second batch. When this batch is inserted into \overline{P}_5 we get

$$\overline{P}_6 = \begin{pmatrix} & & & 1 & \\ 1 & & & & \\ & & \mathbf{1} & & \\ & 1 & & & \\ & & & & 1 \\ & & 1 & & \end{pmatrix}.$$

To get P_6 , we replace the first batch with the 2×2 identity matrix and the last batch with the 2×2 anti-identity matrix, resulting in

$$P_6 = \begin{pmatrix} & & & & 1 & \\ 1 & & & & & \\ & 1 & & & & \\ & & & 1 & & \\ & & & & 1 & \\ & & 1 & & & \\ & & & & & 1 \end{pmatrix}.$$

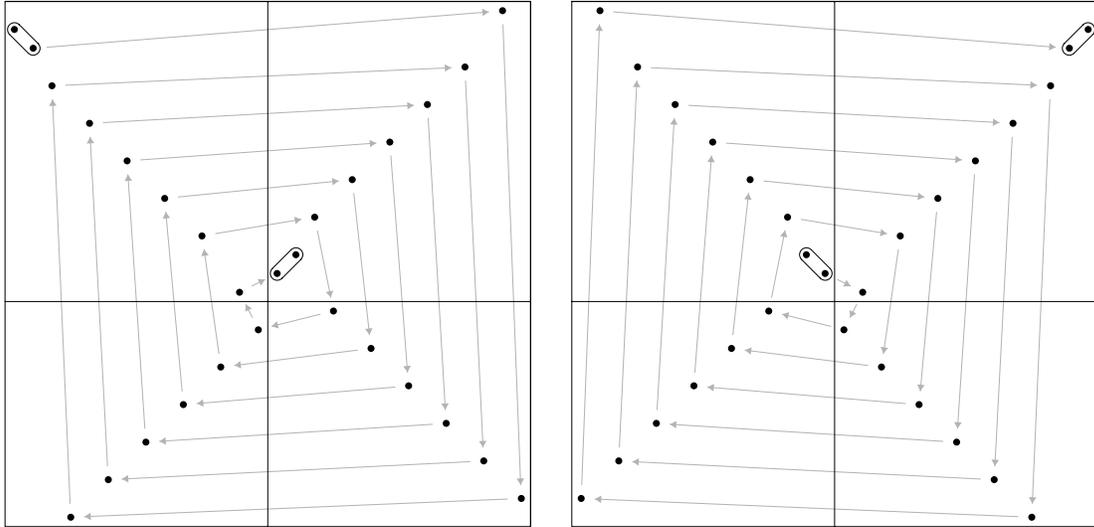


Figure 2: On the left we have a typical element of the Widderschin antichain initialized with bottom-right yearn, on the right it is initialized with top-left yearn. Note that in this case the construction spins inward when initialized as on the left and outward when initialized as on the right, but the two resulting permutations are the same, up to symmetry.

The matrix P_{26} is shown on the left of Figure 2. In the figure we have replaced 1s by dots and drawn an arrow from each batch to the subsequent batch. It should be clear from that picture that we have constructed an antichain almost identical to the Widderschin antichain; the subset $\{P_9, P_{13}, P_{17}, P_{21}, \dots\}$ is – up to symmetry – exactly the Widderschin antichain as we presented it in Section 2.

The matrix shown on the right of Figure 2 shows what we would get had we taken the yearn of the first batch to be top-left instead of bottom-right, and provides an example of our claim that the resulting matrices would be the same, up to symmetry.

Clearly the algorithm is well-defined since we have restricted ourselves to considering only cases where $G(M)$ is a cycle. Almost as clearly, notice that successive batches cycle around $\text{supp}(M)$. Put more precisely, suppose that $G(M)$ is a cycle of length c . Then the m th batch in \overline{P}_n corresponds to the same entry of M that the $(m+c)$ th batch corresponds to.

For the rest of our analysis we will restrict M further, and assume that M contains an even number of -1 s. That this can be done without loss is not completely obvious. Suppose that M has an odd number of -1 s. We form a new matrix M' by replacing the 1s in M by $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, the -1 s by $\begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix}$, and the 0s by the 2×2 zero matrix. It is easy to see that the profile classes of M and M' are identical, but we also need that $G(M')$ contains a unique cycle.

Proposition 4.1. *Let M be a $0/\pm 1$ matrix with an odd number of -1 entries, suppose that $G(M)$ is a cycle, and form M' as described above. Then $G(M')$ contains a unique*

cycle, twice the length of the cycle in $G(M)$.

Proof: First note that there is a natural homomorphism from M' to M that arises by identifying the 2×2 blocks of M' with the cells of M that they came from. This and the fact that every node of $G(M')$ has degree 2 imply that $G(M')$ is either a cycle of twice the length of the cycle in $G(M)$ or two disjoint copies of the cycle in $G(M)$. It is this latter case that we would like to rule out.

Now suppose that M' is $r \times s$ and consider the $r \times s$ matrix S given by $S_{i,j} = (-1)^{i+j}$. For $r = s = 4$, we have

$$S = \begin{pmatrix} 1 & -1 & 1 & -1 \\ -1 & 1 & -1 & 1 \\ 1 & -1 & 1 & -1 \\ -1 & 1 & -1 & 1 \end{pmatrix}.$$

We can form M' by changing entries of S into 0s. A nice property of S is that the length of any walk in S in which diagonal steps are prohibited can be computed modulo 2 by multiplying the entries of the start and end of the walk. If this product is 1 then then length of the walk is even, and if the product is -1 then the length of the walk is odd. Clearly M' also has this property for walks that begin and end at non-zero entries.

Now suppose that $G(M')$ contains two disjoint cycles and choose one of these cycles. Clearly this cycle must contain one entry from each nonzero 2×2 block of M , so the product of the matrix entries used in the cycle is -1 , since M had an odd number of -1 s. Now pair up the entries that lie on the same row of M' . For any such pair, their product tells us whether the distance between them is odd or even. Multiplying all such products together tells us whether the sum of all horizontal distances traversed by the cycle is odd or even. Clearly, this result must be even. However, this product will be the product of all entries used in the cycle, which we have assumed is -1 . Therefore this case is impossible, and the proposition is true. \diamond

Hence we may assume that M contains an even number of -1 s. This assumption will simplify our proofs, but it is worth noting that M and M' give rise to the same antichains; for example, see Figure 3.

Under this assumption we can conclude that any two batches that correspond to the same entry of M share the same yearn. Let us consider the vertical yearn only. Note that it changes only when two successive batches correspond to cells of M in the same column but with opposite signs. Now if we sum over the columns of M the number of -1 entries in each column we must get an even number. Each time two successive batches correspond to cells of M in the same column with the same sign, we either add 0 or 2 to our sum. In the case that these batches correspond to cells of opposite sign, a 1 is contributed. Therefore the number of times that the vertical yearn changes during an entire cycle through $\text{supp}(M)$ must be even. The horizontal case is completely symmetric.

Because of this, one might think of set of batches that correspond to an entry of M as ever more successfully progressing in the direction of their common yearn.

We now prove the major technical lemma we will need to establish that our algorithm produces antichains.

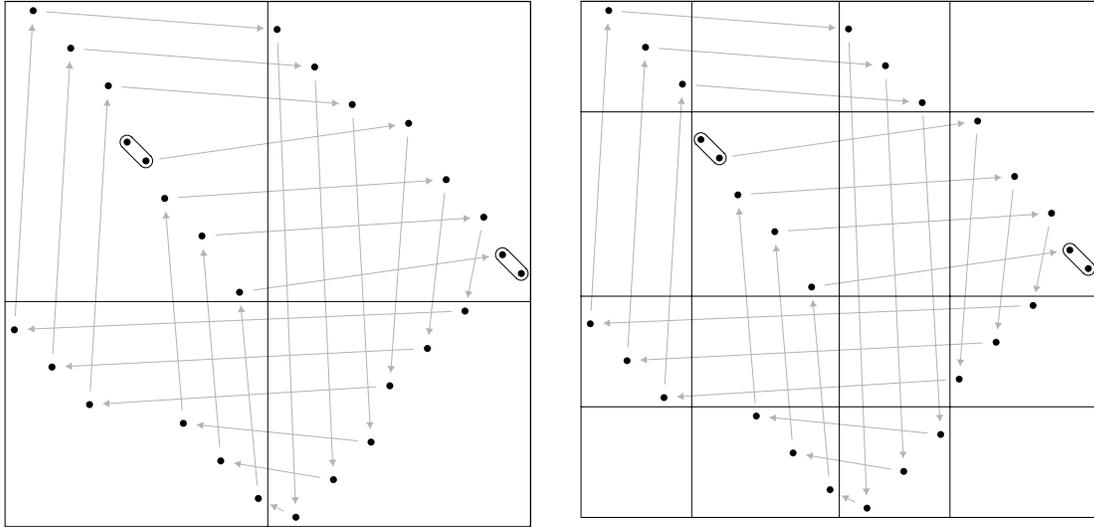


Figure 3: The permutation on the left is a permutation constructed by our algorithm to lie in $\text{Prof} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$. On the right, we have a permutation constructed by our algorithm to lie in the profile class of

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

Unlike our previous examples, in this case the first batch was taken to correspond to the matrix entry $(2, 2)$. Notice that the two permutations are the same.

Lemma 4.2. *Suppose that M is a $0/\pm 1$ matrix with an even number of -1 s and that $G(M)$ is a cycle of length c . Then for $n \geq (c + 1)c^2 + 1$, P_n has a unique M -partition.*

Proof: We begin by proving that under these hypotheses \overline{P}_n has a unique M -partition, from which the claim for P_n will follow rather easily. First note that there is at least one M -partition of \overline{P}_n . This M -partition comes from the correspondence between the batches of \overline{P}_n and the entries of M . Now suppose that we have another M -partition of \overline{P}_n . We will use the verb “allocate” to differentiate this partition from the naturally arising partition just mentioned. So, each batch of \overline{P}_n corresponds to an entry of M (this correspondence coming from the construction) and is allocated to an entry of M (this allocation coming from the other M -partition we have been given).

Since M has precisely c nonzero entries and $n \geq (c + 1)c^2 + 1$, we can find at least $c + 2$ batches that correspond to the same cell of M and are allocated to the same cell of M (note that at this point we cannot assume that these two cells are the same). We will call this set of batches isotope 1. Because they are allocated to the same cell of M , the terms of isotope 1 form either an increasing or a decreasing sequence. Suppose that the lowest numbered batch in isotope 1 is batch number x and that the highest numbered

batch is batch number $x + tc$ (so by our assumptions about the cardinality of isotope 1, $t \geq c + 1$). It should be clear from the principle that successive batches that correspond to the same entry of M are increasingly successful in attaining their common yearn that isotope 1 also contains the batches $x, x + c, x + 2c, \dots, x + tc$.

We know that all the batches in isotope 1 share the same yearn. Without loss, we will assume that they are all top-yearning and that batch $x + 1$ corresponds to an entry of M on the same row as the entry that the batches of isotope 1 correspond to. Then batch $x + 1$ is lower than batch x , batch $x + c + 1$ is higher than batch x but lower than batch $x + c$, and in general batch $x + rc + 1$ is lower than batch $x + rc$ but higher than batch $x + (r - 1)c$ for all $r \in [t - 1]$. So, the batches numbered $x + rc + 1$ for $r \in [t - 1]$ horizontally separate the isotope 1 batches. Therefore these batches must all be allocated to a nonzero entry of M on the same row as the entry that the isotope 1 batches are allocated to. However, these two isotopes must not be allocated to the same entry of M on account of their non-monotonicity, and thus because $G(M)$ is a cycle, there is only one entry of M that the batches numbered $x + rc + 1$, $r \in [t - 1]$, may be assigned to. Let isotope 2 denote the set of all batches that correspond to and are allocated to the same entries of M as these batches.

We now proceed in this manner, defining isotopes numbered 3 through $c + 1$, each either vertically or horizontally separating the last. In general isotope i will contain at least $c + 2 - i$ batches. Suppose that isotope $i - 1$ contains all the batches numbered $x + rc + i - 1$ where $r \in [s, t]$. Then, excepting the $i = 2$ case in which the existence of batch $x + tc + 1$ is uncertain, isotope i will contain the terms $x + rc + i - 1$ where $r \in [s + 1, t]$. Thus we can be guaranteed that isotope $c + 1$, the last isotope we will construct, is non-empty.

These isotopes must cycle around M , so the batches of isotope $c + 1$ are allocated to the same entry of M as the batches of isotope 1. The isotopes must also contain a sequence of successive batches. Furthermore, it is possible to determine the relative vertical and horizontal placement of the cells to which the isotopes are allocated, from which it follows that the batches in the isotopes are allocated to the cells to which they correspond.

It remains only to consider the batches that do not lie in the isotopes. Note that if $c + 1$ consecutive batches are allocated to the cells they correspond to, then also the batches immediately succeeding and preceding this sequence, if they exist, are allocated to the cells they correspond to. For proof, suppose that such a sequence $y, y + 1, \dots, y + c$ of batches is given and that, without loss, batches y and $y + 1$ are allocated to the same row and are both top-yearning. Then batch $y + c + 1$ lies vertically between batches y and $y + c$, so it must be allocated to the same row as these batches (which is the same row that batch $y + 1$ is allocated to). Furthermore, it cannot horizontally separate batches y and $y + c$, so it may not be allocated to the same column as these batches. Since all rows of M contain precisely two non-zero entries, this means that batch $y + c + 1$ must be allocated to the same cell as batch $y + 1$, and by our construction, this is the cell that batch $y + c + 1$ corresponds to.

Therefore every batch in \overline{P}_n must be allocated to the same cell that it corresponds to, and thus \overline{P}_n has a unique M -partition.

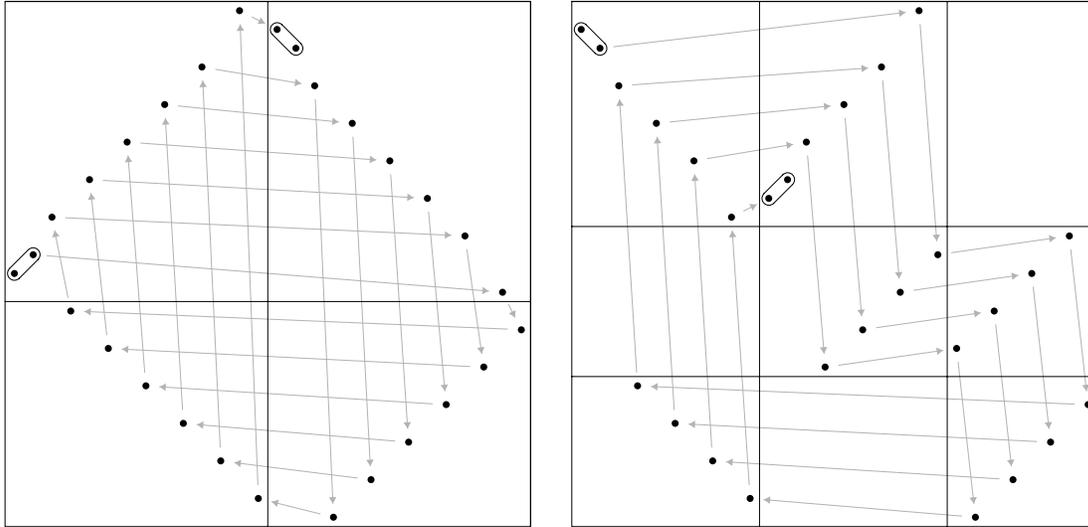


Figure 4: The permutation on the left is an element of the “quasi-square antichain,” introduced in [12] and readily constructed by our algorithm. The permutation on the right comes from a matrix whose graph is a 6-cycle.

Seeing that the same holds for P_n is trivial. Consider the first batch, which is expanded to form a 2×2 matrix when we go from \overline{P}_n to P_n . Of the two nonzero entries in this 2×2 matrix, one of them lies both horizontally and vertically between the other nonzero entry and batch $c+1$; call this entry interior, and then do the same for the last batch. Removing the two interior entries gives \overline{P}_n back, and we know that it has a unique M -partition. But reinserting the interior entries cannot offer us any more possibilities for partitioning. \diamond

The main technical step now complete, we are ready to prove that the permutations we have constructed do indeed form antichains.

Theorem 4.3. *Let M be a $0/\pm 1$ matrix. If $G(M)$ contains a cycle, then $\text{Prof}(M)$ contains an infinite antichain, given by the permutations coming from P_n for n sufficiently large.*

Proof: As we have already remarked, we may assume that M contains an even number of -1 s and that $G(M)$ is nothing but a cycle. Let us assume this cycle is of length $c+1$, and that $n > m$ are both at least c and large enough so that P_m and P_n have unique M -partitions (that we may make this assumption is the content of Lemma 4.2).

We would like to show that P_m and P_n are incomparable. Quite trivially, $P_n \not\leq P_m$, because P_n is larger than P_m , so it suffices to show that $P_m \not\leq P_n$. Suppose to the contrary that $P_m \leq P_n$. Then there is at least one submatrix of P_n that reduces to P_m . In this manner we get a one-to-one map from $\text{supp}(P_m)$ into $\text{supp}(P_n)$, or, as we will think of it, a map from the batches of P_m into the batches of P_n (with the first and last batches of

P_m possibly being mapped into more than one batch of P_n). We begin by making two claims about this mapping:

- (i) if batch i of P_m is mapped into batch j of P_n for some $i \in [2, m - 1]$, then batch $i + 1$ of P_m is mapped into a batch of P_n of numbered at most $j + 1$, and
- (ii) batch 1 of P_m must be mapped into batch 1 of P_n .

We begin with the proof of (i). We may assume without loss that batches i and $i + 1$ of P_m correspond to cells of M that share a row, and that both batches are top-yearning. First note since P_m and P_n have only one M -partition each, batch $i + 1$ must be mapped to a batch of P_n with number congruent to $j + 1$ modulo c . Furthermore, since batches i and $i + 1$ are both top-yearning and correspond to cells in the same row, batch $i + 1$ lies below batch i , so it must be mapped to a batch below batch j . These two restrictions leave only the possibilities we have allowed for.

Now we have to prove claim (ii). Clearly batch 1 of P_m cannot be mapped into the last batch of P_n , so if the claim does not hold then this batch is mapped into two batches of P_n , say $r + 1$ and $r + sc + 1$ (by the uniqueness of M -partitions, these two batches must be congruent to 1 modulo c). Let us suppose that the first batch of P_m has top-right yearn, and that the first and second batches of P_m correspond to cells of M that share a row. Now consider where batch 2 may be mapped to. By the same argument we used in (i), batch 2 must be mapped to a batch this lies below the batch $r + 1$, so it must be mapped to a batch with number at most $r + 2$. We now follow the implication of (i) all the way around the cycle, to see that batch $c + 1$ of P_m must be mapped into a batch with number at most $r + c + 1$. This is a contradiction because either $r + c + 1 = r + sc + 1$ (our mapping was supposed to be one-to-one) or $r + c + 1 < r + sc + 1$ and thus batch $c + 1$ is mapped to a batch that horizontally separates the two entries that batch 1 was mapped to.

Having established (i) and (ii) we are almost done. The first batch of P_m must be mapped to the first batch of P_n , so the second batch of P_m must be mapped to the second batch of P_n , and so on, until we conclude that the $m - 1$ st batch of P_m must be mapped to the $m - 1$ st batch of P_n . Now we have no options for the last batch. Suppose without loss that the $m - 1$ st and m th batches of P_m correspond to row-sharing cells of M , and that the $m - 1$ st batch is top-yearning. Then the m th batch (which consists of two non-zero entries) must lie entirely below the $m - 1$ st batch. This means that the m th batch of P_m must be mapped to a batch of number at most m in P_n . Additionally, of course, the m th batch of P_m may not be mapped into a batch that any other batch of P_m has been mapped into, so we have reached a contradiction, proving the theorem. \diamond

5 Exotic Fundamental Antichains

If we have an antichain of permutations A , then we may form infinitely more antichains from it by direct sums (or skew sums, or in several other ways). For instance, $\{1324 \oplus a :$

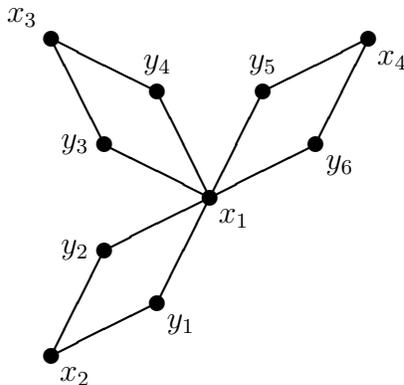


Figure 5: $G \begin{pmatrix} -1 & 1 & -1 & 1 & -1 & 1 \\ 1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 \end{pmatrix}$

forms an antichain, which lies in the closure of

$$\{P_{10}\} \cup \{P_n : 13 \leq n \equiv 1 \pmod{4}\}$$

but is not a subset of this latter antichain. Therefore this antichain is not fundamental, and in particular, $\{P_n : n \geq 9\}$ is not fundamental.

In general, suppose that M is a $0/\pm 1$ matrix for which $G(M)$ is precisely a cycle of length c . If we fix some integer $d \in [c]$, the set

$$\{P_n : n \text{ is sufficiently large and } n \equiv d \pmod{c}\}$$

can be shown to form a fundamental antichain.

Up to this point, all fundamental antichains in the literature and all antichains produced by our algorithm as we have described it are periodic in some sense. We aim in this section to convince the reader that our construction from the last section can be generalized to construct exotic fundamental antichains without this periodicity.

In the description of the construction we assumed that $G(M)$ was a cycle. Suppose instead that we let $G(M)$ contain two or more cycles intersecting at a single vertex. For example, let us take

$$M = \begin{pmatrix} -1 & 1 & -1 & 1 & -1 & 1 \\ 1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 \end{pmatrix},$$

so $G(M)$ will be the graph depicted in Figure 5.

Let us take the first batch to correspond to cell $(1, 1)$ and to have bottom-right yearn. Immediately we are faced with a predicament: there are five other entries on the first row

of M , so which should we choose for the second batch? This situation can be rectified by supplying additional input to the algorithm. Let us supply a word w on the letters 0, 1, and 2 where 0 means that the next several batches should trace out the left cycle, (1, 1), (1, 2), (2, 2), and (2, 1), 1 means that the next several batches should trace out the middle cycle, (1, 3), (1, 4), (3, 4), and (3, 3), and 2 means that the next several batches should trace out the cycle in columns 5 and 6. Then our construction is once again well-defined, and a slight adaptation of the proofs in the last section would show that it still produces antichains.

To produce an aperiodic antichain we need only select an aperiodic word as w . We define the (infinite) binary Thue-Morse word, t , by $t = \lim_{n \rightarrow \infty} u_n$ where $u_0 = a$, $v_0 = b$, and for $n \geq 1$, $u_n = u_{n-1}v_{n-1}$ and $v_n = v_{n-1}u_{n-1}$. For example,

$$u_6 = abbabaabbaababbabaababbaabbabaab.$$

Now we replace each occurrence of abb by 2, each occurrence of ab by 1 (after replacing the occurrences of ab), and each remaining occurrence of a by 0 to get the word w on the letters 0, 1, 2. Applying these substitutions to u_6 , we obtain the word

$$21020121012021.$$

It is known (see, for example, Lothaire [11]) that the word w is square-free. An element of an antichain produced in this manner is shown in Figure 6.

To get an aperiodic fundamental antichain from this construction, we need only make sure to take elements P_n for which the last batch always corresponds to the same cell of M . For example, if we take all permutation matrices P_n (for n sufficiently large) produced by the operation described above, they will form an antichain, but not a fundamental one. Instead if we take all elements P_n where n is sufficiently large and the last batch corresponds to entry (1, 1), this will be a fundamental antichain.

Even though w contains three letters, with a little care we can use it to build an antichain in the profile class of matrix whose graph has only two cycles, for instance, the matrix

$$M = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 0 \end{pmatrix}.$$

To do this we interpret the letters of w differently. If we encounter a 0, we go around the cycle we just looped around in the same direction (clockwise or counterclockwise). In the case of a 1, we go around the other cycle, but keep the direction of the last cycle. If we see a 2, we switch cycles and direction.

For example, suppose we begin by traveling around the right-most cycle of M in a clockwise direction, passing through the entries (2, 2), (1, 2), (1, 3), and (2, 3) in that order. Now we read the first letter of w . Since it is a 2, we go around the left-most cycle of M counterclockwise, passing through (2, 3), (3, 1), and (3, 2). The next letter of w is a 1, so we return to the right-most cycle of M , but in the counterclockwise direction this time, resulting in the walk (2, 2), (2, 3), (1, 3), (1, 2). Since the fourth letter of w is a 2, we go on to walk around the left-most cycle of M in a clockwise direction.

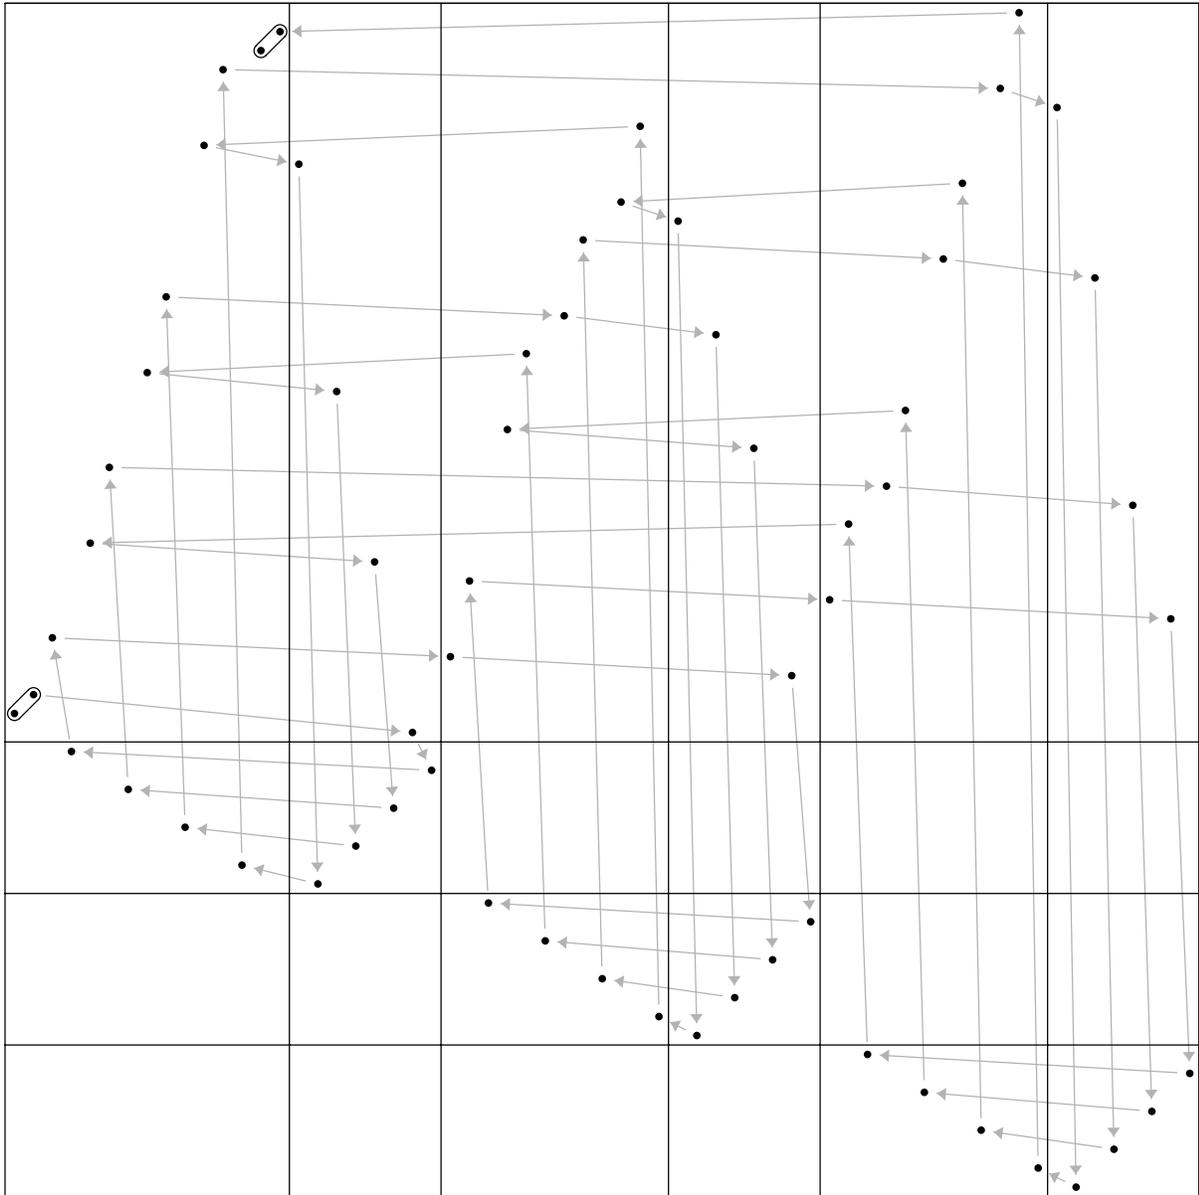


Figure 6: An element of an aperiodic infinite antichain constructed from the Thue-Morse word.

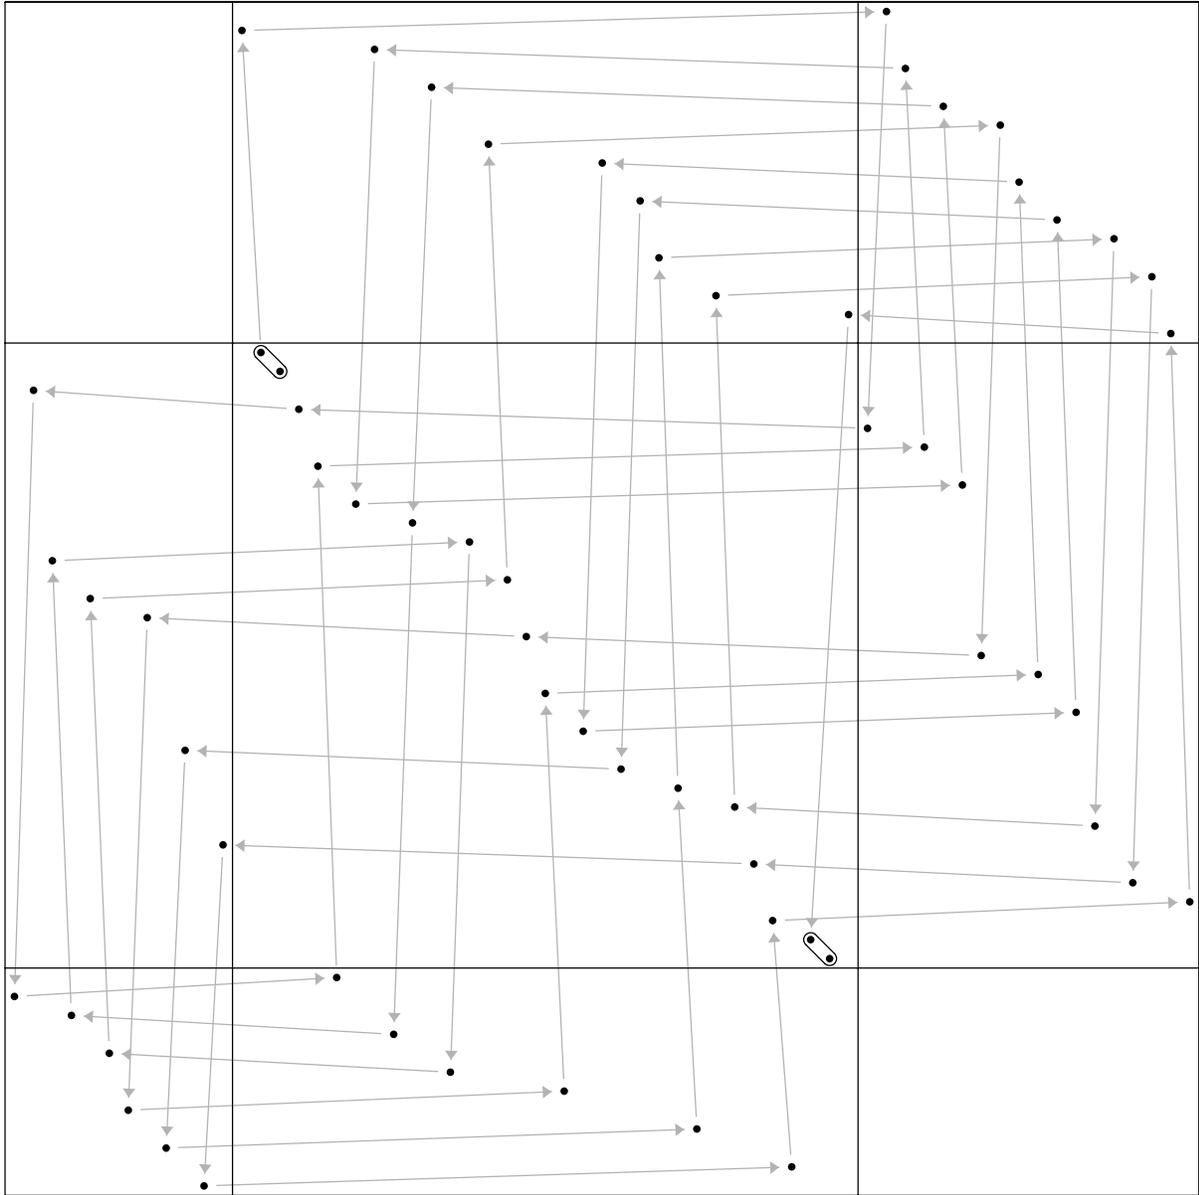


Figure 7: An element of an aperiodic infinite antichain constructed from a matrix with two cycles and the Thue-Morse word.

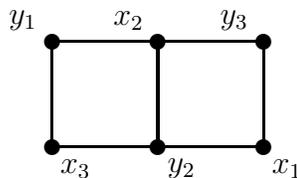


Figure 8: $G \begin{pmatrix} 0 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 0 \end{pmatrix}$

An element of the the antichain constructed in this way is shown in Figure 7.

We have so far discussed only one manner in which antichains produced by our construction might fail to be fundamental, but in this more general setting there are a couple more subtleties. Let us again consider the matrix

$$\begin{pmatrix} 0 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 0 \end{pmatrix}.$$

If our stroll through the entries of this matrix contains a sequence like $(2, 2), (2, 3), (1, 3), (2, 3), (2, 2)$, then the resulting antichain will not be fundamental. And if our stroll does not contain infinitely many cycles, then the resulting sequence of permutations will not even form an antichain.

6 Concluding Remarks

Recently and independently, Albert and Atkinson [1] and Murphy [12] have introduced another method for proving that closed sets of permutations are partially well-ordered. An *interval* of $p \in S_n$ is a segment $p(i)p(i+1) \dots p(j)$, where $1 \leq i \leq j \leq n$, such that $\{p(i), \dots, p(j)\}$ forms a set of consecutive integers. Every permutation in S_n contains trivial intervals of length 1 and n . If p contains no non-trivial intervals then it is said to be *interval-free* or *simple*. For example, 35142 is simple but 25341 is not.

Using the full version of Higman's Theorem from [8], it is not hard to show that an analysis of the simple permutations in a closed set can sometimes be enough to show that the set is partially well-ordered.

Theorem 6.1. [1, 12] *Let X be a closed set of permutations. If X contains only finitely many simple permutations, then X is partially well-ordered.*

An analogous theorem for tournaments exists, and has been used in that context to show that some closed sets of permutations are partially well-ordered. The reader is referred to Latka [9] for an example of this.

In fact, if the hypotheses of Theorem 6.1 happen to hold, then they can be established by computer due to the following theorem, which has been proven in the special case

of permutations by Albert and Atkinson [1] and Murphy [12] and in the more general context of binary relational systems by Schmerl and Trotter [14].

Theorem 6.2. [1, 12, 14] *Every simple permutation of length $n > 2$ contains a simple permutation of length $n - 1$ or $n - 2$.*

Our first aim is to show that our Corollary 3.4 is not a special case of Theorem 6.1. Consider the set of permutations $D = d_1, d_2, \dots$ given by

$$d_k = 2, 4, \dots, 2k \mid 1, 3, \dots, 2k - 1.$$

(Here, as before, the vertical bar is included merely to make the permutation easier to parse, and has no mathematical meaning.) These permutations are all simple, and $\{M_{d_1}, M_{d_2}, \dots\} \subset \text{Prof} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$. Hence Corollary 3.4 (in fact, even Theorem 2.1) can be used to show that $\text{Prof} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$ is partially well-ordered, whereas Theorem 6.1 cannot draw this conclusion.

There are also situations in which Theorem 6.1 can show that a set is partially well-ordered when Corollary 3.4 cannot. To describe these we need to define the *wreath product* of permutations. Let p be an n permutations and let q_1, \dots, q_n be permutations of any positive length. Then $p \wr (q_1, \dots, q_n) = p_1 \dots p_n$ where p_1, \dots, p_n are all intervals such that

- (i) $\text{red}(p_i) = q_i$ for all $i \in [n]$, and
- (ii) if a_i is a term of p_i for all $i \in [n]$ then $\text{red}(a_1 \dots a_n) = p$.

For example, $12 \wr (q_1, q_2) = q_1 \oplus q_2$ and $213 \wr (q_1, q_2, q_3) = (q_1 \ominus q_2) \oplus q_3$.

Now let X denote the largest closed set whose set of simple permutations is precisely $\{1, 12, 21, 3142\}$. Such a set exists because the union of any number of closed sets is again a closed set. This set is partially well-ordered by Theorem 6.1, but it contains the infinite sequence of permutations

$$\begin{aligned} z_1 &= 3, 1, 4, 2, \\ z_2 &= 11, 9, 12, 10 \mid 3, 1, 4, 2 \mid 15, 13, 16, 14 \mid 7, 5, 8, 6, \\ &\dots \\ z_k &= z_{k-1} \wr (3142, 3142, \dots, 3142). \end{aligned}$$

It is perhaps easiest to see the pattern in this sequence by looking at the matrices, for

- [4] M. D. Atkinson and T. Stitt, Restricted permutations and the wreath product, *Discrete Math.* **259** (2002), 19-36.
- [5] G. L. Cherlin and B. J. Latka, Minimal antichains in well-founded quasi-orders with an application to tournaments, *J. Combin. Theory Ser. B* **80** (2000), 258–276.
- [6] Z. Füredi and P. Hajnal, Davenport-Schinzel theory of matrices, *Discrete Math.* **103** (1992), 233-251.
- [7] J. Gustedt, “Algorithmic Aspects of Ordered Structures,” Ph.D. dissertation, Technische Universität Berlin, 1992.
- [8] G. Higman, Ordering by divisibility in abstract algebra, *Proc. London Math. Soc.* **2** (1952), 326-336.
- [9] B. J. Latka, Tournaments that omit N_5 are well-quasi-ordered, preprint.
- [10] R. Laver, Well-quasi-orderings and sets of finite sequences, *Math. Proc. Camb. Philos. Soc.* **79** (1976), 1-10.
- [11] M. Lothaire, “Combinatorics on Words,” Encyclopedia of Math., Vol. 17, Addison Wesley, Reading, MA, 1983.
- [12] M. M. Murphy, Ph.D. dissertation, University of St. Andrews, 2002.
- [13] V. R. Pratt, Computing permutations with double-ended queues, parallel stacks and parallel queues, *Proc. ACM Symp. Theory of Computing* **5** (1973), 268-277.
- [14] J. Schmerl and W. Trotter, Critically indecomposable partially ordered sets, graphs, tournaments and other binary relational structures, *Discrete Math.* **113** (1993), 191-205.
- [15] R. Simion and F. W. Schmidt, Restricted permutations, *European J. Combin.* **6** (1985), 383–406.
- [16] D. Spielman, M. Bóna, An infinite antichain of permutations, *Electronic J. Combinatorics* **7** (2000), #N2.
- [17] R. E. Tarjan, Sorting using networks of queues and stacks, *J. of the ACM* **19** (1972), 341-346.
- [18] H. Wilf, The patterns of permutations, *Discrete Math.* **257** (2002), 575-583.