

POSETS OF MATRICES AND PERMUTATIONS WITH FORBIDDEN SUBSEQUENCES

NIGEL RAY AND JULIAN WEST

ABSTRACT. The enumeration of permutations with specific forbidden subsequences has applications in areas ranging from algebraic geometry to the study of sorting algorithms. We consider a ranked poset of permutation matrices whose global structure incorporates the solution to the equivalent problem of enumerating permutations which contain a *required* subsequence. We describe this structure completely for saturated chains of lengths one and two, so settling several new and general instances of the original problem, and conclude with a superficial asymptotic investigation of arbitrary chains whose length is small by comparison with the rank of its constituent permutations. The value of this approach is reflected in the appearance of closed polynomial formulae (related to the Robinson-Schensted correspondence) and of a framework for the systematic analysis of associated combinatorial questions; indeed, we begin by studying a simpler poset of 0-1 sequences as the natural environment in which to introduce our *insertion* and *deletion* operators.

1. INTRODUCTION AND NOTATION

Throughout this work we are concerned with bijections of a finite set. If the set has m elements we refer to the bijections as m -permutations, or permutations of *dimension* m ; with respect to composition of functions they form the symmetric group S_m . We find it most convenient to represent each such permutation π as an $m \times m$ matrix, acting on the vectors of the standard basis for \mathbb{R}^m by permuting them accordingly. The entries are given by

$$(1.1) \quad \pi_{i,j} = \delta_{i,\pi(j)},$$

in terms of the Kronecker delta; thus $\pi_{\pi(j),j} = 1$ for all $1 \leq j \leq m$. Occasionally, we shall interpret π as a list $(\pi(1)\pi(2)\dots\pi(m))$.

We deem that a permutation π *contains* a permutation ν if the matrix π contains the matrix ν as a submatrix; that is, if we may obtain ν from π by deleting rows and columns. If the number of rows and columns deleted is k , we describe ν as having *codimension* k in π . If π does not contain ν , then we say that π *avoids* ν . Given any set P of permutations, possibly of different dimensions, we follow Simion and Schmidt [13] in writing $S_n(P)$ for the set of all permutations of dimension n which avoid every element of P .

The original formulation of these definitions was in terms of sequences of distinct integers, and is given in full in [15], for example. It is considerably more cumbersome, and characterises $S_n(P)$ as those permutations which have the elements of P as forbidden subsequences.

The problem of enumerating the sets $S_n(P)$ for various n and P has received considerable publicity over the last two decades, and applications have emerged in such diverse areas as algebraic geometry, combinatorics, and differential topology. For example, Macdonald [8] (and others) have shown that the *vevilliary* permutations $S_n(2143)$ are relevant to the theory of Schubert polynomials, and therefore to the cohomological structure of flag manifolds; Knuth [5] has demonstrated that the permutations $S_n(231)$ are exactly those which are *stack*

sortable; and a beautiful theorem of Lakshmibai and Sandhya [6] asserts that the permutations $S_n(1324, 2143)$ are those whose corresponding Schubert varieties are smooth. Recent work of Billey and Warrington [1] on (321)-hexagon-avoiding permutations, and of Mansour and Vainshtein [9] on $S_n(132)$, has developed relationships with Kashdan-Lusztig and Chebyshev polynomials respectively. The trend in these investigations has been for P to be small and n large, so that asymptotic formulae are of major interest. For example, Bóna [2] has proved that $S_n(1423) < S_n(1234) < S_n(1324)$ for all suitably large n .

Our aim is to study a certain partially ordered set, whose structure unifies such problems and suggests an alternative approach to their solution. This poset has made several implicit appearances in earlier literature, such as the work of Laver [7] in 1976. For basic information and terminology relating to partial orderings, we refer readers to Stanley's book [14].

Let S_∞ denote the disjoint union of all the symmetric groups S_m for $m \geq 1$, and write $\nu \subseteq \pi$ whenever ν is contained in π . Clearly \subseteq is a partial ordering on S_∞ , with a zero element given by the identity in S_1 . We may then also specify a *rank* function r by setting

$$r(\pi) = m - 1 \quad \text{for all } \pi \in S_m.$$

It will be convenient to reserve the notation \mathcal{S} for the resulting graded poset. Note that a complete determination of the structure of \mathcal{S} implies knowledge of $S_n(P)$ for every n and P .

To any set of permutations P in S_m we may associate the *filter* $T(P)$, defined by

$$(1.2) \quad T(P) = \bigcup_{\pi \in P} \{\nu \in \mathcal{S} : \pi \subseteq \nu\}.$$

Each such filter is a disjoint union of finite subsets $T_k(P)$, made up of those elements of \mathcal{S} in which some π in P has codimension k . Thus $T_k(P)$ is a subset of S_{m+k} ; we write $t_k(P)$ for its cardinality. Clearly $S_{m+k}(P)$ and $T_k(P)$ are complementary in S_{m+k} , so their enumeration problems are equivalent. By analogy with (1.2), we may also associate to P the *ideal* $U(P)$, defined by

$$(1.3) \quad U(P) = \bigcup_{\pi \in P} \{\nu \in \mathcal{S} : \nu \subseteq \pi\}.$$

It is a finite set, and is the disjoint union of subsets $U_k(P)$, made up of those elements of \mathcal{S} which have codimension k in some member of P . Thus $U_k(P)$ is a subset of S_{m-k} ; we write $u_k(P)$ for its cardinality.

Motivated by this change of emphasis, we shall initiate the study of $T_k(\pi)$ in cases where k is small by comparison with the dimension of π . We refer to the determination of all values of $t_k(\pi)$ as the *codimension k problem* for \mathcal{S} . Its solution is tantamount to determining the structure of the saturated chains of length k . The corresponding investigation for $U_k(\pi)$ and $u_k(\pi)$ is of equal interest, but much greater difficulty, as we show by occasional example.

In fact the codimension 1 problem is relatively simple. Its solution is given by the formula

$$(1.4) \quad t_1(\pi) = m^2 + 1 \quad \text{for all } \pi \in S_m,$$

which was originally obtained in 1990 by Bloomberg, but so far as we are aware remains unpublished. Note that this identifies the number of elements covering each π in \mathcal{S} , and therefore indicates a fundamental regularity in the structure of the poset. We rederive Bloomberg's result in Section 7. Since (1.4) is polynomial in m , the same cannot be true for the cardinality of $S_{m+1}(\pi)$, whose value $m! - m^2 - 1$ follows immediately.

In contrast, we note by inspection that $u_1(1234)$, $u_1(1243)$, $u_1(3241)$ and $u_1(2413)$ are given by 1, 2, 3, and 4 respectively, illustrating the observation of Section 7 that

$$1 \leq u_1(\pi) \leq m \quad \text{for all } \pi \in S_m.$$

At first sight, this linear bound is encouraging; however, the evidence that $u_1(\pi)$ attains every possible value must be compared with the dependence of $t_1(\pi)$ on m alone. This provides our first clue that the downward structure of \mathcal{S} is more elusive than its upward counterpart.

As part of our analysis for the codimension 2 problem, we prove that $t_2(\pi)$ is again a polynomial in m , although the polynomial now has degree 4, depends up to a linear term on the choice of π , and is not valid for the special case $m = 0$. Thus, for example, Theorem 8.16 states that

$$t_2(\iota_m) = \frac{1}{2}(m^4 + 2m^3 + m^2 + 2m + 6),$$

where ι_m is the identity m -permutation. In general, we conjecture that $t_k(\pi)$ is a polynomial in m of degree $2k$, that it depends up to a term of degree less than k on π , and that it is only accurate for values of m greater than some small integer. In Section 9 we explain how supporting evidence arises from the Robinson-Schensted correspondence [12]. The appearance of this correspondence offers the tantalising possibility that the representation theory of the symmetric groups might have a rôle to play; otherwise, we have found no way to utilise their group structure.

We shall continue to interpret the elements of \mathcal{S} as matrices throughout, according to (1.1). Indeed, we may extend the constructions above to form the poset \mathcal{M} of *all* matrices under inclusion, of which \mathcal{S} is merely one of many subposets ripe for investigation. Another is \mathcal{Y} , the poset of 0-1 row vectors y , whose structure is considerably simpler and provides the optimal environment in which to introduce certain basic operations. In so doing, we explain how to compute the statistic $t_k(y)$ for all values of k , in contrast to \mathcal{S} .

We begin in Section 2 by investigating appropriate aspects of \mathcal{Y} , introducing the operations of insertion and deletion. We continue the analysis in Section 3, studying synonymy classes of insertions, and related concepts which play an important subsequent rôle. We extend the insertion and deletion operators to \mathcal{S} in Section 4, establishing our matrix notation in the process, and plan the campaign for \mathcal{S} in Section 5, outlining our approach to the general codimension k problem in terms of upper and lower bounds for $T(\pi)$. We continue this study in Section 6, where we develop the notions of superblocs and track matrices. In Section 7 we implement our strategy in codimension 1, which exemplifies the techniques in a suitably straightforward case, and develop the theme in Section 8, where we analyse many aspects of the upward structure in codimension 2. Finally, in Section 9 we discuss certain features of the problem for codimensions $k > 2$, obtaining some elementary polynomial approximations.

We shall regularly consider partitioned versions of matrices M , for which we use the standard form

$$(1.5) \quad \begin{pmatrix} M(0,0) & \dots & M(0,q) \\ \vdots & \vdots & \vdots \\ M(p,0) & \dots & M(p,q) \end{pmatrix}.$$

We refer to the submatrices $M(r,s)$ as *blocks*, and abbreviate each $y(0,s)$ to $y(s)$ in the study of row vectors y .

Our conventions dictate that we write ι_m for the $m \times m$ identity matrix, and it will be convenient to reserve the notation η_m for the corresponding antiidentity. As a permutation, η_m is the *longest* element of S_m , and reverses the order of the m elements upon which it acts.

Several of our arguments concerning elements of \mathcal{S} will proceed on a case by case basis, and it is often possible to reduce the work by using symmetry. To formalise this process we take any permutation matrix π , and define π^{rr} and π^{cr} to be the matrices obtained by respectively reversing the order of its rows and its columns; we define π^{tr} to be its transpose, as usual. Thus

$$\iota_m^{rr} = \eta_m, \quad \eta_m^{cr} = \iota_m \quad \text{and} \quad \eta_m^{tr} = \eta_m,$$

for example. Of course, rr , cr , and tr generate an action of the dihedral group D_4 on \mathcal{S} . This action induces a bijection between $T_k(\pi)$ and $T_k(\pi^d)$ for any d in D_4 , whence $t_k(\pi)$ and $t_k(\pi^d)$ are equal. So far as the poset \mathcal{Y} is concerned the only nontrivial such symmetry which survives is cr , reversing the order of the entries in each vector and generating a cyclic subgroup C_2 of D_4 .

For each m , we always denote the set of natural numbers $\{1, 2, \dots, m+1\}$ by $[m+1]$. Whenever we select k elements from $[m+1]$, allowing repetitions, we insist that the resulting multiset R on $[m+1]$ be written in nondecreasing order, say as $\{r_1, \dots, r_k\}$.

We are especially grateful to Richard Stanley for his continuing encouragement, and to Grant Walker and Reg Wood for providing helpful and illuminating examples. Nevertheless, preparation of our final text has taken nine years longer than it ought, and further study of the poset \mathcal{S} has already been undertaken by others; for example see [3], where a new infinite antichain is constructed.

2. BINARY SEQUENCES, INSERTION AND DELETION

In this section we consider the poset \mathcal{Y} . We deduce basic properties of the insertion and deletion operators; although not strictly necessary for the study of \mathcal{Y} , they are an important ingredient in our investigation of the analogous operators for \mathcal{S} in later sections.

For each $m \geq 0$, we write Y_m for the set of sequences of length m whose entries are either 0 or 1, following the convention that Y_0 contains only the empty sequence \emptyset . Thus Y_m has cardinality 2^m , and we may express an arbitrary element y as $y_1 y_2 \dots y_m$; we often refer to y as a *binary m -sequence*. As hinted in Section 1, the disjoint union Y_∞ of all the Y_m inherits the partial ordering \subseteq , induced by $x \subseteq y$ whenever x is a subsequence of y . Thus the empty sequence is a zero element, and we may specify a rank function r by setting

$$r(y) = m \quad \text{for all } y \in Y_m.$$

We refer to the resulting graded poset as \mathcal{Y} .

As with \mathcal{S} , each sequence y in Y_m determines the filter $T(y)$, and we let $t_k(y)$ denote the cardinality of the subset $T_k(y)$, consisting of all binary sequences x such that y has codimension k in x . Of course, $T_k(y)$ is a subset of Y_{m+k} . We refer to the determination of $t_k(y)$ as the *codimension k problem* for \mathcal{Y} , and obtain its solution in this section and the next; this is tantamount to describing the saturated chains of length k in \mathcal{Y} .

We recall that the entire poset \mathcal{Y} admits the group of symmetries C_2 , generated by the involution cr . This induces a bijection between $T_k(y)$ and $T_k(y^{cr})$, so that $t_k(y) = t_k(y^{cr})$.

Every binary m -sequence y has the set $[m+1]$ as its *grid* $G(y)$, interpreted as the set of gaps between successive elements y_i , and we refer to any member of $G(y)$ as a *site* of y . Our

aim is to attack the codimension k problem by inserting binary digits at k sites of a generic y , since any element of $T_k(y)$ may clearly be so constructed.

Let x and y respectively be binary k - and m -sequences, and let C be any k -element multiset of $[m+1]$, written in nondecreasing order as $\{c_1, c_2, \dots, c_k\}$. We define a new $m+k$ -sequence whose $(c_i + i - 1)$ -th entry is x_i and whose remaining m entries are made up of the entries of y in the natural order. We refer to the resulting binary $m+k$ -sequence as $\text{Ins}(x; C; y)$, the *insertion* of x into y at C . Implicitly, we are exploiting the natural correspondence

$$(2.1) \quad v_i = c_i + i - 1.$$

between a k -subset V of $[m+k]$ and a k -multiset C of $[m+1]$. We write V as C^+ and C as V^\bullet when emphasising this bijective relationship. We also consider an associated sequence $v(h)$, where v is the surjection $[m+k] \rightarrow [2]$ defined by $v(c_i + i - 1) = 2$ for each $i = 1, \dots, k$ and $v(h) = 1$ otherwise. It is convenient to refer to the multiset C as a k -*multigrad* (either of y , or $[m+1]$, as the context dictates), and to refer to $(x; C)$ as a k -*pair*; if C is a singleton $\{c\}$, we may write the pair as $(x; c)$.

It is possible, but not necessary for our present purposes, to view the insertion operators as specific elements of an appropriately defined category of operators. As such, we expect them to have basic properties of associativity, commutativity and invertibility, which we now express in terms of the corresponding multigrads. We begin with associativity, which is notationally trickiest.

Lemma 2.2 (Associativity Lemma). *Let x be any binary m -sequence, and let y and z be binary $(k-l)$ - and l -sequences respectively; then for every l -multigrad B of $[k-l+1]$ and k -multigrad C of $[m+1]$, we may define an l -multigrad B^\triangleright of $[k+m-l+1]$ and a $(k-l)$ -multigrad C^\triangleright of $[m+1]$ such that*

$$(2.3) \quad \text{Ins}(\text{Ins}(z; B; y); C; x) = \text{Ins}(z; B^\triangleright; \text{Ins}(y; C^\triangleright; x))$$

in Y_{k+m} . Conversely, for every l -multigrad D of $[k+m-l+1]$ and $k-l$ -multigrad E of $[m+1]$, we may define an l -multigrad D^\triangleleft of $[k-l+1]$ and a k -multigrad E^\triangleleft of $[m+1]$ such that

$$(2.4) \quad \text{Ins}(\text{Ins}(z; D^\triangleleft; y); E^\triangleleft; x) = \text{Ins}(z; D; \text{Ins}(y; E; x))$$

in Y_{k+m} .

Proof. The left hand side of (2.3) is characterised by the surjection $[k+m] \rightarrow [3]$, given by

$$(2.5) \quad v(h) = \begin{cases} 1 & \text{if } h = c_i + i - 1 \text{ and } c_i = b_j + j - 1 \text{ for some } i, j \\ 2 & \text{if } h = c_i + i - 1 \text{ and } c_i \neq b_j + j - 1 \text{ for some } i, \text{ any } 1 \leq j \leq l \\ 3 & \text{if } h \neq c_i + i - 1 \text{ for any } 1 \leq i \leq k; \end{cases}$$

the h th element of the lefthand side of (2.3) is x_q , y_q , or z_q , as $v(h)$ is the q th 3, 2, or 1 respectively. Now define B^\triangleright and C^\triangleright by

$$\begin{aligned} B^\triangleright &= \{c_{b_1} + b_1 - 1, \dots, c_{b_j+j-1} + b_j - 1, \dots, c_{b_l+l-1} + b_l - 1\} \quad \text{and} \\ C^\triangleright &= \{c_1, \dots, c_{b_1-1}, c_{b_1+1}, \dots, c_{c_j+j-2}, c_{c_j+j}, \dots, c_{b_l+l-2}, c_{b_l+l}, \dots, c_k\} \end{aligned}$$

respectively. Then a similar check reveals that the right hand side of (2.3) is determined by the same surjection (2.5), as required.

The entire procedure may be reversed by defining D^\triangleleft by

$$D^\triangleleft = \{i_1, \dots, i_j, \dots, i_l\}$$

where $e_{i_j-1} \leq d_j - i_j + 1 \leq e_{i_j}$ for $1 \leq j \leq l$, and E^\triangleleft by

$$E^\triangleleft = \{e_1, \dots, e_{i_1-1}, d_1 - i_1 + 1, e_{i_1}, \dots, e_{i_j-1}, d_j - i_j + 1, e_{i_j}, \dots, \dots, e_{i_l-1}, d_l - i_l + 1, e_{i_l}, \dots, e_{k-l}\};$$

(2.4) may then be checked by the same method. \square

Note that the construction of B^\triangleright and C^\triangleright depends on both B and C ; so we write $(B, C)^\triangleright$ for the pair $(B^\triangleright, C^\triangleright)$, and say that it is obtained from (B, C) by *forward* association. Similar remarks apply to $(D, E)^\triangleleft$, which we say is obtained from (D, E) by *backward* association. These associations are, of course, related by the formulae

$$(2.6) \quad ((B, C)^\triangleright)^\triangleleft = (B, C) \quad \text{and} \quad ((D, E)^\triangleleft)^\triangleright = (D, E).$$

The Associativity Lemma may be applied repeatedly so as to exchange a string of insertions between z and x .

We now turn to commutativity, considering x and C as above.

Lemma 2.7 (Commutativity Lemma). *Given any binary k -sequence w , we may define an m -multigrad C^\perp in $[k+1]$ such that*

$$(2.8) \quad \text{Ins}(w; C; x) = \text{Ins}(x; C^\perp; w)$$

in Y_{k+m} .

Proof. The surjection of $[k+m]$ onto $[2]$ corresponding to the right hand side of (2.8) is obtained from that on the left hand side by exchanging the rôles of 1 and 2. If we define C^\perp by

$$C^\perp = \{1, \dots, 1, \dots, i, \dots, i, \dots, k+1, \dots, k+1\},$$

where 1 appears c_1 times, i appears $c_i - c_{i+1}$ times for $1 < i < k+1$, and $k+1$ appears $m - c_k$ times, then replacing C by C^\perp achieves the same effect, as sought. \square

The Commutativity Lemma indicates the existence of a certain duality between the codimension k problem for m -sequences, and the codimension m problem for k -sequences.

To express invertibility properties, we are led to the notion of a deletion operator. For any k element subset V of $[m]$, and any binary m -sequence y , we define the *deletion* $\text{Del}(V; y)$ of entries V from y to be the binary $(m-k)$ -sequence obtained by deleting from y all elements y_i with $i \in V$; it admits a standard partitioning into blocks bounded by the deleted elements. The deleted elements themselves form a binary k sequence $y|V$, the *restriction* of y to V , and both $\text{Del}(V; y)$ and $y|V$ patently lie in $U(y)$.

We consider w , C and y as in Lemma 2.7 above.

Lemma 2.9 (Invertibility Lemma). *If we define C^- to be $(C^\perp)^+$, then we have*

$$\text{Del}(C^+; \text{Ins}(w; C; y)) = y \quad \text{and} \quad \text{Del}(C^-; \text{Ins}(w; C; y)) = w$$

in Y_m . *Given any k element subset V of $[m]$, we also have*

$$(2.10) \quad \text{Ins}(y|V; V^\bullet; \text{Del}(V; y)) = y$$

in S_m .

It often happens that $\text{Del}(V; y)$ agrees with $\text{Del}(W; y)$, even though $y|V \neq y|W$; this is one way in which distinct insertions in z may yield the same sequence y . Our main effort in studying \mathcal{Y} (and subsequently \mathcal{S}) is directed towards resolving such ambiguities.

We finally consider the symmetry cr , with x and C as in Lemma 2.2.

Lemma 2.11 (Symmetry Lemma). *Given any binary k -sequence w , we may define a k -multigrind \overline{C} such that*

$$(\text{Ins}(w; C; x))^b = \text{Ins}(w^b; \overline{C}; x^b)$$

in Y_{k+m} .

Proof. If we define \overline{C} by

$$\overline{C} = \{m + 2 - c_k, \dots, m + 2 - c_{k-i}, \dots, m + 2 - c_1\},$$

the required equation follows immediately. \square

Using the bijection (2.1) between multisets C of $[m+1]$ and subsets V of $[m+k]$, the subset \overline{V} corresponding to the multiset \overline{C} is constructed by replacing each v in V with $(m+k) - v$.

It will occasionally prove convenient to write a generic element of C^+ , C^\perp , C^- , or \overline{C} as c_i^+ , c_i^\perp , c_i^- or \overline{c}_i respectively.

3. THE STRUCTURE OF \mathcal{Y}

In this section we compute $t_k(y)$ for any binary m -sequence y ; our main aim is to establish methodology for the study of \mathcal{S} in later sections. We begin by investigating $T_k(y)$.

It is clear (but also follows from the Invertibility Lemma 2.9) that any z in $T_k(y)$ may be obtained from y by the insertion of an appropriate k -pair $(w; C)$, so it is convenient to introduce the set $\Psi_{k,m}$ of all such pairs. Insertion in y then defines a map

$$(3.1) \quad f_y: \Psi_{k,m} \longrightarrow T_k(y)$$

which is a surjection, but fails to be an injection, by (2.10) and the remarks following. Note that the cyclic group C_2 acts on $\Psi_{k,m}$ according to the rule $(w; C)^{cr} = (w^{cr}; \overline{C})$.

Once we have introduced f_y , it is natural to define two k -pairs $(w; C)$ and $(x; D)$ to be y -synonymous, written $(w; C) \underset{y}{\sim} (x; D)$, whenever $\text{Ins}(w; C; y)$ and $\text{Ins}(x; D; y)$ are the same.

Then f_y extends to a bijection from the set of y -synonymity classes to $T_k(y)$, and we aim to enumerate these classes by determining the conditions under which two arbitrary pairs are synonymous. Note that the action of cr transforms a y -synonymity Ω into a y^{cr} -synonymity Ω^{cr} .

Let us order the k -multigrinds C lexicographically, writing $C \prec D$ to denote the fact that C precedes D . We then select from any y -synonymity class the unique pair whose multigrind is greatest; this is its *canonical representative*, and is well defined because $w = x$ whenever $(w; C) \underset{y}{\sim} (x; D)$ and $C = D$. We therefore consider the subset $\text{Can}_k(y)$ of $\Psi_{k,m}$ consisting of all canonical representatives, and observe that the map f_y restricts to a bijection between $\text{Can}_k(y)$ and $T_k(y)$. Our quest is therefore to enumerate $\text{Can}_k(y)$.

We label a pair $(w; C)$ as y -active if no element w_j equals y_{c_j} for any $1 \leq j \leq k$, and let $\text{Act}_k(y)$ denote the subset of $\Psi_{k,m}$ consisting of all active pairs; neither $\text{Can}_k(y)$ nor $\text{Act}_k(y)$ is C_2 -invariant.

Proposition 3.2. *For any choice of y and k , we have that*

$$\text{Can}_k(y) = \text{Act}_k(y)$$

as subsets of $\Psi_{k,m}$.

Proof. We prove that a k -pair $(w; C)$ is inactive if and only if it fails to be canonical.

Assume first that $(w; C)$ is y -inactive, so that $w_j = y_{c_j}$ for some $1 \leq j \leq k$ and

$$(3.3) \quad \text{Ins}(w_j; c_j; y) = \text{Ins}(w_j; c_j + 1; y).$$

We choose the minimal such j , and write z for $\text{Del}(\{j\}; w)$ and B for $\{j\}^\perp$; thus (2.10) implies that $\text{Ins}(z; B; w_j) = w$. Applying a forward association of the form (2.3) yields

$$(3.4) \quad \text{Ins}(\text{Ins}(z; B; w_j); C; y) = \text{Ins}(z; B^\triangleright; \text{Ins}(w_j; C^\triangleright; y)),$$

where $C^\triangleright = c_j$. So we may appeal to (3.3) and rewrite the righthand side of (3.4) as $\text{Ins}(z; B^\triangleright; \text{Ins}(w_j; c_j + 1; y))$. We then apply a backward association of the form (2.4) to obtain

$$\text{Ins}(\text{Ins}(z; B; w_j); C; y) = \text{Ins}(\text{Ins}(z; D; w_j); E; y),$$

in which $(D, E) = (B^\triangleright, c_j + 1)^\triangleleft$. It is now simple to check from the definitions that $e_i = c_i$ for $1 \leq i < j$ and that $e_j = c_j + 1$, whence $C \prec E$. Since $(w; C) \underset{y}{\sim} (\text{Ins}(x; D; w_j); E)$, we conclude that $(w; C)$ is not canonical, as required.

Assume secondly that $(w; C)$ fails to be canonical, so that for some $C \prec E$ we have $\text{Ins}(w; C; y) = \text{Ins}(x; E; y)$ in Y_{k+m} . We denote their common value by v , and choose j such that $e_i = c_i$ for $1 \leq i < j$ and $e_j > c_j$. Then $x_i = w_i$ for $1 \leq i < j$, and v_{c_j+j-1} is given by w_j and y_{c_j} respectively. Thus $w_j = y_{c_j}$, and $(w; C)$ is y -inactive, as sought. \square

As will be the case for \mathcal{S} , the main difficulty is with multiple insertions; inserting a single character into a string is more straightforward. For example,

$$\text{Ins}(110, \{3, 3, 3\}, 10111) = \text{Ins}(101, \{4, 4, 4\}, 10111) = \text{Ins}(011, \{5, 6, 6\}, 10111)$$

have the common value 10110111. Clearly $(101, \{4, 4, 4\})$ is not active, because it involves inserting a 1 immediately before a 1, whereas $(110, \{3, 3, 3\})$ is less obviously inactive because the inserted 1s are buffered by an intervening 0; only $(011, \{5, 6, 6\})$ is active, and canonical to boot.

To compute $t_k(y)$ it remains to enumerate $\text{Act}_k(y)$.

Proposition 3.5. *For any y in Y_m , the cardinality of $\text{Act}_k(y)$ is given by*

$$a_k(y) = \sum_{j=0}^k \binom{m+k}{j}.$$

Proof. Whenever a k -pair $(w; C)$ is active, all elements of w are predetermined except those which are inserted at the last site $m+1$. So we partition the set of k -multigrids according to the number of times j which any site other than the last is selected, where $1 \leq j \leq k$. Since there are $\binom{m}{k}$ such multigrids (appealing to the multichoose notation of [14]), and the last $k-j$ entries of w are arbitrary, we deduce that

$$a_k(y) = \sum_{j=0}^k 2^{k-j} \binom{m}{k}.$$

The required formula follows by standard manipulation of binomial coefficients, as described in [11], for example. \square

This result indicates a fundamental regularity in the poset \mathcal{Y} ; the cardinality of each $T_k(y)$ is independent of y in Y_m , and is given by a polynomial of degree k in m whose leading terms are $(m^k + \frac{1}{2}k(k+3)m^{k-1})/k!$.

4. INSERTION AND DELETION IN \mathcal{S}

In this section we extend our insertion and deletion operators to the poset \mathcal{S} , establishing notational conventions as we proceed. We follow Section 2 closely.

For any m -permutation matrix π , its *grid* $G(\pi)$ is the $(m+1) \times (m+1)$ matrix of ordered pairs of natural numbers defined by

$$g_{i,j}(\pi) = (i, j), \quad \text{where } 1 \leq i, j \leq m+1.$$

Clearly $G(\pi)$ is a subset of the plane which is independent of the choice of π , and should be construed as the set of intersections of the $m+1$ horizontal and vertical lines separating the rows and columns of π . As such, it will occasionally be referred to as *the* $(m+1)$ -grid. We call any element of $G(\pi)$ a *site* of π , and describe the collection of sites for which i or j equals 0 or m as the *boundary* of $G(\pi)$. We refer to any site (i, j) for which both $\pi_{i,j-1}$ and $\pi_{i,j}$ are zero as *active* in π , and to the remaining sites as *inactive*. Since each nonzero element of π spawns 2 inactive sites, there are $2m$ in total, and π therefore has $m^2 + 1$ active sites.

Our insertion and deletion operators are defined analogously to those for \mathcal{Y} except we must now work with rows and columns separately. We must also take the whole dihedral group D_4 into account when discussing symmetry. We write our insertion operator as $\text{Ins}(\rho; R, C; \pi)$, where ρ and π are any k - and m -permutations, and where (R, C) is a pair of k element multisets of $[m+1]$. This operator interleaves ρ and π by inserting $\rho_{i,j}$ at the site (r_i, c_j) of π , for all $1 \leq j \leq k$, and we speak of the *insertion* of ρ into π at (R, C) . It is convenient to refer to the pair (R, C) as a *k-multigrad*, and to the triple $(\rho; R, C)$ as a *k-triple*. Sometimes, for example when k is small, it is helpful to write (R, C) as $(r_1, \dots, r_k, c_1, \dots, c_k)$. Of course a 1-multigrad of π is just a site.

We note that (R, C) determines a partitioning

$$(4.1) \quad \begin{pmatrix} \pi(0, 0) & \pi(0, c_1) & \dots & \pi(0, c_k) \\ \pi(r_1, 0) & \pi(r_1, c_1) & \dots & \pi(r_1, c_k) \\ \vdots & \vdots & & \vdots \\ \pi(r_k, 0) & \pi(r_k, c_1) & \dots & \pi(r_k, c_k) \end{pmatrix}$$

of π , and that whenever R or C contains a boundary site or repeated element then the corresponding blocks of (4.1) are empty. So long as we introduce the convention

$$(4.2) \quad r_0 = c_0 = 0 \quad \text{and} \quad r_{k+1} = c_{k+1} = \infty,$$

we may describe the blocks of (4.1) by

$$\pi(r_i, c_j) = (\pi_{t,u}) \quad \text{with} \quad r_i \leq t < r_{i+1} \quad \text{and} \quad c_j \leq u < c_{j+1}.$$

If we write τ for $\text{Ins}(\rho; R, C; \pi)$, we may describe it elementwise by

$$(4.3) \quad \tau_{p,q} = \begin{cases} \rho_{i,j} & \text{if } p = r_i^+ \text{ and } q = c_j^+, \\ \pi_{p-i, q-j} & \text{if } r_i^+ < p < r_{i+1}^+ \text{ and } c_j^+ < q < c_{j+1}^+, \\ 0 & \text{otherwise.} \end{cases}$$

We now turn to the associativity, commutativity, invertibility, and symmetry properties of our insertion operators in \mathcal{S} . These follow directly from their counterparts in \mathcal{Y} , and proofs may easily be read off from Section 2.

Lemma 4.4. *Let π be an arbitrary m -permutation, and let τ and σ be arbitrary $(k-l)$ - and l -permutations respectively. Given any l -multigrad (Q, B) of τ and any k -multigrad (R, C) of π , we have*

$$\text{Ins}(\text{Ins}(\sigma; Q, B; \tau); R, C; \pi) = \text{Ins}(\sigma; Q^\triangleright, B^\triangleright; \text{Ins}(\tau; R^\triangleright, C^\triangleright; \pi))$$

in S_{m+k} . Conversely, given any l -multigrad (S, D) in the $(m+k-l+1)$ -grid and any $(k-l)$ -multigrad (T, E) of π , we have

$$\text{Ins}(\text{Ins}(\sigma; S^\triangleleft, D^\triangleleft; \tau); T^\triangleleft, E^\triangleleft; \pi) = \text{Ins}(\sigma; S, D; \text{Ins}(\tau; T, E; \pi))$$

in S_{m+k} .

Lemma 4.5. *Given any ρ, R, C and π as above, we have*

$$\text{Ins}(\rho; R, C; \pi) = \text{Ins}(\pi; R^\perp, C^\perp; \rho)$$

in S_{m+k} .

To formalise the notion of deleting elements from an m -permutation matrix π , suppose that V and W are k element subsets of $[m]$. Then the *deletion* $\text{Del}(V, W; \pi)$ of rows V and columns W from π is the $(m-k) \times (m-k)$ matrix obtained by deleting all elements $\pi_{i,j}$ with either i in V or j in W ; it admits a standard partitioning into blocks bounded by the deleted elements. By construction, $\text{Del}(V, W; \pi)$ is itself a permutation matrix if and only if there is a k -permutation $\rho(V, W)$ such that the entry 1 occurring in column w_i of π also occurs in row $v_{\rho(i)}$, for each $1 \leq i \leq k$. Henceforth, we insist that this condition holds; it is equivalent to recognising that the pair (V, W) is uniquely specified by either one of its components. So $\text{Del}(V, W; \pi) \subseteq \pi$ by definition, and $\text{Del}(V, W; \pi)$ lies in $U(\pi)$ by (1.3).

Lemma 4.6. *Given any ρ, R, C and π as above, we have*

$$\text{Del}(R^+, C^+; \text{Ins}(\rho; R, C; \pi)) = \pi \quad \text{and} \quad \text{Del}(R^-, C^-; \text{Ins}(\rho; R, C; \pi)) = \rho$$

in S_m . Given any k element subsets V and W of $[m]$, we also have

$$\text{Ins}(\rho(V, W); V^\bullet, W^\bullet; \text{Del}(V, W; \pi)) = \pi$$

in S_m .

The third equation may be rephrased to state that, given any permutation ν of codimension k in π , there is a k -tuple $(\rho; R, C)$ such that π may be written as $\text{Ins}(\rho; R, C; \nu)$. In many cases the choice of V and W will not be unique, and therefore nor will the choice of ρ and (R, C) . But because V (or W) determines $\rho(V, W)$ and W (or V), so R (or C) always determines ρ and C (or R).

We conclude by considering the action of the dihedral group.

Lemma 4.7. *Given any ρ, R, C and π as above, we have*

$$\begin{aligned} (\text{Ins}(\rho; R, C; \pi))^{rr} &= \text{Ins}(\pi^a; \overline{R}, C; \rho^{rr}), & (\text{Ins}(\rho; R, C; \pi))^{cr} &= \text{Ins}(\pi^b; R, \overline{C}; \rho^{cr}) \\ \text{and} & & (\text{Ins}(\rho; R, C; \pi))^{tr} &= \text{Ins}(\pi^t; C, R; \rho^{tr}) \end{aligned}$$

in S_{m+k} .

5. CANONICAL TRIPLES AND SUBSETS

In this section we outline our strategy for computing $t_k(\pi)$ for π and any natural number k . In particular, we describe an upper bound for the set $T_k(\pi)$. Our approach to both problems is motivated by the discussion of \mathcal{Y} in Section 3.

By Lemma 4.6, any ν in $T_k(\pi)$ may be obtained from π by the insertion of a k -triple $(\rho; R, C)$, so it is convenient to introduce the set $\mathcal{Y}_{k,m}$ of all such triples. Insertion in π then defines a surjection

$$f_\pi: \mathcal{Y}_{k,m} \longrightarrow T_k(\pi).$$

Although this map fails to be an injection, it is still instructive to obtain a crude upper bound for $t_k(\pi)$ by enumerating $\mathcal{Y}_{k,m}$. We obtain

$$|\mathcal{Y}_{k,m}| = k! \left(\binom{m+1}{k} \right)^2,$$

and therefore deduce that $t_k(\pi)$ is bounded above by a polynomial of degree $2k$ in m , which for future reference we write as

$$(5.1) \quad \frac{1}{k!} (m^{2k} + k(k+1)m^{2k-1})$$

modulo lower powers of m . Since the coefficients of the powers of m in this expression are themselves functions of k , 5.1 may only be used to approximate $t_k(\pi)$ when k is sufficiently small by comparison with m ; for example, if $k! < m$, then the terms of the polynomial certainly decrease with the necessary speed.

Note that the dihedral group D_4 acts on $\mathcal{Y}_{k,m}$ according to the rules

$$(5.2) \quad \begin{aligned} (\rho; R, C)^{rr} &= (\rho^{rr}; \overline{R}, C), & (\rho; R, C)^{cr} &= (\rho^{cr}; R, \overline{C}) \\ \text{and} \quad (\rho; R, C)^{tr} &= (\rho^{tr}; C, R), \end{aligned}$$

enabling arguments by symmetry to be made on triples.

Following Section 3, we define two k -triples $(\rho; R, C)$ and $(\sigma; S, D)$ to be π -synonymous whenever $\text{Ins}(\rho; R, C; \pi)$ and $\text{Ins}(\sigma; S, D; \pi)$ have the same value τ in S_{m+k} . Assuming that R precedes S lexicographically, we write $(\rho; R, C) \sim_\pi (\sigma; S, D)$; if $S \prec R$ then we reverse the order of the triples. We refer to τ as the *value* of the synonymy, and to $(\rho; R, C)$ as its *lefthand* triple. Thus f_π extends to a bijection from the set of π -synonymy classes to $T_k(\pi)$, and we wish to enumerate these classes by codifying the conditions under which two triples are π -synonymous. The action of any d in D_4 transforms a π -synonymy Ω into a π^d -synonymy Ω^d , although we must interchange the constituent triples whenever $S^d \prec R^d$.

Proposition 5.3. *Suppose given that $(\rho; R, C) \sim_\pi (\sigma; S, D)$; then the two triples are equal if either $R = S$ or $C = D$.*

Proof. We have that

$$\text{Ins}(\rho; R, C; \pi) = \text{Ins}(\sigma; S, D; \pi)$$

in S_{m+k} . Applying Lemma 4.6 (and subsequent comments), we deduce that R determines C and ρ , and S determines D and σ . Thus if $R = S$, then $C = D$ and $\rho = \sigma$, as sought.

The case of $C = D$ follows by symmetry, under the action of tr in D_4 . \square

According to Proposition 5.3, we may select from any π -synonymy class Ω the unique triple whose row multiset is lexicographically greatest; this is the *canonical representative* of Ω , and our conventions ensure that it appears on the right hand side of any synonymy. We

write $\text{Can}_k(\pi)$ for the subset of $\mathcal{Y}_{k,m}$ consisting of all canonical representatives, and observe that the map f_π restricts to a bijection between $\text{Can}_k(\pi)$ and $T_k(\pi)$. In order to enumerate $\text{Can}_k(\pi)$, we begin by establishing a more delicate upper bound than provided by (5.1).

We label a triple $(\rho; R, C)$ as π -*active* (or *active*, in the usual situation where π is understood) whenever every site of the form $(r_{\rho(g)}, c_g)$ is active in π for $1 \leq g \leq k$, and we write $\text{Act}_k(\pi)$ for the subset of $\mathcal{Y}_{k,m}$ consisting of all such triples. We say that a π -synonymity $(\rho; R, C) \underset{\pi}{\sim} (\sigma; S, D)$ is active if $(\rho; R, C)$ is an active triple. If either a triple or a π -synonymity fails to be active, we call it *inactive*.

Proposition 5.4. *For any choice of π and k , we have that*

$$\text{Can}_k(\pi) \subseteq \text{Act}_k(\pi)$$

as subsets of $\mathcal{Y}_{k,m}$.

Proof. The proof proceeds by an appropriate modification of Proposition 3.2. We take an inactive k -triple (ρ, R, C) and prove that it is not canonical.

By definition, either $\pi_{r_{\rho(g)}, c_g}$ or $\pi_{r_{\rho(g)}, c_g - 1}$ equals 1 for some $1 \leq g \leq k$, so that

$$(5.5) \quad \text{Ins}(\iota_1; r_{\rho(g)}, c_g; \pi) = \text{Ins}(\iota_1; r_{\rho(g)} + 1, c_g \pm 1; \pi).$$

We choose such a g to minimise $\rho(g)$, and then write $\mu = \text{Del}(\rho(g), g; \rho)$. From Lemmas 4.5 and 4.6, we find that $\text{Ins}(\mu; Q, B; \iota_1) = \rho$, where $Q = \rho(g)^\perp$ and $B = g^\perp$. Now we apply 4.4 to write

$$(5.6) \quad \text{Ins}(\text{Ins}(\mu; Q, B; \iota_1); R, C; \pi) = \text{Ins}(\mu; Q^\triangleright, B^\triangleright; \text{Ins}(\iota_1; R^\triangleright, C^\triangleright; \pi));$$

but $R^\triangleright = r_{\rho(g)}$ and $C^\triangleright = c_g$, so we may appeal to (5.5) and rewrite the right hand side of (5.6) as $\text{Ins}(\mu; Q^\triangleright, B^\triangleright; \text{Ins}(\iota_1; r_{\rho(g)} + 1, c_g \pm 1; \pi))$. Thus

$$\text{Ins}(\text{Ins}(\mu; Q, B; \iota_1); R, C; \pi) = \text{Ins}(\text{Ins}(\mu; S, T; \iota_1); U, V; \pi),$$

where $(S, U) = (Q^\triangleright, r_{\rho(g)} + 1)^\triangleleft$ and $(T, V) = (B^\triangleright, c_g \pm 1)^\triangleleft$. So by the minimality of $\rho(g)$, we deduce that U has the form $\{r_1, \dots, r_{\rho(g)-1}, r_{\rho(g)} + 1, \dots\}$. Therefore $R \prec U$, and $(\rho; R, C) \underset{\pi}{\sim} (\text{Ins}(\mu; S, T; \iota_1); U, V)$; so $(\rho; R, C)$ is not canonical, as required. \square

We remark that the situation for \mathcal{S} is more complex than that for \mathcal{Y} , and invite readers to construct examples of active triples in codimension 2 which are *not* canonical. In general, our strategy will be to develop a controlled procedure for reducing $\text{Act}_k(\pi)$ to $\text{Can}_k(\pi)$ by discarding all lefthand triples of active synonymities.

The action of cr on $\mathcal{Y}_{k,m}$ yields bijections between $\text{Can}_k(\pi)$ and $\text{Can}_k(\pi^{cr})$, and $\text{Act}_k(\pi)$ and $\text{Act}_k(\pi^{cr})$, for any permutation π . In general, however, no such bijections exist for rr or tr , and arguments by symmetry must therefore be handled with care.

We now define $a_k(\pi)$ to be the cardinality of $\text{Act}_k(\pi)$, and proceed to calculate $a_1(\pi)$ and $a_2(\pi)$ explicitly.

Proposition 5.7. *For any $\pi \in S_m$ and $m \geq 1$, we have that*

$$a_1(\pi) = m^2 + 1 \quad \text{and} \quad a_2(\pi) = \frac{1}{2}(m^4 + 2m^3 + m^2 + 6m + 2).$$

Proof. In codimension 1, we are concerned with triples $(\iota_1; r, c)$ for which (r, c) is active in π ; the formula for $a_1(\pi)$ follows at once. We note for reference that there are 2 inactive sites in each row of the grid except the last, where there are none. Likewise, there are 2 inactive sites in each column, except the rightmost and leftmost, which each have 1.

In codimension 2, we are concerned with triples $(\rho; r_1, r_2, c_1, c_2)$, where ρ is either ι_2 or η_2 . We concentrate on the choice of insertion sites. If we select any two distinct active sites, we obtain a unique active triple whenever neither of the equations $r_1 = r_2$ and $c_1 = c_2$ holds; for we insist that the two nonzero elements are inserted at the chosen sites. However, if exactly one of the equations holds then *either* choice for ρ yields an active triple, and we must count the selection once more to take account of this. For all but the last of the $m + 1$ rows, and all but the first and last of the $m + 1$ columns, this yields $\binom{m-1}{2}$ extra triples; the last row and the extreme columns yield a further $\binom{m+1}{2}$ and $2\binom{m}{2}$ respectively. If both of the above equations hold then the insertion sites coincide and we may again choose either value of ρ , thereby yielding an additional $2(m^2 + 1)$ active triples. Summing all cases gives

$$\binom{m^2 + 1}{2} + (2m - 1)\binom{m - 1}{2} + \binom{m + 1}{2} + 2\binom{m}{2} + 2(m^2 + 1),$$

which simplifies to the formula given. □

We shall consider a lower bound for $\text{Can}_k(\pi)$ in Section 9 below.

6. SUPERBLOCKS AND TRACK MATRICES

In this section we study the properties of a generic π -synonymity Ω between k -triples $(\rho; R, C)$ and $(\sigma; S, D)$, assuming throughout that $R \prec S$ and Ω has value τ . Our goal is the determination of all active synonymities, which with the aid of Proposition 5.4 will begin the computation of $t_k(\pi)$ from $a_k(\pi)$. We focus on certain combinatorial structures imposed by Ω on the matrices π and τ .

We define the *pattern* $\wp(\Omega)$ to be an ordered pair of sequences of multisets of + and – signs. The first sequence is derived by working down the $(m + 1)$ -grid, and recording a + or – whenever a row appears in R or S respectively; every term is therefore a singleton unless R and S have common or repeated elements. Since $R \prec S$, the first term to display an unequal distribution of signs will contain a majority of + signs. The second sequence is derived from the columns of the grid by considering C and D similarly. The dihedral group acts on $\wp(\Omega)$ in the obvious fashion, with rr and cr reversing the orders of the row and column sequences respectively, and tr interchanging the sequences themselves; however, we must modify the result by a final interchange of + and – signs for any d with $S^d \prec R^d$.

Now consider π , and an arbitrary entry $\pi_{p,q}$. We define the integers $e(p)$ and $f(q)$ to be the number of elements of R and C which do not exceed p and q respectively; we define $g(p)$ and $h(q)$ similarly, with reference to S and D respectively. Thus $e(p)$ and $f(q)$ record the number of + signs up to the appropriate entries in $\wp(\Omega)$, and $g(p)$ and $h(q)$ similarly record the number of – signs. We therefore have

$$\begin{aligned} r_{e(p)} \leq p < r_{e(p)+1}, & \quad c_{f(q)} \leq q < c_{f(q)+1}, \\ s_{g(p)} \leq p < s_{g(p)+1}, & \quad \text{and} \quad d_{h(q)} \leq q < d_{h(q)+1}. \end{aligned}$$

Applying (1.5), the $2k$ -multigrid $(R \cup S, C \cup D)$ partitions π into blocks

$$(6.1) \quad \begin{pmatrix} P(0, 0) & \dots & P(0, 2k) \\ \vdots & \vdots & \vdots \\ P(2k, 0) & \dots & P(2k, 2k) \end{pmatrix},$$

which we refer to as the *difference blocks* of Ω . We observe from the definitions that the values of $e(p)$, $g(p)$, $f(q)$ and $h(q)$ are constant throughout each difference block $P(a, b)$, and

satisfy

$$a = e(p) + g(p) \quad \text{and} \quad b = f(q) + h(q).$$

As a result, we may assign *difference coordinates*

$$\gamma(a) = e(p) - g(p) \quad \text{and} \quad \delta(b) = f(q) - h(q)$$

to the block $P(a, b)$, and therefore also to its entries. Note that $-k \leq \gamma(a)$, $\delta(b) \leq k$ for all $1 \leq a, b \leq 2k$, and that $\gamma(0)$, $\delta(0)$, $\gamma(2k)$ and $\delta(2k)$ are zero. Clearly $\gamma(a)$ and $\delta(b)$ may be obtained from $\wp(\Omega)$ by adding the appropriate signs of the row and column sequences respectively.

On inserting the elements of ρ and the elements of σ into (6.1), we obtain partitions τ' and τ'' of τ . Each difference block $P(a, b)$ will in general occupy distinct positions in the two partitions, which therefore overlap and interfere. The combinatorics of this relationship is fundamental to our description of Ω , and we proceed to make it more explicit. It is helpful to associate to $\pi_{p,q}$ the integers $t(p)$ and $u(q)$, defined to be the number of elements of R^+ and C^+ which do not exceed $p + g(p)$ and $q + h(q)$ respectively; we define $v(p)$ and $w(q)$ similarly, with reference to S^+ and D^+ respectively. Thus

$$\begin{aligned} r_{t(p)}^+ &\leq p + g(p) < r_{t(p)+1}^+, & c_{u(q)}^+ &\leq q + h(q) < c_{u(q)+1}^+, \\ s_{v(p)}^+ &\leq p + e(p) < s_{v(p)+1}^+, & d_{w(q)}^+ &\leq q + f(q) < d_{w(q)+1}^+. \end{aligned}$$

These integers measure the relative distribution of each multigrind in terms of the other *after* the insertions have taken place. Their importance stems from the following simple, but fundamental formulae.

Lemma 6.2 (Replication Lemma). *With the data and notation above, we have that*

$$\pi_{p,q} = \begin{cases} \sigma_{v(p),w(q)} & \text{if } s_{v(p)}^+ = p + e(p) \text{ and } d_{w(q)}^+ = q + f(q) \\ \pi_{p+\alpha(p),q+\beta(q)} & \text{if } s_{v(p)}^+ < p + e(p) < s_{v(p)+1}^+ \text{ and } d_{w(q)}^+ < q + f(q) < d_{w(q)+1}^+ \\ 0 & \text{otherwise,} \end{cases}$$

where $\alpha(p) = e(p) - v(p)$ and $\beta(q) = f(q) - w(q)$; similarly,

$$\pi_{p,q} = \begin{cases} \rho_{t(p),u(q)} & \text{if } r_{t(p)}^+ = p + g(p) \text{ and } c_{u(q)}^+ = q + h(q) \\ \pi_{p-\psi(p),q-\omega(q)} & \text{if } r_{t(p)}^+ < p + g(p) < r_{t(p)+1}^+ \text{ and } c_{u(q)}^+ < q + h(q) < c_{u(q)+1}^+ \\ 0 & \text{otherwise,} \end{cases}$$

where $\psi(p) = t(p) - g(p)$ and $\omega(q) = u(q) - h(q)$.

Proof. We apply (4.3) twice, first with respect to τ' and then with respect to τ'' . Thus $\tau'_{p+e(p),q+f(q)} = \pi_{p,q}$ and

$$\tau''_{p+e(p),q+f(q)} = \begin{cases} \sigma_{v(p),w(q)} & \text{if } s_{v(p)}^+ = p + e(p) \text{ and } d_{w(q)}^+ = q + f(q) \\ \pi_{p+e(p)-v(p),q+f(q)-w(q)} & \text{if } s_{v(p)}^+ < p + e(p) < s_{v(p)+1}^+ \text{ and } d_{w(q)}^+ < q + f(q) < d_{w(q)+1}^+ \\ 0 & \text{otherwise,} \end{cases}$$

thereby establishing the first equation. The second follows in similar style, by considering $\tau_{p+g(p),q+h(q)}$. \square

We call Lemma 6.2 the Replication Lemma because it quantifies the way in which elements of π are propagated under the influence of Ω . If we select a difference block $P(a, b)$ in τ' , the first equation describes the way in which it is partitioned into subblocks with respect to τ'' ; on the other hand, if we construe the same $P(a, b)$ as a difference block in τ'' , the second equation describes the way in which it is partitioned into subblocks with respect to τ' . We emphasise that these two partitions into subblocks always coexist, and that Lemma 6.2 describes *all* the conditions imposed on π by Ω . For obvious reasons, we refer to $\alpha(p)$ and $\beta(q)$ as the *forward shift coordinates*, and $\psi(p)$ and $\omega(q)$ as the *backward shift coordinates* of Ω . Each of them is constrained to lie between $-k$ and k , and is only defined when the appropriate strict inequalities hold.

The following properties will be useful, and are a direct consequence of the definitions.

Lemma 6.3.

- (1) If $\pi_{p,q}$ lies in the difference block $P(a, b)$, then both $\alpha(p)$ and $\psi(p)$ are defined and zero whenever $\gamma(a)$ is zero, and are either not defined or nonzero whenever $\gamma(a)$ is nonzero; similarly, both $\beta(q)$ and $\omega(q)$ are defined and zero whenever $\delta(b)$ is zero, and are either not defined or nonzero whenever $\delta(b)$ is nonzero.
- (2) For all values of p and q , we have that

$$\begin{aligned} \alpha(p - \psi(p)) &= \psi(p), & \psi(p + \alpha(p)) &= \alpha(p), \\ \beta(q - \omega(q)) &= \omega(q) & \text{and} & \quad \omega(q + \beta(q)) = \beta(q) \end{aligned}$$

whenever the relevant quantities are defined.

- (3) The values of $t(p)$, $v(p)$, $u(q)$ and $w(q)$ are nondecreasing, and each increase is by precisely 1.

We continue to investigate Ω by partitioning (6.1) into V -, W -, X - and Y -matrices according to the zeros of the difference (or the shift) coordinates, obtaining

$$(6.4) \quad \begin{pmatrix} V(0,0) & Y(0,1) & V(0,2) & Y(0,3) & \dots & V(0,2u) \\ X(1,0) & W(1,1) & X(1,2) & W(1,3) & \dots & X(1,2u) \\ V(2,0) & Y(2,1) & V(2,2) & Y(2,3) & \dots & V(2,2u) \\ X(3,0) & W(3,1) & X(3,2) & W(3,3) & \dots & X(3,2u) \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ V(2t,0) & V(2t,1) & V(2t,2) & V(2t,3) & \dots & V(2t,2u) \end{pmatrix}.$$

Each entry in (6.4) is partitioned as a maximal union of adjacent blocks $P(a, b)$ satisfying $\gamma(a) = 0$ and $\delta(b) = 0$ for each V -matrix. $\gamma(a) \neq 0$ and $\delta(b) \neq 0$ for each W -matrix, $\gamma(a) \neq 0$ and $\delta(b) = 0$ for each X -matrix, $\gamma(a) = 0$ and $\delta(b) \neq 0$ for each Y -matrix. We refer to these as the *superblocks* of Ω , and identify them as being of class V , W , X , or Y . We may apply the Replication Lemma 6.2 to characterise all superblocks of class X , Y and V .

Proposition 6.5. *All superblocks of class X and Y are zero, whereas superblocks of class V are unrestricted submatrices of π .*

Proof. Choose a superblock $P(a, b)$ of class X in π , and consider a constituent element $\pi_{p,q}$. By definition $\delta(b) = 0$, so $\beta(q)$ and $\omega(q)$ are both zero by Lemma 6.3(1), and the Replication Lemma 6.2 implies that $\pi_{p,q} = \pi_{p+\alpha(p),q}$ or 0. But $\gamma(a)$ is nonzero, so either $\alpha(p)$ is defined and nonzero or else $\pi_{p,q} = 0$, by Lemma 6.3(1) again. Since π is a permutation matrix, we deduce that $\pi_{p,q} = 0$ in all cases, and therefore the entire superblock is zero. A similar argument reversing the rôles of the coordinates applies to superblocks of class Y .

So far as superblocks of class V are concerned the Replication Lemma 6.2 identifies $\pi_{p,q}$ only with itself, and therefore imposes no restrictions. \square

We learn from Proposition 6.5 that the essential features of Ω are concentrated in the superblocks of class W , to which we now turn.

Given any m -permutation matrix ν , it is helpful to rewrite the nonzero entry $\nu_{\nu(q),q}$ as ν_q , for every $1 \leq q \leq m$. Thus in Lemma 6.2, the nonzero elements $\pi_{p+\alpha(p),q+\beta(q)}$ and $\pi_{p-\psi(p),q-\omega(q)}$ are abbreviated to $\pi_{q+\beta(q)}$ and $\pi_{q-\omega(q)}$ respectively.

Theorem 6.6. *The nonzero elements of any superblock $W(j,k)$ may be partitioned into sequences $(\pi_{q_1}, \dots, \pi_{q_\ell})$, to each of which are associated unique elements $\rho_{q'}$ in ρ and $\sigma_{q''}$ in σ ; furthermore, $\rho_{q'}$ and $\sigma_{q''}$ are each inserted at a site in $W(j,k)$ when forming τ' and τ'' , and the sequences*

$$(\rho_{q'}, \pi_{q_1}, \dots, \pi_{q_\ell}) \quad \text{and} \quad (\pi_{q_1}, \dots, \pi_{q_\ell}, \sigma_{q''})$$

become identified.

Proof. We begin by choosing any nonzero π_{q_i} in $W(j,k)$, and assuming that both

$$(6.7) \quad r_t^+ < \pi(q_i) + g < r_{t+1}^+ \quad \text{and} \quad c_u^+ < q_i + h < c_{u+1}^+$$

and

$$(6.8) \quad s_v^+ < \pi(q_i) + e < s_{v+1}^+ \quad \text{and} \quad d_w^+ < q_i + f < d_{w+1}^+$$

are true for some t, u, v , and w . We use the Replication Lemma 6.2 to define

$$\pi_{q_{i-1}} = \pi_{q_i - \omega(q_i)} \quad \text{and} \quad \pi_{q_{i+1}} = \pi_{q_i + \beta(q_i)},$$

and then to iterate the procedure on $\pi_{q_{i-1}}$ and $\pi_{q_{i+1}}$, noting from Lemma 6.3(2) that we create only two new elements at each step. The sequence cannot extend indefinitely in either direction since all our matrices are finite, and using Lemma 6.2 again we see that the only method of termination is for equalities to appear in either or both of (6.7) and (6.8). When this occurs in (6.7), we decree that we have reached π_{q_1} , which is then identified in τ' with some $\rho_{q'}$; and when it occurs in (6.8), we decree that we have reached π_{q_ℓ} , which is similarly identified in τ'' with some $\sigma_{q''}$.

We must now show that whenever π_q lies in $W(j,k)$ (and is not at the end of a sequence) then $\pi_{q-\omega(q)}$ and $\pi_{q+\beta(q)}$ also lie in $W(j,k)$; we prove the latter, and begin by assuming it to be false. Both $\alpha(p)$ and $\beta(q)$ are nonzero, and we may take them to be positive without loss of generality. Then either $r_m = s_m = p_*$ for some $p < p_* < p + \alpha(p)$ and $e(p) < m \leq k$, or $c_n = d_n = q_*$ for some $q < q_* < q + \beta(q)$ and $f(q) < n \leq k$, or both. Again without loss of generality, we assume the latter; since $d_w \leq q + \beta(q)$ by definition of $w(q)$, we deduce that $n \leq w(q)$. Thus $\beta(q) = f(q) - w(q)$ is negative, a contradiction. The proof that $\pi_{q-\omega(q)}$ also lies in $W(j,k)$ is entirely similar.

It remains to consider the sites of insertion of $\rho_{q'}$ and $\sigma_{q''}$. We consider the latter, assuming that $\sigma_{q''}$ is inserted outside $W(j,k)$. Again we may restrict attention to the column coordinate q_ℓ , noting from Lemma 6.2 that $\sigma_{q''}$ is inserted at the site (s_v, d_w) , where $d_w = q_\ell + \beta(q_\ell) - 1$. By a similar argument to the above, we deduce that $f(q_\ell) < w(q_\ell)$, and again contradict the fact that $\beta(q_\ell)$ is positive. The proof that $\rho_{q'}$ is also inserted in $W(j,k)$ is entirely similar. \square

We refer to the sequence $(\pi_{q_1}, \dots, \pi_{q_\ell})$ as a *track* of length ℓ associated to Ω , and to each pair $(\pi_{q_i}, \pi_{q_{i+1}})$ as a *segment*. The *source* and *target* of the track are the sites of insertion of the source element $\rho_{q'}$ and target element $\sigma_{q''}$ respectively.

It is obvious that any track length ℓ must satisfy $1 \leq \ell \leq m$, and that the total number of tracks associated to Ω cannot exceed k . Any ρ_x which fails to be a source element must coincide, as an entry in τ under the respective insertions, with some σ_y which fails to be a target element. They are therefore canonically paired, and each pair may be construed as defining a track of length zero. With this convention, the total number of tracks is exactly k .

We deem a segment $(\pi_{q_i}, \pi_{q_{i+1}})$ to have *type* $[y : x]$, where $y = \alpha(\pi(q_i)) = \psi(\pi(q_{i+1}))$ and $x = \beta(q_i) = \omega(q_{i+1})$. This notion is suggested by considering the matrix π as an array of dots positioned at the centres of the cells $(\pi(q), q)$, and interpreting each segments as a vector of slope y/x connecting the two dots. For small values of k the types $[y : x]$ vary gently along any track; as the codimension increases, the possibility grows of more erratic variation within the bounds

$$(6.9) \quad 1 \leq |\alpha(\pi(q_i))|, |\beta(q_i)|, |\psi(\pi(q_{i+1}))|, |\omega(q_{i+1})| \leq k.$$

The type of a track is obtained by concatenating the types of its individual segments, and quantifies aspects of its appearance within the matrix π . When consecutive segments have the same type, as often happens in longer tracks, we may use exponential notation as an abbreviation; thus, for example, $[1 : 2][1 : 2][1 : 1][1 : 1][2 : 1]$ becomes $[1 : 2]^2[1 : 1]^2[2 : 1]$, or even $1^4 2 : 2^2 1^3$. The latter highlights certain similarities with types such as $12^4 : 2^4 1 = [1 : 2][2 : 2]^3[2 : 1]$ or $1^3 2^2 : 2^3 1^2 = [1 : 2]^3[2 : 1]^2$, but obscures differences amongst individual segments.

We are now in a position to describe the global structure imposed on π by the existence of a synonymy Ω . This will help us to analyse the cases of small k in subsequent sections.

Consider the partition (6.4) of π into superblocks, and write $w(\Omega)$ for the matrix obtained by deleting all blocks of class V , X and Y . We call $w(\Omega)$ the *track matrix* of the synonymy, noting that it retains the original decomposition into constituent difference blocks. Track matrices play a major rôle in the remainder of our work, and we need to consider their basic properties. We retain the notation of Theorem 6.6.

Proposition 6.10. *For any synonymy Ω :*

- (1) $w(\Omega)$ is a u -permutation matrix for some $1 \leq u \leq m$, and its nonzero elements may be partitioned into j tracks of positive length, where $1 \leq j \leq u$;
- (2) the track which intersects the first row of $w(\Omega)$ has its sink below its source;
- (3) the permutation $w(\Omega)$ admits a synonymy whose track matrix is $w(\Omega)$ itself.

Proof. Since π is a permutation matrix, every row deleted in forming $w(\Omega)$ contains a single nonzero entry, concentrated in the blocks of class V by virtue of Proposition 6.5; the same remarks apply to the deleted columns, which are therefore equal in number. Thus $w(\Omega)$ is $u \times u$, for some $0 \leq u \leq m$, and contains d nonzero entries; it is therefore a permutation matrix. Since $u = 0$ only when the synonymy is equality, we may take $u \geq 1$. The partition of $w(\Omega)$ into tracks, of which there can be at most u , then follows from Theorem 6.6, so proving (1).

Then (2) is an immediate consequence of our insistence that $R \prec S$.

Theorem 6.6 also implies that the original π -synonymy Ω is realised on the grid of $w(\Omega)$. It may be expressed as

$$(6.11) \quad \text{Ins}(\rho^*; R^*, C^*; w(\Omega)) = \text{Ins}(\sigma^*; S^*, D^*; w(\Omega)),$$

where ρ^* and σ^* are the j -permutations obtained from ρ and σ by restriction to the sources and targets of the tracks in $w(\Omega)$, and (R^*, C^*) and (S^*, D^*) are the $(u + 1)$ -multigrids

obtained from (R, C) and (S, D) by deleting the appropriate sites. The result is therefore a $w(\Omega)$ -synonymy with track matrix $w(\Omega)$, confirming (3). \square

When we consider $w(\Omega)$ as a permutation matrix in its own right, the shift coordinates of the synonymy described by Proposition 6.10(3) are, of course, in agreement with those of Ω for the corresponding rows and columns of π ; we shall use this fact without further comment below, since it is an important aid to simplifying notation.

Any u -permutation which is the track matrix of some synonymy (not necessarily of dimension m or codimension k) we call a j -track matrix of dimension u , where j is the number of tracks. When u is large compared with j such matrices are rare in S_u , and exhibit distinctive distributions of nonzero elements. We defer examples until the next two sections, where the structure of 1- and 2-track matrices will be determined.

Observe that a matrix may occur simultaneously as both an i -track matrix and a j -track matrix for distinct values of i and j ; for example ι_2 is both the 1-track matrix of the ι_2 -synonymy $(\iota_1; 0, 0) \sim (\iota_1; 2, 2)$, and the 2-track matrix of the ι_2 -synonymy $(\eta_2; 0, 0, 2, 2) \sim (\eta_2; 2, 2, 0, 0)$.

We say that a j -track matrix ζ is *composite* if it admits a nontrivial standard partitioning into blocks which are either zero, or else are themselves track matrices; we insist that each of the j tracks of ζ appears in its entirety in one of the constituent blocks. If ζ is not composite, we call it *prime*. Thus every composite ζ has a unique primary decomposition into prime factors $\zeta(i)$, each of which is a j_i -track matrix for some j_i satisfying $\sum_i j_i = j$.

For all symmetries d in D_4 , we note that $w(\Omega)^d$ agrees with $w(\Omega^d)$ and that ζ^d has prime factors $\zeta(i)^d$ for $1 \leq i \leq t$.

7. CODIMENSION 1

We now reveal the structure of \mathcal{S} in codimension 1. The arguments are simple and explicit, but we use the methodology we have introduced above in order to prepare the ground for subsequent sections. We continue to work with a generic m -permutation π .

We begin by considering a typical π -synonymy Ω , which in this case is between triples of the form $(\iota_1; r, c)$ and $(\iota_1; s, d)$, with $r < s$; we may assume that it is nondegenerate, since if $r = s$ (for example), then $c = d$ by Proposition 5.3 and Ω is the identity. The pattern of Ω will therefore be one of

$$(7.1) \quad (+-, +-), \quad (+-, -+),$$

which differ by the action of cr in D_4 .

We shall find it helpful to express track matrices diagrammatically, using lines of appropriate slope to indicate tracks of nonzero elements, and arrows to indicate their direction from source to target.

Proposition 7.2. *The 1-track matrices corresponding to the patterns 7.1 are*

$$\left(\begin{array}{c} \searrow \\ \swarrow \end{array} \right) \quad \text{and} \quad \left(\begin{array}{c} \swarrow \\ \searrow \end{array} \right)$$

respectively.

Proof. By Proposition 6.10, we have that $w(\Omega)$ is a u -permutation matrix (for some $1 \leq u \leq m$) whose elements make up a single track. For the pattern $(+-, +-)$, we deduce that the forward shift coordinates $\alpha(p)$ and $\beta(q)$ are both 1 for $1 \leq p, q < u$; so the track has type

$[1 : 1]^u$. Similarly, for $(+ -, - +)$ we have that $\alpha(p)$ is 1 for $1 \leq p < u$ and $\beta(q)$ is -1 for $1 < q \leq u$; so the track has type $[1 : -1]^u$. The matrices are therefore as indicated. \square

Corollary 7.3. *For any m -permutation π , all active 1-triples are canonical.*

Proof. Suppose that $(\iota_1; r, c)$ is noncanonical, so that a synonymy $(\iota_1; r, c) \sim (\iota_1; s, d)$ exists with $r < s$. Its track matrix is described by Proposition 7.2, ensuring that $\pi_{r,c} = 1$ or $\pi_{r,c-1} = 1$ respectively. Therefore $(\iota_1; r, c)$ is inactive, as required. \square

The proof of Corollary 7.3 illustrates how we may often use identical arguments to deal with the synonymies Ω and Ω^{cr} . It is analogous to the second half of the proof of Proposition 3.2; but Corollary 7.3 is false in all higher codimensions.

Theorem 7.4. *For any m -permutation π in \mathcal{S} , we have that*

$$t_1(\pi) = m^2 + 1$$

Proof. Corollary 7.3 combines with Proposition 5.4 to yield $\text{Act}_1(\pi) = \text{Can}_1(\pi)$. The formula follows from Proposition 5.7. \square

We conclude this section with remarks about the downward structure of \mathcal{S} in codimension 1. If π is partitioned into blocks so that each block is either ι_k , η_k or zero, we call this a *partition of π into entities*. If a partition into entities is not a strict refinement of some other partition into entities, we call it *maximal*. We invite the reader to verify that every permutation matrix π admits a unique maximal partition $E(\pi)$ into entities, and that there is an obvious bijection between the nonzero blocks of $E(\pi)$ and the ideal $U_1(\pi)$. Clearly the number of such blocks is between 1 and m ; we therefore assert that $1 \leq u_1(\pi) \leq m$. It is easy to see that these bounds are attained when $m > 3$.

8. CODIMENSION 2

We begin our analysis of the codimension 2 case by following the lead of (7.1), and classifying the patterns of codimension 2 π -synonymies Ω for an m -permutation π .

We assume initially that such a synonymy $(\rho; R, C) \underset{\pi}{\sim} (\sigma; S, D)$ is nondegenerate, in that no element of R , C , S , and D is repeated, and neither R and S nor C and D have any element in common. Then $\wp(\Omega)$ is one of

$$(8.1) \quad \begin{array}{ll} (+ - + -, + - + -) & \text{(A1)}, & (+ - + -, + + --) & \text{(B1)}, \\ (+ - - +, + - + -) & \text{(A2)}, & (+ - - +, + + --) & \text{(B2)}, \\ (+ - - +, - + + -) & \text{(A3)}, & & \\ (+ - - +, + - - +) & \text{(A4)}, & (+ + --, + + --) & \text{(C)}, \end{array}$$

or their *translates* by d in D_4 ; as in codimension 1, the action of d must be followed by an interchange of signs whenever $S^d \prec R^d$. We note that each of the patterns (A) has a *caesura* in the form of a zero shift coordinate between the middle symbols in both rows and columns. The patterns (B) have a caesura in the row direction only, whilst (C) have a caesura in neither direction.

As explained in Section 5, our goal is to reduce $\mathcal{Y}_{2,m}$ to $\text{Can}_2(\pi)$ by discarding the lefthand triples of nontrivial synonymies (which are therefore noncanonical). We use the classification (8.1) to progress from simpler cases to the more complex, recording such synonymies as we find them, and taking care to avoid those whose lefthand triples have already been identified.

The procedure may be made algorithmic, but we content ourselves with enumerating the results.

We began this process in Proposition 5.4, by proving that an inactive triple is always lefthand. The track matrix of the relevant synonymy is 1×1 , and consists of a single 1. In codimension 1, Corollary 7.3 confirms that no further lefthand triples are associated with 1-tracks of greater length, and we extend this to codimension 2 and patterns (A) in Corollary 8.3 below. It is precisely because these simplest synonymies account for almost all noncanonical triples that we were led to introduce the subset $\text{Act}_2(\pi) \subseteq \mathcal{T}_{2,m}$ in Section 5. The situation for patterns (B) is more delicate, and we show in Corollary 8.8 that all but $m - 1$ of the associated lefthand triples have already been identified by the restriction to $\text{Act}_2(\pi)$. Pattern (C) is the most complex case, in which *every* synonymy involves a previously unrecorded lefthand triple. Although small, we shall see in Corollary 8.14 that the number of these varies between 0 and $m - 1$, depending on the permutation π . We deal with degenerate versions of all patterns (8.1) as they arise.

For each pattern we begin by studying the relevant 2-track matrices, which we assume to be $u \times u$. We determine the restrictions imposed by the patterns on the types of their constituent tracks, and consider the choices available for the inserted 2-permutations; at most four are possible, as ρ and σ vary over ι_2 and η_2 . After each Proposition we provide illustrative examples, and encourage readers to sketch the associated matrices as an aid to maximising understanding.

Proposition 8.2. *The 2-track matrices corresponding to pattern (A) synonymies consist of four difference blocks, and are composite; they have two prime factors, both of which are entities.*

Proof. Given any such synonymy Ω , the presence of the caesurae ensures that $w(\Omega)$ consists of four blocks whose difference coordinates are ± 1 in every case (or zero in the case of a degeneracy). The forward shift coordinate α always takes the value 1 in the upper two blocks, and ± 1 in the lower two blocks, whereas β is ± 1 in the left pair and the right. The upper track therefore has type $[1:\pm 1]^j$ (where j may be zero or u in the case of a degeneracy), and the lower track has type $[\pm 1:\pm 1]^{u-j}$. Thus $w(\Omega)$ is composite, and has the stated structure; any of the four possible combinations for ρ and σ may occur. \square

Corollary 8.3. *There are no active synonymies of patterns (A) or their translates for any m -permutation π .*

Proof. This follows by applying the arguments of Corollary 7.3 to the upper track. \square

The proof of Proposition 8.2 notes how degeneracies such as $\{+-\} + -$ simply reduce the number of tracks to one (or zero!). The same is true for degeneracies such as $+ \{-+\} -$; but in the columns of (A1), the latter coincides with the degeneracy $+ \{+-\} -$ in the columns of (B1). We have therefore already taken care of *all* such degenerate examples of patterns (B), and so may insist that $c_2 < d_1$ in our discussion of (B1) and (B2) below. The discussion is unaffected by remaining degeneracies such as $c_1 = c_2$, as can be seen in Example 8.6 below.

Example 8.4. A (4123)-synonymy of pattern (A1) is given by


$$(\eta_2; \{1, 4\}, \{1, 2\}) \sim (\eta_2; \{4, 5\}, \{2, 5\}),$$

since the insertions have common value $(\overline{5}\underline{6}\overline{1}23\underline{4})$ (where the left- and righthand insertions are over- and underlined respectively). The track matrix is (4123), and the upper and lower tracks


have types $[1:1]^2$ and $[1:1]^0$ respectively. The pattern $(+\{-+\}-, +\{-+\}-)$ is degenerate, and the synonymy is inactive. The righthand triple is canonical.

To continue our analysis of patterns (B) and (C), we generalise the diagrammatic portrayal of track matrices introduced in Section 7. We may assume by Proposition 6.10 that the uppermost track points *downwards*. We may further assume that it points *rightwards* by applying cr whenever necessary; Proposition 7.2 exemplifies the ease with which proofs may be adapted to allow for this symmetry.

Proposition 8.5. *The 2-track matrices corresponding to pattern (B) synonymities consist of six difference blocks and are prime; those corresponding to (B1) have one of the forms*



and those corresponding to (B2) have one of the forms



or their translates by rr .

Proof. Given any such synonymity Ω , the presence of the caesura ensures that $w(\Omega)$ consists of six difference blocks $P(a, b)$.

For pattern (B1) these have $\gamma(a) = 1$, and $\delta(b) = 1, 2$, and 1 from left to right; the condition $c_2 < d_1$ ensures that the block with $\delta(b) = 2$ is nonempty, and therefore that $c_2 < d_1 - 1$. The forward shift coordinate α always has value 1 , whereas β takes values $1, 2$ and 1 from left to right, thereby creating segments of type $[1:1]$ and $[1:2]$. All four combinations for ρ and σ are possible within these constraints, and the first four diagrams correspond to the pairs (ι_2, ι_2) , (η_2, η_2) , (η_2, ι_2) , and (ι_2, η_2) respectively. Each case is prime.

For pattern (B2), the difference blocks have $\gamma(a) = 1$ and -1 , with $\delta(b)$ as before; again $c_2 < d_1 - 1$. In this case α has value 1 in the upper three blocks and -1 in the lower, with β as before, thereby creating segments of types $[1:1]$ and $[1:2]$ for the upper chain and of types $[-1:1]$ and $[-1:2]$ for the lower chain. All four combinations for ρ and σ are again possible, and the fifth and sixth diagrams correspond to the pairs (ι_2, ι_2) and (η_2, ι_2) ; their translates under rr correspond to (η_2, η_2) and (ι_2, η_2) respectively. Again, each case is prime. \square

Translate of patterns (B1) and (B2) produces similar diagrams. We note in particular that the parallel segments of the first four diagrams have type $[1:-2]$ after cr , $[2:1]$ after tr and $[-1, 2]$ after rr , and that the wedge of converging arrows in the final two diagrams is reorientated by both cr and tr .

Care is required to interpret the diagrams of Proposition 8.5 correctly with respect to the fringe effects at the boundaries of the difference blocks. Thus an arrow of type $[1:2]$ signifies a track which begins in a block with upper left corner either $\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}$ or $\begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$. These two possibilities occur in pairs, one above the other, wherever the tracks of type $[1:2]$ are paired in the diagrams; the order in which they occur is determined by the choice of insertion. Similar remarks apply to the tracks of type $[1:-2]$ and $[\pm 2:1]$.

Example 8.6. A (415263)-synonymity of pattern (B1) is given by

$$(\eta_2; \{1, 4\}, \{1, 1\}) \sim (\eta_2; \{4, 7\}, \{7, 7\}),$$

since the insertions have common value $(\overline{51627384})$. The track matrix is (415263), and the upper and lower tracks both have types $[1 : 2]^2$. The pattern $(+\{-+\}-, \{++\}\{-\})$ is degenerate, and the source of the lower track is active; even so, the source of the upper is inactive, and so is the synonymity. The righthand triple is canonical.

Example 8.7. A (7651432)-synonymity of pattern (B2) is given by

$$(\eta_2; \{1, 8\}, \{1, 3\}) \sim (\eta_2; \{2, 4\}, \{6, 6\}),$$

since the insertions have common value $(\overline{987162543})$. The track matrix is (7651432), and the upper and lower tracks have types $[1 : 2]^0$ and $[-1 : 1]^2[-1 : 2][-1 : 1]^2$ respectively. The synonymity is active, and the righthand triple is canonical.

We require the following corollary to Proposition 8.5.

Corollary 8.8. *There are no active synonymities of pattern (B1) or its translates; synonymities of pattern (B2) and its translates yield exactly $m - 1$ active lefthand triples for each m -permutation π .*

Proof. Since neither track of an active synonymity can begin with a segment of type $[1 : 1]$, the four diagrams of pattern (B1) in Proposition 8.5 are inactive unless $c_1 = c_2$. Even with this degeneracy, one of the two tracks containing segments of type $[1 : 2]$ must begin with the submatrix $\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}$, ensuring that the lefthand triple is always inactive. A similar argument applies equally well after the application of cr . Applying rr reorients both tracks upwards, and must be followed by an interchange of signs; the results are the same as applying cr , up to reordering the diagrams. Applying tr creates vertically oriented tracks, for which insertion at an inactive site is again unavoidable.

Turning to patterns (B2), the fifth diagram of Proposition 8.5 is inactive unless $c_1 = c_2 = q$ for some $1 \leq q < m - 1$ (where $q \neq m - 1$ because $c_2 < d_1 - 1$). If $\pi(q) < \pi(q + 1)$, the upper track must pass through π_q (and the lower track through π_{q+1}), thereby locating its source at π_q , which is inactive. If $\pi(q) > \pi(q + 1)$, then the diagram corresponds to the synonymity

$$(8.9) \quad (\eta_2; \pi(q + 1), \pi(q) + 1, q, q) \underset{\pi}{\sim} (\eta_2; \pi(q + 1) + 1, \pi(q), q + 2, q + 2),$$

in which the source π_{q+1} of the upper track is clearly active. If the 1 in row $\pi(q) + 1$ does not lie in column $q - 1$, then the source $(\pi(q) + 1, q)$ of the lower track is also active. If the 1 does lie there, it must be the uppermost element of a maximal antiidentity η_i for some $1 \leq i \leq q - 1$, and the associated synonymity

$$(8.10) \quad (\eta_2; \pi(q + 1), \pi(q) + i + 1, q - i, q) \underset{\pi}{\sim} (\eta_2; \pi(q + 1) + 1, \pi(q), q + 2, q + 2)$$

is active, and corresponds to the *sixth* diagram of pattern (B2). In either case, we have located an active left hand triple.

Applying rr to these arguments merely interchanges the role of the fifth and sixth diagrams, and reveals no new synonymities. There are no active synonymities at all under tr , for similar reasons to those for pattern (B1).

If we apply cr , however, we find synonymities in which the convergent tracks point to the left, and η_2 and η_i are replaced by ι_2 and ι_i respectively. These are active and lefthand when $\pi(q) < \pi(q + 1)$, and inactive when $\pi(q) > \pi(q + 1)$. Each pair of adjacent columns of π

therefore corresponds to precisely one active lefthand triple, and there are $m - 1$ such pairs in all. \square

The righthand triple of (8.9) is only canonical when both tracks terminate in column $q + 1$; otherwise it doubles as the active *lefthand* triple of the synonymymity associated with the adjacent columns $q + 1$ and $q + 2$! In the latter situation, the upper and lower tracks have at least one segment of type $[1:2]$ and $[-1:2]$ respectively. Similar remarks apply to (8.10). It is instructive to reexamine Examples 8.6 and 8.7 in the light of this proof.

The existence of the active synonymymities identified by Corollary 8.8 depends on the horizontal wedge of converging arrows in the fifth and sixth diagrams. The analysis of pattern (C) synonymymities is quite different because both tracks are oriented downwards.

We consider π synonymymities $(\rho; R, C) \sim (\sigma; S, D)$ where $r_1 \leq r_2 < s_1 \leq s_2$ and $c_1 \leq c_2 < d_1 \leq d_2$, noting that the central inequalities are strict because a degeneracy returns us to patterns (B). The absence of caesurae ensures that the track matrix is a contiguous submatrix of π , consisting of nine difference blocks in which the shift coordinates α and β take values 1, 2 and 1 as we pass down the three rows and along the three columns respectively. For convenience, we label the heights of the blocks by $h_1 = r_2 - r_1$, $h_2 = s_1 - r_2$ and $h_3 = s_2 - s_1$, and the widths by $w_1 = c_2 - c_1$, $w_2 = d_1 - c_2$ and $w_3 = d_2 - d_1$. Degeneracies in which one or more of h_1 , h_3 , w_1 and w_3 are zero do not affect the arguments, as can be seen in Example 8.12 below.

Our description of the associated chain matrices breaks with earlier practice, and considers only those corresponding to *active* synonymymities. We encourage readers to sketch the diagrams for themselves, armed with the philosophy that pattern (C) displays codimension 2 phenomena in their purest form.

Proposition 8.11. *The 2-track matrices of active synonymymities of pattern (C) are prime; the synonymymities take one of the forms*

- (1) $(\eta_2; R, C) \sim (\eta_2; S, D)$ with tracks of type $1^{h_1}2^{h-1}:2^{w-1}1^{w_3}$ and $2^{h-1}1^{h_3}:1^{w_1}2^{w-1}$,
- (2) $(\eta_2; R, C) \sim (\iota_2; S, D)$ with tracks of type $1^{h_1}2^{h-1}:2^{w-1}$ and $2^{h-1}1^{h_3}:1^{w_1}2^{w_1}1^{w_3}$,

or the transpose of (2), where $h_2 = 2h$ and $w_2 = 2w$ in (1), $h_2 = 2h$ and $w_2 = 2w + 1$ in (2), and $h_2 = 2h + 1$ and $w_2 = 2w$ in its transpose.

Proof. It follows from the Replication Lemma 6.2 that $\pi_{r_1, c_1} = \rho_{1,1}$, and we deduce that $\rho = \eta_2$ if the synonymymity is active. This ensures that the sources of the upper and lower tracks are (r_1, c_2) and (r_2, c_1) respectively. We consider the two possibilities for σ .

When $\sigma = \eta_2$, the upper and lower tracks have targets (s_1, d_2) and (s_2, d_1) respectively, because the tracks cannot cross. The restrictions on α and β then ensure that the upper track has type $1^{h_1}2^{h-1}:2^{w-1}1^{w_3}$, and the lower has type $2^{h-1}1^{h_3}:1^{w_1}2^{w-1}$. The result is prime.

When $\sigma = \iota_2$ we must consider two further cases, distinguished by the upper track having target (s_1, d_1) or (s_2, d_2) . In the former situation the central difference blocks has height $2h$ and width $2w + 1$, so that the upper track has type $1^{h_1}2^{h-1}:2^{w-1}$ and the lower has type $2^{h-1}1^{h_3}:1^{w_1}2^{w_1}1^{w_3}$. The latter situation is similar, and is obtained by transposition; the upper track has type $1^{h_1}2^{h_1}1^{h_3}:2^{w-1}1^{w_3}$ and the lower has type $2^{h-1}:1^{w_1}2^{w-1}$. Both are prime. \square

Translations of active pattern (C) synonymymities by cr are also active (using identical arguments), but translations by rr have both tracks oriented upwards, and are always inactive.

In using our condensed notation for track types, we note that the sums of the exponents on either side of the $:$ sign must agree. Proposition 8.11 therefore imposes constraints such as $h_1 - w = w_3 - h$ in (1), $w - h_1 = h$ in (2), and $h - w_1 = w$ in its transpose. The

exact configuration of the individual segments of the two tracks depends on value of these quantities. In (1), for instance, the upper track has type

$$\begin{aligned} [1:2]^{h_1} [2:2]^{w-h_1-1} [2:1]^{w_3} & \quad \text{if } h_1 - w < 0, \\ [1:2]^{h_1-1} [1:1] [2:1]^{w_3-1} & \quad \text{if } h_1 - w = 0, \\ [1:2]^{w-1} [1:1]^{h_1-w+1} [2:1]^{h-1} & \quad \text{if } h_1 - w > 0. \end{aligned}$$

We have no need to elaborate on these details, but hope that our next two examples will imbue readers with sufficient confidence to follow the analysis to its conclusion.

We extend our diagrammatic conventions by marking segments of type $[2:2]$ with a double arrowhead, and creating extra symbols by writing $k < j' < k'$ for any $1 \leq j < k \leq 9$.

Example 8.12. An active (61823457)-synonymity of pattern (C) is given by

$$(\eta_2; \{1, 5\}, \{1, 1\}) \sim (\eta_2; \{9, 9\}, \{5, 9\}),$$


since the insertions have common value $(\overline{61}821'34579)$. The upper and lower tracks have types $[1:2][1:1]^3[2:1]$ and $[2:2]$ respectively. The pattern $(++\{--\}, \{++\}--)$ is degenerate, and the righthand triple is canonical.

Example 8.13. An active (3'2'1'864927153)-synonymity is given by

$$(\iota_2; \{1, 2\}, \{12, 13\}) \sim (\eta_2; \{11, 12\}, \{2, 6\}),$$

since the insertions have common value $(5'4'3'1'862'4927\overline{153})$. The upper and lower tracks have types $[1:-2][2:-2][2:-1]^3[1:-1]$ and $[2:-1][2:-2]^2$ respectively. The pattern $(++--, --++)$ is the translate of (C) by cr , and the righthand triple is not canonical.

The track matrices for these examples have diagrams



illustrating the downward pointing, lens-shaped configurations which characterise active synonymities of pattern (C). It is worth comparing these with the corresponding permutation matrices

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

We require the following corollary to Proposition 8.11.

Corollary 8.14. *Synonymities of pattern (C) and its translates yield j active lefthand triples for each m -permutation π , where j depends on π and satisfies $0 \leq j \leq m - 1$.*

Proof. The initial segment of the upper track of any such synonymity Ω has type $[1:\pm 2]$.

We begin with the case $[1:2]$, and consider the requirements for a particular π_q to be the source element of an upper track. If it is, then $r_1 = \pi(q)$ and $c_2 = q - 1$. Following the track in the manner of the Replication Lemma 6.2, we must observe a sequence of 1s forming

segments of type $[1:2]$ (possibly changing to $[1:1]$) until the shift coordinate α increases from 1 to 2. This determines r_2 and c_1 , since the row skipped at the changeover is $r_2 + 1$, and $r_2 + 1 = \pi(c_1)$; we may therefore read off the entire lefthand triple from π_q , and determine Ω by inspection. We must also have that $\pi(q - 1) > \pi(q)$, since π_{q-1} will belong to the lower track.

The case $[1:2]$ corresponds to the situation $\pi(q - 1) < \pi(q)$, in which π_{q-1} is the potential source element; we repeat the procedure above, under the action of the symmetry cr . In either case the requirements for a track may fail at any stage, so that columns $q - 1$ and q yield no such Ω . We therefore have at most $m - 1$ active lefthand triples in all. \square

Example 8.15. To illustrate this proof, we identify the active synonymities of pattern (C) for the permutation $\pi = (2413)$. We deal with each adjacent pair of columns in turn, and find two synonymities in all.

We note first that $\pi(1) < \pi(2)$, so the candidate for the initial segment of an upper track has type $[1:-2]$, with source $(2, 3)$. Attempting to follow the track to the left we find no more 1s, forcing the target to be $(4, 1)$. This involves skipping row 3, whose 1 in column 4 therefore lies in the lower track. The source of this track must be $(2, 5)$, and attempting to follow *this* track to the left violates the constraints of Proposition 8.11.

Secondly, we have that $\pi(2) > \pi(3)$, so the candidate for the initial segment has type $[1:2]$, with source $(1, 2)$. Again there is an immediate change from shift coordinate 1 to shift coordinate 2, so we look for a lower track with source $(1, 1)$. This leads us to a 2-track matrix (π itself, in fact) which corresponds to the active synonymity

$$(\eta_2; \{1, 1\}, \{1, 2\}) \sim (\eta_2; \{5, 5\}, \{4, 5\}).$$

Thirdly, we have that $\pi(3) < \pi(4)$, so the initial candidate is $[1:-2]$, with source $(1, 5)$. There are now two segments before a row is skipped, namely row 3; as in the first case, the lower source must therefore be $(2, 5)$. Once more π is revealed as a 2-track matrix, and corresponds to the active synonymity

$$(\iota_2; \{1, 2\}, \{5, 5\}) \sim (\iota_2; \{4, 5\}, \{1, 1\}).$$

In example 8.15, π functions as a 2-track matrix in two distinct ways; this is unusual, and we believe the case to be unique up to degeneracies. Although 2-track matrices of pattern (C) are highly constrained, however, they are not uncommon, as shown by the example ι_2 , for which $h = w = 1$ and $h_1 = h_3 = w_1 = w_3 = 0$.

We may now prove our main result.

Theorem 8.16. *For each m -permutation π , we have that $t_2(\pi)$ is given by*

$$\frac{1}{2}(m^4 + 2m^3 + m^2 + 4m + 4 - 2j)$$

for some $0 \leq j \leq m - 1$; moreover, if $m \geq 5$ then every possible value of j is realised by some choice of π .

Proof. In Corollaries 8.8 and 8.14 we have enumerated triples in $\text{Act}_k(\pi)$ which are not canonical. The formula then follows by modification of Proposition 5.7.

To conclude, we must exhibit m -permutations realising the promised values of j ; since our construction is inductive we begin with the case $m = 5$. We may readily check that the permutations ι_5 , (14523) , (14253) , (41253) , and (25314) have j values 4, 3, 2, 1, and 0 respectively, and we note in addition that the first three of these matrices have 1 in position $(1, 1)$. By our remarks above, we may therefore construct 6-permutations with j values 5, 4,

and 3, and with 1 in the $(1, 1)$ position. If we similarly enlarge the last two 5-permutations, we obtain (152364) and (136425) , whose j values may be checked to remain unaltered at 1 and 0. Rotating the second of these through $3\pi/2$ (via the symmetry $pr = tr \cdot rr$) yields (364251) with j value 0, which is unchanged by the insertion of 1 at site $(1, 1)$. A final check reveals that (152634) has j value 2.

We may therefore express our inductive assumption in two parts; firstly, that for each $m \geq 7$ and $1 \leq i \leq m - 1$ there is an $(m - 1)$ -permutation $\pi(m - 1, i)$ whose j value is i and which has 1 in the $(1, 1)$ position, and secondly, that there is an $(m - 1)$ -permutation $\pi(m - 1, 0)$ whose j value is 0, which has 1 in the $(1, 1)$ position, and whose j value is unaltered by applying pr and inserting 1 at site $(1, 1)$. We complete the inductive step by defining

$$\pi(m, i + 1) = \text{Ins}(1; 1, 1; \pi(m, i)) \quad \text{if } 0 \leq i \leq m,$$

and $\pi(m, 0) = \text{Ins}(1; 1, 1; \pi(m, 0)^{pr})$, needing only to confirm that the j value of $\pi(m, 0)$ remains 0 after applying pr and inserting 1 at $(1, 1)$. This fact follows by inspection of the original case (136425) (when $m = 6$), and by recording the spacing of the nonzero elements which are inserted at each subsequent step. \square

We note that when $m = 2$ or 3, then j takes the single value 1 or 2 respectively, and that when $m = 4$, then j takes one of the values 2 or 3 (as supported by Example 8.15).

9. THE GENERAL CASE

We conclude with a few general remarks about the codimension k structure of \mathcal{S} when $k > 2$, establishing numerical bounds of the form $h_k(\pi) \leq t_k(\pi) \leq a_k(\pi)$ and applying the Robinson-Schensted correspondence to give a more accurate prescription for $t_k(\iota_m)$.

In Proposition 5.7 we obtained an explicit formula for $a_2(\pi)$. We shall not attempt to describe $a_k(\pi)$ in such detail for any larger k , but there are still two important points to make. The first is motivated by the fact that both $a_1(\pi)$ and $a_2(\pi)$ depend only on m .

Lemma 9.1. *For all m -permutations π , we have that $a_k(\pi) = a_k(\iota_m)$.*

Proof. By definition, the matrix π is obtained from ι_m by permuting the rows. This permutation acts on the horizontal lines of the grid $G(\iota_m)$ (line p is mapped to line $\pi(p)$ for $p \leq m$, and line $m + 1$ is fixed), carrying it to $G(\pi)$ in such a way as to map inactive sites to inactive sites. Each ι_m -active triple $(\rho; R, C)$ is therefore transformed into a π -active triple by the action induced on R , setting up a bijection between $\text{Act}_k(\iota_m)$ and $\text{Act}_k(\pi)$. \square

Our second point concerns $a_k(\pi)$ for $k > 2$.

Proposition 9.2. *For any m -permutation π , the value of $a_k(\pi)$ is polynomial in m of the form*

$$\frac{1}{k!} (m^{2k} + k(k-1)m^{2k-1})$$

modulo lower powers of m .

Proof. We follow the proof of Proposition 5.7 precisely, considering k -triples $(\rho; R, C)$. If we select any k -element subset of the $m^2 + 1$ active sites, we obtain a unique active triple whenever neither R nor C contains repeated elements; for we insist that the k nonzero elements are inserted at the chosen sites. However, if exactly one row or exactly one column is repeated, then inserting the corresponding adjacent rows or columns of π in either order yields an

active triple, and we must count the selection once more to take account of this. We have then enumerated

$$(9.3) \quad \binom{m^2 + 1}{k} + (2m - 1) \binom{m - 1}{2} \binom{m^2 - 1}{k - 2}$$

active triples, and may incorporate the remainder by similar (but increasingly complicated) considerations. At each stage a polynomial function of m is added, and none has degree exceeding $2k - 2$. Our result then follows by evaluating the two terms of highest degree in (9.3). \square

It is important to compare this formula with the cruder bound obtained in (5.1). Here too the coefficients of the powers of m are functions of k , and we may only use Proposition 9.2 to approximate $a_k(\pi)$ when similar restrictions are placed on the value of k with respect to m .

We now turn to the problem of bounding $\text{Can}_k(\pi)$ from below, assuming from the start that $2k^2 \leq m$.

For each $1 \leq h \leq m$, we define the subset $E_h(\pi) \subset G(\pi)$ to consist of those sites (i, j) which satisfy

$$(9.4) \quad \pi(h) - k + 1 \leq i \leq \pi(h) \quad \text{and} \quad h - k + 1 \leq j \leq h + k;$$

these are the $2k^2$ sites (or less, near the boundary) which are closest to, and above, the nonzero element $\pi_{\pi(g),g}$. If the nonzero elements of π are widely scattered then the sets $E_h(\pi)$ will be disjoint, and their union $E(\pi)$ will contain $2k^2$ sites in each row and column of $G(\pi)$ which is sufficiently far from the boundary. If, on the other hand, the nonzero elements of π occur in clusters (as in the identity, for example), then the sets $E_h(\pi)$ will overlap, and $E(\pi)$ will be of lesser cardinality. In either event, we refer to the sites in $E(\pi)$ as *excluded*.

We label a π -active triple $(\rho; R, C)$ as π -*hyperactive* (or *hyperactive*, in the usual situation where π is understood) whenever there is at most one value of g , with $1 \leq g \leq k$, for which the site $(r_{\rho(g)}, c_g)$ is excluded. We write $\text{Hyp}_k(\pi)$ for the subset of $\text{Act}_k(\pi)$ consisting of all hyperactive triples, and denote its cardinality by $h_k(\pi)$.

Proposition 9.5. *For any choice of π , we have that*

$$\text{Hyp}_k(\pi) \subseteq \text{Can}_k(\pi)$$

as subsets of $\mathcal{T}_{k,m}$.

Proof. Let $(\rho; R, C)$ be a π -hyperactive triple. We assume that it takes part in some synonymy Ω , whose tracks must have segments in which the absolute value of each component is less than or equal to k by appeal to (6.9). At most one element of ρ can be a source, namely ρ_g where $(r_{\rho(g)}, c_g)$ is the single permissible excluded site. Therefore $w(\Omega)$ is either empty (so Ω is trivial), or else consists of a single track, and is an entity by Proposition 6.10. In other words, the site $(r_{\rho(g)}, c_g)$ must have the form $(\pi(j) + \epsilon, j + \delta)$ for some $1 \leq j \leq m$ and $0 \leq \epsilon, \delta \leq 1$; but $(\rho; R, C)$ is active by definition, so the only two possibilities are those with $\epsilon = 1$. In both these cases Ω decreases R lexicographically, and $(\rho; R, C)$ must be canonical, as required. \square

By construction, $\text{Hyp}_1(\pi)$ and $\text{Act}_1(\pi)$ coincide for any m -permutation π ; so when k is 1, Proposition 9.5 may be considered to be a restatement of Corollary 7.3.

In order to obtain a significant estimate for $h_k(\pi)$, it is helpful to introduce two distinct procedures for enlarging the set of excluded sites. They both depend on our initial assumption that $2k^2 \leq m$, and increase the cardinality of $E(\pi)$ to $2k^2m$. For the first procedure

we adjoin excluded sites to each row of $G(\pi)$ in turn, until every row except the last contains precisely $2k^2$ elements; we leave the last row untouched, with no excluded sites. For the second procedure we adjoin excluded sites to each nonextremal column of $G(\pi)$ in turn, until every such column contains exactly $2k^2$ elements; we share the remaining $2k^2$ extremal sites between the first and last columns (equally, if we enjoy symmetry). In either case we may make the adjunctions at random, or else more systematically by using the lexicographic ordering of the sites.

Proposition 9.6. *For any m -permutation π , the value of $h_k(\pi)$ is bounded below by a polynomial of the form*

$$\frac{1}{k!}(m^{2k} + k(k-1)m^{2k-1})$$

modulo lower powers of m .

Proof. We define certain disjoint sets H_ℓ of hyperactive triples $(\rho; R, C)$, for $1 \leq \ell \leq 4$.

To construct H_1 and H_2 , we first enlarge $E(\pi)$ by either of the above procedures. For each triple in H_1 we then select a k -element subset of the $m^2 - 2(k^2 - 1)m + 1$ nonexcluded sites; if neither R nor C contains repetitions, we fix ρ by insisting that its nonzero entries are inserted at the selected sites, otherwise we choose a permissible ρ at random. For each triple in H_2 we select a single active but excluded site together with a $(k-1)$ -element subset of the nonexcluded sites, and choose ρ by the same process as before. The cardinality of $H_1 \cup H_2$ is therefore

$$(9.7) \quad \binom{m^2 - 2(k^2 - 1)m + 1}{k} + 2(k^2 - 1)m \binom{m^2 - 2(k^2 - 1)m + 1}{k-1}.$$

To construct H_3 , we enlarge $E(\pi)$ by the *first* of the above procedures. For each triple in H_3 we then select a pair of nonexcluded sites from any one of the first m rows, and a $(k-2)$ -element subset from the nonexcluded sites in the unused $m-1$ rows. To fix ρ , we consider whether the resulting R and C have already featured in H_1 ; if so, we choose the remaining possible ρ , and if not, we choose either of the two possibilities at random. The cardinality of H_3 is therefore

$$(9.8) \quad m \binom{m+1-2k^2}{2} \binom{m^2 - 2k^2m + 2k^2 - 1}{k-2}.$$

To construct H_4 , we enlarge $E(\pi)$ by the *second* of the above procedures. For each triple in H_4 we then select a pair of nonexcluded sites from any one of the nonextremal columns, and a $(k-2)$ -element subset from the nonexcluded sites in the unused $m-2$ columns. Fixing ρ as in H_3 , we obtain the cardinality of H_4 as

$$(9.9) \quad (m-1) \binom{m+1-2k^2}{2} \binom{m^2 - (2k^2+1)m + 2(2k^2-1)}{k-2}.$$

Combining (9.7), (9.8), and (9.9), we deduce that the cardinality of $\cup_{\ell=1}^4 H_\ell$ is polynomial in m , and of the required form. \square

We may combine our estimates for $t_k(\pi)$ in the following fashion.

Corollary 9.10. *For any m -permutation π , we have that*

$$t_k(\pi) = \frac{1}{k!}(m^{2k} + k(k-1)m^{2k-1}) + O(m^{2k-2})$$

when k is sufficiently small by comparison with m .

Proof. Combine the bounds $h_k(\pi) \leq t_k(\pi) \leq a_k(\pi)$ with Propositions 9.2 and 9.6. \square

We note in passing that Proposition 9.6 provides the lower bound

$$\frac{1}{2}(m^4 + 2m^3 - 195m^2 - 400m + 192) < h_2(\pi)$$

for $m > 7$. This should be compared with Proposition 5.7 and Theorem 8.16.

We conclude by explaining how the value of $t_k(\iota_m)$ may be computed from the Robinson-Schensted correspondence [12], which establishes a bijection between permutations π of dimension d and ordered pairs (P, Q) of standard Young tableaux having order d and shape the same. Under this correspondence, the length of the longest increasing subsequence of π is encoded as the number of columns in the tableaux. We therefore obtain a bijection between $T_k(\iota_m)$ and the set of pairs of tableaux having order $m + k$ and m or more columns.

Proposition 9.11. *The value of $t_k(\iota_m)$ is a polynomial of the form*

$$\frac{1}{k!}(m^{2k} + k(k-1)m^{2k-1})$$

modulo lower powers of m .

Proof. We first suppose that $k < m$, and let t'_k be the number of permutations in $T_k(\iota_m)$ for which ι_m is the largest identity submatrix; by Robinson-Schensted, these permutations correspond to pairs of standard Young tableaux of order $m + k$ with m elements in the first row. But according to the classic formula of Frame, Robinson, and Thrall [4] there are $(m+k)!/L_\lambda$ such tableaux, where λ is the partition of k provided by the remaining rows, and L_λ denotes the product of all hook lengths in the initial diagram. In consequence, we have that

$$(9.12) \quad t'_k = \sum_{\lambda} (m+k)!^2 / L_{\lambda}^2.$$

Applying the same considerations to the entire permutation group S_k , we deduce that

$$(9.13) \quad \sum_{\lambda} 1/H_{\lambda}^2 = 1/k!,$$

where H_{λ} denotes the product of the hook lengths in the appropriate diagram of order k . Thus

$$L_{\lambda} = (m+d_0)(m-1+d_1)\dots(m-r+d_r)(m-r-1)!H_{\lambda},$$

where the diagram for λ has columns of length $d_0 \geq d_1 \geq \dots \geq d_r$, with $\sum d_i = k$. For each λ we therefore have

$$(9.14) \quad \begin{aligned} (m+k)!/L_{\lambda} &= \left(\prod_{\substack{j=-r \\ j \neq d_i - i}}^k (m+j) \right) / H_{\lambda} \\ &= (m^k + k(k-1)m^{k-1}/2 + \dots) / H_{\lambda} \end{aligned}$$

where the remaining summands are smaller powers of m . Combining (9.12), (9.13), and (9.14), we deduce that t'_k is a polynomial of the form $(m^{2k} + k(k-1)m^{2k-1})/k!$ modulo lower powers of m . A parallel analysis of permutations with largest identity submatrix ι_{m+s} (for any $1 < s \leq k$) shows that they are enumerated by a polynomial in m of degree $2(k-s)$.

Whenever $k \geq m$ the extra terms introduced are also lesser powers of m , and the result follows. \square

A more detailed investigation along these lines is possible, and asymptotic formulae for various related quantities have been obtained by Regev [10]. As usual, Proposition 9.11 only provides a useful approximation when k is suitably small by comparison with m .

The importance of invoking the Robinson-Schensted correspondence is clear in codimension 2, for which our analysis in Theorem 8.16 confirms that $t_2(\pi)$ attains its minimum value when π is ι_m (amongst other possibilities). However, it is extremely interesting to observe from [15] (as amplified in [2]) that this is no longer the case for $t_k(\pi)$ when $k \geq 3$; for example $t_3(\iota_4) = 2,279$, yet $t_3(1324) = 2,278$. Caution is therefore required in placing Proposition 9.11 in an appropriate context, although it remains our only method for precise calculation of any $T_k(\pi)$ for $k > 2$.

REFERENCES

- [1] Sara Billey and Gregory Warrington. Kashdan-Lusztig polynomials for 321-hexagon-avoiding permutations. arXiv:math.CO/0005052, 5 May 2000.
- [2] Miklós Bóna. Permutations avoiding certain patterns; the case of length 4 and some generalizations. *Discrete Mathematics*, 175:55–67, 1997.
- [3] Miklós Bóna and Daniel Spielman. An infinite antichain of permutations. *Electronic Journal of Combinatorics*, 7(1), 2000.
- [4] J. S. Frame, G. deB. Robinson, and R. M. Thrall. The hook graphs of the symmetric group. *Canadian Journal of Mathematics*, 6:316–324, 1954.
- [5] Donald E. Knuth. *The Art of Computer Programming*, volume 1. Addison-Wesley, 1973.
- [6] V. Lakshmibai and B. Sandhya. Criterion for smoothness of Schubert varieties in $SL(N)/B$. *Proceedings of the Indian Academy of Science - Mathematical Sciences*, 100:45–52, 1990.
- [7] Richard Laver. Well-quasi-orderings and sets of finite sequences. *Mathematical Proceedings of the Cambridge Philosophical Society*, 79:1–10, 1976.
- [8] Ian G. Macdonald. *Notes on Schubert Polynomials*, volume 6 of *Publications du Laboratoire de Combinatoire et d'Informatique Mathématique*. Université du Québec à Montréal, 1991.
- [9] Toufik Mansour and Alek Vainshtein. Restricted permutations, continued fractions, and Chebyshev polynomials. arXiv:math.CO/9912052v2, 14 Mar 2000.
- [10] Amitai Regev. Asymptotic values for degrees associated with strips of Young diagrams. *Advances in Mathematics*, 41:115–136, 1981.
- [11] John Riordan. *Combinatorial Identities*. Krieger, 1979.
- [12] C. Schensted. Longest increasing and decreasing subsequences. *Canadian Journal of Mathematics*, 13:179–191, 1961.
- [13] Rodica Simion and Frank W. Schmidt. Restricted permutations. *European Journal of Combinatorics*, 6:383–406, 1985.
- [14] Richard P. Stanley. *Enumerative Combinatorics 1*. Wadsworth, 1986.
- [15] Julian West. *Permutations with forbidden subsequences; and, stack-sortable permutations*. PhD thesis, MIT, 1990.

UNIVERSITY OF MANCHESTER, MANCHESTER M13 9PL, ENGLAND
E-mail address: nige@ma.man.ac.uk

MALASPINA UNIVERSITY-COLLEGE, UNIVERSITY OF VICTORIA, CANADA
E-mail address: julian@math.uvic.ca