On packing densities of permutations

M. H. Albert, M. D. Atkinson, C. C. Handley
Department of Computer Science, University of Otago
D. A. Holton, Department of Mathematics and Statistics
University of Otago
W. Stromquist, Berwyn, Pennsylvania, USA

Abstract

The density of a permutation pattern π in a permutation σ is the proportion of subsequences of σ of length $|\pi|$ that are isomorphic to π. The maximal value of the density is found for several patterns π, and asymptotic upper and lower bounds for the maximal density are found in several other cases. The results are generalised to sets of patterns and the maximum density is found for all sets of length 3 patterns.

1991 *Mathematics Subject Classification* 05A15, 05A16

Keywords: pattern containment, permutations, layered permutations, packing density.

1 Introduction

The permutation 24153 has three subsequences 243, 253, 153 that are all *isomorphic* to (ordered in the same way as) the permutation 132. We can regard 132 as defining a pattern that occurs three times within 24153. There have been many investigations of this sort of pattern containment (for example [1, 5, 6, 8]) but most of them have focussed on characterising (or counting) the permutations which contain no instances of some particular pattern. In this paper we consider an opposite phenomenon: for some given permutation π, we shall want to know which permutations of a given length contain π the largest number of times (or, more precisely, which permutations of a given length contain the pattern defined by π the largest number of times). For example, the pattern (defined by) 2413 is contained in 35817246 a total of 17 times and no other permutation of length 8 has a greater "density" of 2413 occurrences.

Studies of this type were first suggested by Herbert Wilf in his 1992 address to the SIAM meeting on Discrete Mathematics. Although Wilf's lecture evidently inspired a great deal of work on pattern containment, very little on the question above has hitherto been published.

The most significant contribution that we know of is the doctoral thesis of Alkes Price [4] who proved a number of results about *layered* patterns (of which more below). Price's thesis is also a useful source of references to unpublished work by earlier researchers. Our work builds on these studies and solves some of the problems raised by these earlier workers. We address also the more general situation of maximising (within permutations of a given length) the number of occurrences from a particular *set* of patterns rather than a singleton set.

Let Π denote a set of patterns all of the same length k. For any permutation σ of length n let $\nu(\Pi, \sigma)$ be the number of subsequences of σ that are isomorphic to a permutation in Π (for brevity, we call these Π-subsequences of σ). Then it is natural to consider

$$d(\Pi, \sigma) = \frac{\nu(\Pi, \sigma)}{\binom{n}{k}}$$

which is the proportion of Π-subsequences in σ. In other words $d(\Pi, \sigma)$ is the probability that a subsequence of length k of σ is isomorphic to a pattern in Π. We are interested in how large this probability can be and we therefore define

$$\delta_n(\Pi) = \max\{d(\Pi, \sigma)\}$$

where the maximum is taken over all permutations of length n. This type of situation arises in other combinatorial settings (see, for example, [2] Theorem 4) and so the following result may be regarded as part of the combinatorial folklore.

Proposition 1.1 *If $k < n$ then $\delta_{n-1}(\Pi) \geq \delta_n(\Pi)$.*

PROOF: Let σ be a permutation of length n for which $d(\Pi, \sigma) = \delta_n(\Pi)$. Consider the following procedure for randomly choosing a subsequence of σ of length k. First choose (uniformly at random) a symbol s of σ to omit. Second choose (also uniformly at random) a subsequence of $\sigma - s$ of length k. Clearly this procedure produces subsequences of σ of length k uniformly at random and so

$$d(\Pi, \sigma) = \frac{1}{n} \sum_{s \in \sigma} d(\Pi, \sigma - s)$$

Therefore, for some s, $d(\Pi, \sigma - s) \geq d(\Pi, \sigma)$. But then

$$\delta_{n-1}(\Pi) \geq d(\Pi, \sigma - s) \geq d(\Pi, \sigma) = \delta_n(\Pi)$$

as required. ■

It follows immediately that, as $n \to \infty$, the sequence $(\delta_n(\Pi))$ approaches a limit $\delta(\Pi)$ which is called the *packing density* of Π. In cases where $\Pi = \{\pi\}$ is a singleton set it is convenient to denote this just by $\delta(\pi)$ and refer to the packing density of the permutation π. From the way in which $\delta(\Pi)$ has been defined it is evident that $\delta(\Pi) \leq \delta_n(\Pi)$ for all $n > k$. Much of this paper is concerned with evaluating $\delta(\Pi)$ for various sets Π. To do this it is helpful to know the values $\delta_n(\Pi)$ and therefore it is of interest to know the permutations σ for which $d(\Pi, \sigma) = \delta_n(\Pi)$.

If Π contains $12\ldots k$ then, clearly, $d(\Pi, 12\ldots n) = 1$ and so $\delta(\Pi) = 1$. Another obvious remark is that packing density is invariant under the usual 8 symmetries of permutation patterns [8]. As a first application of this consider $\delta(\pi)$ where π has length 3. By symmetry we can assume that $\pi = 123$ or $\pi = 132$. In the former case $\delta(\pi) = 1$. For the latter case Galvin, Kleitmann, and Stromquist (independently but unpublished) showed that $\delta(132) = 2\sqrt{3} - 3$. Price extended this to $\pi = 1k(k-1)\cdots 2$ and also handled a number of other cases where Π contains a single permutation including the case $\Pi = \{2143\}$. He concentrated exclusively on the case of layered permutations and relied heavily on an unpublished result of Stromquist [7] (which we prove and generalise in section 2).

Definition 1.2 A *layered* permutation is one which has a partition into segments $\sigma_1, \sigma_2, \ldots, \sigma_r$ where $\sigma_1 < \sigma_2 < \ldots < \sigma_r$ (that is, the elements of each σ_i are less than those of σ_{i+1}), and where the elements within each σ_i are consecutive and decreasing. The segments σ_i are called the *layers* of the permutation.

The main theorem of [7] is that, when Π consists of layered permutations, there is a layered permutation σ that maximises $\nu(\Pi, \sigma)$. We give the proof of this theorem in section 2 and give a condition on Π which ensures that *all* maxima of $\nu(\Pi, \sigma)$ occur at layered permutations. This result is very useful because it reduces the search for permutations of maximal density to a numerical optimisation problem and we show how to solve this in several cases.

Also in section 2 we shall compute $\delta(\Pi)$ for several sets Π of layered permutations. This touches on what appears to be an important dichotomy: for some Π the number of layers in an optimal permutation σ of length n is unbounded as $n \to \infty$, and for other Π this number is bounded. Price gave examples of singleton sets $\Pi = \{\pi\}$ of both types. Specifically, he proved that if π has two layers of sizes 1 and $k-1$ then $\Pi = \{\pi\}$ is of the unbounded type; on the other hand, if π has either two layers each of size greater than 1, or layers all of the same size $b > 1$, then $\Pi = \{\pi\}$ is of the bounded type. We shall generalise these results and show that the issue is quite sensitive to the presence or absence of layers of size 1.

The non-layered case appears to be much more difficult but, in section 3, we give some lower and upper bounds for the packing density of $\{1342\}$ and $\{2413\}$. We also compute $\delta(\Pi)$ for any set of permutations of length 3.

In section 4 we list some open problems.

2 Layered patterns

Although this section is about the special case of layered patterns we need to begin by generalising the situation in two ways.

- We drop the requirement that the members of Π have the same length,

- We allow Π to be a multiset where each $\pi_i \in \Pi$ occurs w_i times ($w_i > 0$). Then each occurrence of π_i in a permutation σ contributes w_i to the total number $\nu(\Pi, \sigma)$.

We shall be interested in the maximal value of the Π-enumerator

$$\nu(\Pi, \sigma) = \sum_{i=1}^{m} w_i \nu(\pi_i, \sigma)$$

as σ varies over permutations of length n. Those permutations for which the Π-enumerator is maximal will be called Π-*maximal.*

As an inductive aid we define a permutation π to be *layered on top* if it ends with a non-empty segment of the form $p, p-1, p-2\ldots$ where p is the maximum symbol of π. Even though π itself may not be layered it is convenient to call this segment the *final layer* of π.

Proposition 2.1 *Let Π be a multiset of permutations all of which are layered on top. Then among the Π-maximal permutations of each length there will be one that is layered on top. Furthermore if the final layer of every $\pi \in \Pi$ has size greater than 1 then every Π-maximal permutation is layered on top.*

Proof: In this proof we shall use the representation of a permutation as a *planar* poset, i.e. as a set of points in the plane ordered by dominance ($(x_1, y_1) \leq (x_2, y_2)$ if $x_1 \leq x_2$ and $y_1 \leq y_2$). Given a permutation $\sigma = s_1, \ldots, s_n$ the corresponding planar poset is the set $\{(i, s_i)\}$. Conversely, given a planar poset as a set of points in the plane with distinct x-coordinates and distinct y-coordinates we can replace them by an equivalent set using the unique order preserving bijections

$$\begin{array}{rcl} \{x_1, x_2, \ldots, x_n\} & \longrightarrow & \{1, 2, \ldots, n\} \\ \{y_1, y_2, \ldots, y_n\} & \longrightarrow & \{1, 2, \ldots, n\} \end{array}$$

Now if we list the pairs by increasing order of first coordinate, the second coordinates define a permutation of $\{1, 2, \ldots, n\}$.

The planar poset of a permutation that is layered on top has the form shown in Figure 1.

Consider an arbitrary planar poset. We say that two maximal points (x_1, y_1) and (x_2, y_2) of the poset are *adjacent* if there are no points in the poset whose first components lie between x_1 and x_2 and no points whose second components lie between y_1 and y_2. Of course, if the first (respectively second) components comprise the set $\{1, 2, \ldots, n\}$ this says that x_1 and x_2 (respectively y_1 and y_2) differ by 1; but, in practice, this assumption is not very useful.

Furthermore, we say that a pair of maximal points is *connected* if the pair is in the transitive closure of the adjacency relation. It is easy to see that a poset is layered on top if and only if every pair of maximal points is connected.

Let σ be a Π-maximiser that is not layered on top and let P be its associated planar poset. Consider any two maximal points u and v of P that are not

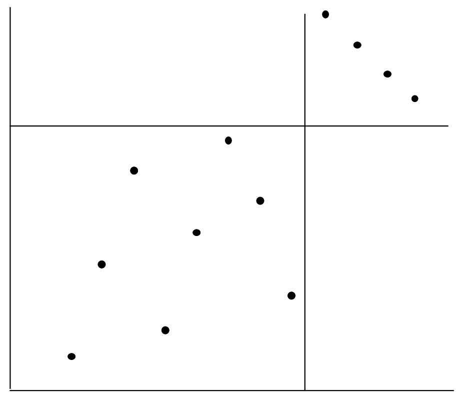

Figure 1: A permutation layered on top

connected. We categorise the Π-subposets of P according to which of u, v they contain by defining

$\nu(\Pi, P, u, v)$, the number of Π-subposets of P containing both u and v
$\nu(\Pi, P, \bar{u}, v)$, the number of Π-subposets of P containing v but not u
$\nu(\Pi, P, u, \bar{v})$, the number of Π-subposets of P containing u but not v
$\nu(\Pi, P, \bar{u}, \bar{v})$, the number of Π-subposets of P containing neither u nor v

Obviously

$$\nu(\Pi, P) = \nu(\Pi, P, u, v) + \nu(\Pi, P, \bar{u}, v) + \nu(\Pi, P, u, \bar{v}) + \nu(\Pi, P, \bar{u}, \bar{v})$$

is the total number of Π-subposets of P.

We now consider two posets P_1 and P_2 obtained by slight 'editing' of the poset P. In P_1 the point v has been moved so that it is adjacent to u, while in P_2 the point u has been moved so that it is adjacent to v. An example of planar posets P and P_1 which are related in this way is shown in Figure 2. In the example, maximal points are shown as squares and other points are shown as circles.

Clearly

$$\nu(\Pi, P, \bar{u}, \bar{v}) = \nu(\Pi, P_1, \bar{u}, \bar{v}) = \nu(\Pi, P_2, \bar{u}, \bar{v})$$

since the subposets of P, P_1, P_2 consisting of points other than u, v are equal.

Similarly

$$\nu(\Pi, P, u, \bar{v}) = \nu(\Pi, P_1, u, \bar{v})$$

since the subposets of P, P_1 consisting of points other than v are equal, and

$$\nu(\Pi, P, \bar{u}, v) = \nu(\Pi, P_2, \bar{u}, v)$$

for a similar reason.

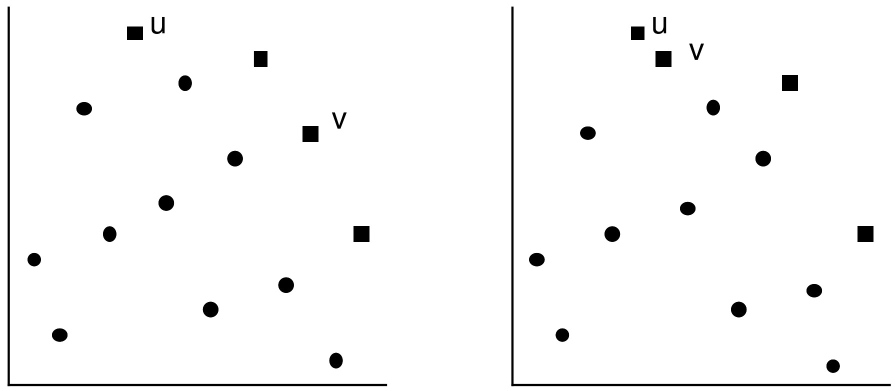

Figure 2: Planar posets P and P_1

Also

$$\nu(\Pi, P_1, \bar{u}, v) = \nu(\Pi, P_1, u, \bar{v})$$

and

$$\nu(\Pi, P_2, \bar{u}, v) = \nu(\Pi, P_2, u, \bar{v})$$

since each of P_1, P_2 have an automorphism that exchanges u, v while fixing the other points.

We also have

$$\nu(\Pi, P, u, v) \leq \nu(\Pi, P_1, u, v)$$

and

$$\nu(\Pi, P, u, v) \leq \nu(\Pi, P_2, u, v)$$

These inequalities follow from the layered on top condition by the following argument. Consider a Π-subposet Q of P with $u, v \in Q$. Since u and v are maximal in P they must correspond to maximal points of Q. Since Q is layered on top its maximal points are all greater than every non-maximal point. But, in P_1, the new position of u ensures that it has the same comparabilities as v does. Therefore there is a Π-subposet Q_1 of P_1 on the same set of points as Q. A similar argument holds for P_2.

Two equalities now follow from the maximality of $\nu(\Pi, P)$. First, we must have $\nu(\Pi, P, \bar{u}, v) = \nu(\Pi, P, u, \bar{v})$. For, suppose the left hand side was greater than the right hand side, then

$$\begin{aligned}
\nu(\Pi, P) &= \nu(\Pi, P, u, v) + \nu(\Pi, P, \bar{u}, v) + \nu(\Pi, P, u, \bar{v}) + \nu(\Pi, P, \bar{u}, \bar{v}) \\
&< \nu(\Pi, P, u, v) + \nu(\Pi, P, \bar{u}, v) + \nu(\Pi, P, \bar{u}, v) + \nu(\Pi, P, \bar{u}, \bar{v}) \\
&\leq \nu(\Pi, P_2, u, v) + \nu(\Pi, P_2, \bar{u}, v) + \nu(\Pi, P_2, \bar{u}, v) + \nu(\Pi, P_2, \bar{u}, \bar{v}) \\
&\leq \nu(\Pi, P_2, u, v) + \nu(\Pi, P_2, \bar{u}, v) + \nu(\Pi, P_2, u, \bar{v}) + \nu(\Pi, P_2, \bar{u}, \bar{v}) \\
&= \nu(\pi, P_2)
\end{aligned}$$

This contradicts the maximality of $\nu(\Pi, P)$. It also shows that $\nu(\Pi, P, u, v) = \nu(\Pi, P_1, u, v) = \nu(\Pi, P_2, u, v)$. In particular we have

$$\nu(\Pi, P) = \nu(\Pi, P_1) = \nu(\Pi, P_2)$$

The procedure just described shows that we can transfer maximal points from one component to another and obtain another Π-maximal poset. We can now prove the first part of the proposition. Starting from any Π-maximal poset we transfer maximal points between components until all the maximal points are in one component. The resulting poset is still Π-maximal and is now layered on top.

We turn to the proof of the second part of the proposition where we have the hypothesis that every $\pi \in \Pi$ has top layer of size greater than 1.

Suppose, for a contradiction, that there exist Π-maximisers which are not layered on top and choose one, σ say, with a largest number of maximal points.

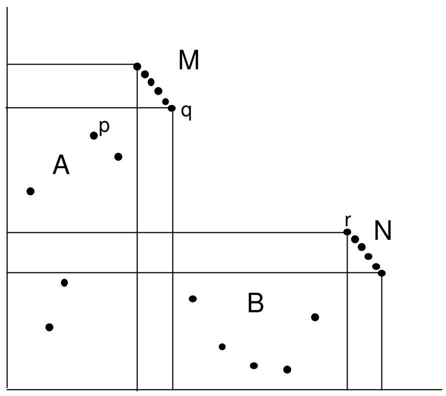

Figure 3: Two components of maximal points

We carry out a sequence of the operations described above until we have ensured that the number of connected components of maximal elements is 2 and we have a poset as shown in Figure 3 where these components are denoted by M and N. Notice that these operations do not reduce the number of maximal elements in a poset.

Consider the two maximal points marked as q and r in Figure 3. Since the set of maximal points does not form a single component there exists at least one point in one of the two regions marked A and B. We shall derive a contradiction from the assumption that region A is non-empty (a similar argument gives a contradiction if B is non-empty).

Among the points of region A choose the point p with largest 'value' (second component). There must be at least one Π-subsequence which uses p (otherwise we could move p elsewhere within the planar poset so that it produced another Π-subsequence).

We consider two cases. In the first case p is matched to some point x of some $\pi \in \Pi$ such that x is not in the final layer T of π. In this case the match of π must match T to some of the points of M. In particular this implies that $|M| \geq |T|$. But now the 'editing' operation that moves r to be adjacent to q (which we know does not decrease the number of Π-subsequences containing at most one of q and r) must introduce some further Π-subsequences containing both q and r. This is a contradiction.

In the second case we can assume that, for every $\pi \in \Pi$, all π-subsequences containing p match p to a point of the final layer of π (and therefore to the largest point of π). However we may now 'edit' the poset in a different way. We move p vertically upwards so that it has highest value in the poset. Since p matched only points of largest value, this cannot decrease the number of Π-subsequences. As it increases the number of maximal points we shall have a contradiction unless the maximal points now form a single component (recall that σ was chosen to have as many maximal points as possible). For the maximal points now to be connected we must have had $A = \{p\}$ and $B = \emptyset$. Now consider some $\pi \in \Pi$ that takes part in one of the matches involving p. Since the top layer of π has more than one point the remainder of the top layer of π is matched to some of the elements of N. But with p in its new position there will be additional matches where the rest of the top layer will be matched to points of $M \cup N$, a final contradiction.

$\blacksquare$

Theorem 2.2 *Let Π be a multiset of layered permutations. Then among the Π-maximal permutations of each length there will be one that is layered. Furthermore if all the layers of every $\pi \in \Pi$ have size greater than 1 then every Π-maximal permutation is layered.*

PROOF: For each permutation $\pi_i \in \Pi$ we let $\pi_i = \phi_i \lambda_i$ where λ_i is the final layer of π_i. By Proposition 2.1 for every $n > 0$ we can find a Π-maximal permutation that is layered on top. Thus, in searching for every Π-maximal permutations of length n we can confine our attention to permutations of the form $\sigma = \theta\lambda$ where λ is the final layer of some fixed length $n - k$ and θ is a permutation of length $k < n$.

The π_i-subsequences of σ have two possible forms

- a π_i-subsequence of θ, and

- a ϕ_i-subsequence of θ together with a λ_i-subsequence of λ (of which there are $\binom{|\lambda|}{|\lambda_i|}$).

This proves that

$$\nu(\pi_i, \sigma) = \nu(\pi_i, \theta) + \binom{|\lambda|}{|\lambda_i|}\nu(\phi_i, \theta)$$

and so

$$\begin{aligned}
\nu(\Pi, \sigma) &= \sum w_i \nu(\pi_i, \sigma) \\
&= \sum w_i \nu(\pi_i, \theta) + \sum w_i \binom{|\lambda|}{|\lambda_i|} \nu(\phi_i, \theta) \\
&= \nu(\Pi', \theta)
\end{aligned}$$

where Π' is the multiset in which each π_i occurs w_i times and each ϕ_i occurs $w_i\binom{|\lambda|}{|\lambda_i|}$ times. Since σ is a Π-maximal permutation, θ is necessarily a Π'-maximal permutation. By induction on n we can find a layered Π'-maximal permutation θ and then σ itself will be layered.

This proves the first part. The second part is proved by a similar argument. In this case every Π-maximal permutation is layered on top and has the form $\theta\lambda$ as above. Then, because Π' also consists of permutations whose layer sizes are greater than 1, we can use induction again to deduce that θ is necessarily layered. ■

Remark 2.3 The hypothesis about layers of size 1 in the second statement of the theorem cannot be omitted. Let Π consist of 43215 alone. Then

$$16, 15, \ldots, 2, 1, 17, 19, 20, 18$$

is a Π-maximiser of length 20 that is not layered on top.

As a first application of this theorem we have

Proposition 2.4 *Let Π consist of the single permutation 1243. Then every Π-maximiser has the form*

$$1\,2\ldots k\;n\;n-1\ldots k+1$$

where $k = \lfloor n/2 \rfloor$ or $\lceil n/2 \rceil$. In addition, $\delta(\Pi) = 3/8$.

PROOF: In obtaining the form of a Π-maximiser σ we may confine our attention to layered permutations which we describe simply by giving their layer sizes. First of all, note that σ must begin with the symbol 1; for we can always move that symbol to first place without destroying any 1243-subsequences, and unless all layers from the second onwards are of size one (impossible as there would be no 1243-subsequences at all) this results in new 1243-subsequences.

Since our aim is to prove that σ begins with k layers of size 1 followed by a single layer of size $n-k$ we suppose for a contradiction that the layer sizes are

$$1^a b c_1 c_2 \cdots c_u d$$

with $a \geq 1$, $b \geq 2$, and $u \geq 0$. We will show that such a permutation does not have a maximal number of 1243-subsequences.

First consider the possibility of replacing the layer of size b by a further increasing segment, giving layer sizes

$$1^{a+b}\, c_1\, c_2 \cdots c_u\, d$$

In making this transformation the only 1243-subsequences we destroy are those whose first two elements come from among the initial a symbols, and whose final two come from the next b symbols, that is we lose:

$$\binom{a}{2}\binom{b}{2}$$

1243-subsequences. On the other hand we gain

$$\binom{b}{2}\left(\binom{c_1}{2}+\binom{c_2}{2}+\cdots+\binom{c_u}{2}+\binom{d}{2}\right)$$

1243-subsequences. Since σ is a Π-maximiser the gain cannot exceed the loss and so it is clear that $a \geq d$.

Now suppose that $u \geq 1$ and consider the possibility of adding the layer of size c_u to the final layer, giving layer sizes

$$1^a\, b\, c_1\, c_2 \cdots c_{u-1}\, (c_u + d).$$

The 1243-subsequences lost are those involving one element of the layer of size c_u (matching the '2' symbol) and a pair of the final layer (matching the symbols '43'). There are

$$(a + b + c_1 + \cdots + c_{u-1})\, c_u \binom{d}{2}$$

such. The 1243-subsequences gained are any involving a pair of elements (in increasing order) from the layers up to the one of size c_{u-1} and a decreasing pair from the last layer, one from the first c_u elements, the other from the final d elements. There are

$$\left(\binom{a}{2} + a(b + c_1 + \cdots + c_{u-1})\right) c_u d$$

of these. If $a \geq d$, an easy algebraic manipulation establishes that the gain is bigger than the loss.

This leaves the case of layer sizes $1^a bd$ with $a \geq d$. We can repeat the transformation just used, with b in place of c_u. If $a > d$ this leads to a net increase in the number of 1243-subsequences so we must have $a = d$. When $a = d$ the transformation neither gains nor loses 1243-subsequences. But then the permutation with layer sizes $1^a(a+b)$ would also maximise the 1243-subsequences; but that permutation has fewer 1243-subsequences than the permutation with layer sizes $1^{a+1}(a+b-1)$.

This contradiction has proved that the layers of σ have sizes $1^a d$. In this case the number of 1243-subsequences is clearly

$$\binom{a}{2}\binom{d}{2}$$

But the maximum value of this expression occurs for the values of a, d claimed in the statement of the proposition. The value of $\delta(1243)$ is now easily seen to be $3/8$. ■

Our second application of Theorem 2.2 is to some sets Π containing two permutations.

Theorem 2.5 *Let $\Pi = \{\alpha, \beta\}$ where α and β each have two layers with sizes k, l and l, k respectively. Then every Π-maximiser of length at least $k + l$ has exactly two layers.*

PROOF: Let σ be a Π-maximiser of length $n \geq k + l$. Certainly, by Theorem 2.2, σ will be layered. Suppose, for a contradiction, that σ has 3 layers or more. Interchanging the sizes of two adjacent layers does not affect the number of Π-subsequences since it simply trades subsequences isomorphic to α for subsequences isomorphic to β and vice versa. Therefore we may assume that the layers of σ come in increasing order of size $a \leq b \leq \ldots c$.

We shall obtain our contradiction by showing that the permutation σ' obtained from σ by replacing the first two layers by a layer of size $a + b$ has more Π-subsequences. In comparing the number of Π-subsequences of σ' with the number in σ we note that we shall have lost Π-subsequences whose elements come from the first two layers of σ; there are

$$L = \binom{a}{k}\binom{b}{l} + \binom{a}{l}\binom{b}{k}$$

of these.

On the other hand σ' has gained some additional Π-subsequences whose elements come from the new first layer, but not entirely from among the first a or final b elements, and from some other layer. There are at least

$$G = \left(\binom{a+b}{k} - \binom{a}{k} - \binom{b}{k}\right)\binom{c}{l} + \left(\binom{a+b}{l} - \binom{a}{l} - \binom{b}{l}\right)\binom{c}{k}$$

of these. We have to prove that the latter expression is larger than the former, or that $G - L$ is positive.

We do this by considering two cases. In the first case we assume one of k and l equals 1 (k say). Then

$$G - L = \binom{a+b}{l} c - (a+c)\binom{b}{l} - (b+c)\binom{a}{l}$$

But note that

$$\begin{aligned}
\binom{a+b}{l} &= \frac{(a+b)(a+b-1)\dots(a+b-l+1)}{l!} \\
&> \frac{(a+b)(a-1)\dots(a-l+1)}{l!} \\
&\quad + \frac{(a+b)(b-1)\dots(b-l+1)}{l!} \\
&= \frac{a+b}{a}\binom{a}{l} + \frac{a+b}{b}\binom{b}{l}
\end{aligned}$$

Therefore $G - L$ is larger than

$$\left(\frac{c(a+b)}{a} - (b+c)\right)\binom{a}{l} + \left(\frac{c(a+b)}{b} - (a+c)\right)\binom{b}{l}$$

which is non-negative.

In the second case we may assume that both k and l are larger than 1. We shall show that the difference between the first term of G and the first term of L is positive (and then by symmetry the same will be true of the second terms). This difference is (since $b \leq c$) at least

$$\left(\binom{a+b}{k} - 2\binom{a}{k} - \binom{b}{k}\right)\binom{c}{l}$$

Considered as a function of b this expression is minimal (assuming that $b \geq a$) when $b = a$. So it is enough to prove that

$$\binom{2a}{k} - 3\binom{a}{k} > 0$$

However, this follows easily as

$$\binom{2a}{k} \geq 2^k \binom{a}{k}$$

■

Remark 2.6 It is not in general true that the two layers in the Π-maximiser σ in the above result differ in size by at most 1. However, for future use we note that when $k = 1, l = 2$, the layer sizes of σ do indeed have this property; this is easily seen by considering the maximal value of

$$\nu(\{132, 213\}, \sigma) = a\binom{b}{2} + b\binom{a}{2}$$

as a and b vary subject to $a+b=n$. It follows easily from this that

$$\delta(\{132,213\}) = \frac{3}{4}$$

In the remainder of this section we shall suppose that Π consists of a single layered permutation π of length m with r layers of sizes $m_1,\ldots,m_r$. We know that every π-maximiser σ is layered and in Price's investigations [4] two different situations emerged. The number of layers of σ might grow without bound as $n\to\infty$ (exemplified by the case $r=2, m_1=1, m_2=k-1$ and analysed thoroughly by Price). On the other hand it might be possible to choose the permutations σ to have a bounded number of layers. In the latter case the bound may be larger than the number of layers of π itself (Price gave the example of $r=6$ and all $m_i=2$ where the best upper bound is 7).

We shall give some results which distinguish the bounded case from the unbounded case. To do this it is necessary to adopt some of the machinery of [4]. We begin by defining

$$p_s(\lambda_1,\ldots,\lambda_s) = \binom{m}{m_1\cdots m_r} \sum_{1\le e_1<\ldots<e_r\le s} \lambda_{e_1}^{m_1}\ldots\lambda_{e_r}^{m_r}$$

This expression has a probabilistic interpretation. Suppose the unit interval is partitioned into s subintervals of length $\lambda_1,\ldots,\lambda_s$ (reading from left to right) and we choose m points from the interval uniformly at random. Then $p_s(\lambda_1,\ldots,\lambda_s)$ is the probability that the leftmost m_1 of these points will all lie in one of the subintervals, the next m_2 will all lie in another of the subintervals, and so on.

This probabilistic interpretation leads one to expect a connection with packing density. For suppose that n is large, and σ is a layered permutation with layer sizes close to $\lambda_1 n, \lambda_2 n,\ldots,\lambda_s n$. Then m distinct points randomly chosen from σ will form a subsequence isomorphic to π exactly if the leftmost m_1 of them lie in a single layer, the next m_2 lie in another layer, and so on. Thus the probability $d(\pi,\sigma)$ of such an event will be close to $p_s(\lambda_1,\ldots,\lambda_s)$[1].

Let

$$p_s = \max p_s(\lambda_1,\ldots,\lambda_s)$$

where the maximum is taken over the set

$$\{(\lambda_1,\ldots,\lambda_s) \mid \lambda_i\ge 0, \sum\lambda_i = 1\}$$

Notice that this set is compact and so the maximum is attained by a point $(\bar\lambda_1,\ldots,\bar\lambda_s)$ in the set. Price did not discuss the uniqueness or otherwise of the maximising $(\bar\lambda_1,\ldots,\bar\lambda_s)$ and this issue remains open. However, we do know that when this theory is generalised to sets of permutations Π the maximising $(\bar\lambda_1,\ldots,\bar\lambda_s)$ is not unique.

[1] These remarks are proven rigorously in [4] but their general tenor is a standard part of the probabilistic method.

Price showed that such a maximal point has an interpretation for packing density. Suppose we consider permutations *with at most s layers* and ask which ones have the greatest density $\tau_s(n)$ of π-subsequences. Then, in fact, $\tau_s(n) \to p_s$ as $n \to \infty$. Furthermore, the permutations which achieve $\tau_s(n)$ may be chosen to have layer sizes which, as $n \to \infty$, approach the proportions $\bar{\lambda}_1, \dots, \bar{\lambda}_s$. Notice that if any of $(\bar{\lambda}_1, \dots \bar{\lambda}_s)$ is zero then $p_{s-1} = p_s$; so, in this case, permutations with $s-1$ layers are, asymptotically, as dense in π-subsequences as those with s layers.

Clearly $p_r \leq p_{r+1} \leq \dots$ and, in fact, $p_s \to \delta(\pi)$. It is not difficult to verify that the bounded case corresponds precisely to when $p_s = \delta(\pi)$ for some index s (and hence for all larger indices).

Theorem 2.7 *A permutation π none of whose layers has size 1 is of the bounded type. In particular, $p_s = \delta(\pi)$ for all $s \geq \frac{1}{2}(m-1)(p_r+1)/p_r$.*

Proof: We use the probabilistic machinery and notation summarised above. Consider any value of s with $s > r$ and maximising values $\bar{\lambda}_1, \dots, \bar{\lambda}_s$ for $p_s = \max p_s(\lambda_1, \dots, \lambda_s)$. Then p_s is the probability of success for the experiment described above. Let $\bar{\lambda}_j$ be any of $\bar{\lambda}_1, \dots, \bar{\lambda}_s$. A successful outcome can arise in two ways:

- None of the m random points lie in the subinterval of length $\bar{\lambda}_j$. The probability that a success of this type happens is
$$(1-\bar{\lambda}_j)^m p_{s-1}(\lambda'_1, \dots, \lambda'_{s-1})$$
where $\lambda'_1, \dots, \lambda'_{s-1}$ arise by omitting $\bar{\lambda}_j$ from $\bar{\lambda}_1, \dots, \bar{\lambda}_s$ and scaling by $1/(1-\bar{\lambda}_j)$. The probability that this occurs is therefore at most $(1-\bar{\lambda}_j)^m p_{s-1}$.
- At least one, and therefore at least two, of the m points fall in the interval of length $\bar{\lambda}_j$ and the probability for this is at most $\binom{m}{2}\bar{\lambda}_j^2$.

Therefore we have

$$\begin{aligned} p_s &\leq (1-\bar{\lambda}_j)^m p_{s-1} + \binom{m}{2}\bar{\lambda}_j^2 \\ &\leq \left(1 - m\bar{\lambda}_j + \binom{m}{2}\bar{\lambda}_j^2\right) p_{s-1} + \binom{m}{2}\bar{\lambda}_j^2 \\ &= p_{s-1} + \bar{\lambda}_j \left(-mp_{s-1} + \binom{m}{2} p_{s-1}\bar{\lambda}_j + \binom{m}{2}\bar{\lambda}_j\right) \end{aligned}$$

and so

$$0 \leq p_s - p_{s-1} \leq \bar{\lambda}_j \left(-mp_{s-1} + \binom{m}{2} p_{s-1}\bar{\lambda}_j + \binom{m}{2}\bar{\lambda}_j\right)$$

Therefore $\bar{\lambda}_j = 0$ or $-p_{s-1} + \frac{1}{2}(m-1)p_{s-1}\bar{\lambda}_j + \frac{1}{2}(m-1)\bar{\lambda}_j \geq 0$. But the second alternative gives

$$\begin{aligned} \bar{\lambda}_j &\geq \frac{2p_{s-1}}{(m-1)(p_{s-1}+1)} \\ &\geq \frac{2p_r}{(m-1)(p_r+1)} \end{aligned}$$

and there are at most

$$\frac{(m-1)(p_r+1)}{2p_r}$$

$\bar{\lambda}_j$'s for which this can be true. ■

By a similar argument applied to adjacent pairs of layers we can prove

Proposition 2.8 *Any permutation which has three layers of sizes $a, 1, b$ with $a, b > 1$ is of the bounded type.*

The last two results together with a great deal of computational evidence have led us to make the

Conjecture 2.9 *Suppose that π is a layered permutation whose first and last layers have sizes greater than 1 and which has no adjacent layers of size 1. Then π is of the bounded type.*

In the other direction we have

Proposition 2.10 *If π is a permutation whose first or last layer is of size 1 then π is of the unbounded type.*

PROOF: Again we adopt all the notation given before the proof of Theorem 2.7. Suppose that the first layer of π has size 1 (the other case is proved in the same way). Then, as $m_1 = 1$, the coefficient of λ_1 in any $p_s(\lambda_1, \dots, \lambda_s)$ occurs to the first power only and it follows readily that every summand of $p_{s-1}(\lambda_1 + \lambda_2, \lambda_3, \dots, \lambda_s)$ is also a summand of $p_s(\lambda_1, \lambda_2, \dots, \lambda_s)$. Thus, for positive $\lambda_1, \lambda_2, \dots, \lambda_s$ we have

$$p_{s-1}(\lambda_1 + \lambda_2, \lambda_3, \dots, \lambda_s) < p_s(\lambda_1, \lambda_2, \dots, \lambda_s)$$

with strict inequality since the right hand side has terms of the form $\lambda_1 \lambda_2^{m_2} \dots$ which are not present on the left hand side.

We need to prove that $p_{s-1} < p_s$ for all $s \geq r$. This is certainly true for $s = r$ as $p_r > p_{r-1} = 0$. We make the inductive assumption that $p_{s-2} < p_{s-1}$. Then a maximising point $(\bar{\lambda}_2, \dots, \bar{\lambda}_s)$ for $p_{s-1}(\lambda_2, \dots, \lambda_s)$ must have all non-zero components (as a zero component could be omitted to obtain a maximising point for $p_{s-2}(\lambda_2,\ldots))$. But then

$$\begin{aligned} p_{s-1} &= p_{s-1}(\bar{\lambda}_2,\ldots,\bar{\lambda}_s) \\ &= p_{s-1}(\frac{1}{2}\bar{\lambda}_2+\frac{1}{2}\bar{\lambda}_2,\ldots,\bar{\lambda}_s) \\ &< p_s(\frac{1}{2}\bar{\lambda}_2,\frac{1}{2}\bar{\lambda}_2,\ldots,\bar{\lambda}_s) \\ &\leq p_s \end{aligned}$$

as required. ∎

3 Non-layered patterns

All patterns of length 3 are equivalent to either 123 or 132 and, as stated in the introduction, answers to the packing density questions for these permutations are already known. Of the patterns of length 4 the only permutations not handled by the techniques of the previous section are equivalent to either 1342 and 2413. We have the following two results.

Proposition 3.1 $0.19657 \leq \delta(1342) \leq 2/9$.

PROOF: To prove the upper bound let f_n denote the maximal value of $\nu(1342,\sigma)$ as σ varies over permutations of length n. Consider a permutation σ attaining this bound and write $\sigma = \alpha n \beta$. Occurrences of 1342 within σ are of two kinds: those contained entirely within $\alpha\beta$, and those which contain the symbol n.

In the first kind, as $\alpha\beta$ is a permutation of length $n-1$, we shall have at most f_{n-1} 1342-subsequences.

For the second type we need to give an upper bound for the number of subsequences of length 3 which have two symbols $a_1, a_2 \in \alpha$ and a symbol $b \in \beta$ with $a_1 < b < a_2$. We consider how to choose α, β so as to maximise the number of such suitable (a_1, a_2, b) triples. Clearly, if α contains a pair of symbols that are not in increasing order we cannot decrease the number of suitable triples by interchanging them; therefore we may presume that α is increasing. Similarly we may assume that β is increasing.

We now argue that the symbols of β may be taken to be consecutive. If that were not the case we could write $\beta = \lambda\mu\rho\tau$ where μ and ρ are each sequences of consecutive symbols. In addition we would have $\alpha = \gamma\delta\epsilon$ where $\gamma < \mu < \delta < \rho < \epsilon$. Let $\gamma, \delta, \epsilon, \mu, \rho$ have lengths g, d, e, m, r respectively. The number of triples a_1, a_2, b which have $b \in \mu\rho$ is easily seen to be

$$g(d+e)m + (g+m)er = ge(m+r) + d(gm+er)$$

For fixed $m+r$ this is greatest when either $m = 0$ or $r = 0$ (according to whether $g \leq e$. In other words we may take one of μ and ρ to be empty.

Now that we know that β consists of consecutive values we can put $\alpha\beta = \alpha_1\alpha_2\beta$ with $\alpha_1 < \beta < \alpha_2$ and compute the number of suitable triples as

$|\alpha_1||\alpha_2||\beta|$. This expression is maximal when all three factors are equal. So the number of suitable triples is at most $n^3/27$.

Therefore $f_n \leq f_{n-1} + n^3/27$ and so

$$\delta_n(1342) \leq \frac{\sum_{i=1}^{n} i^3/27}{\binom{n}{4}}$$

and the right hand side tends to $2/9$ as $n \to \infty$.

To get the lower bound we consider a sequence of permutations σ_n (with $n = |\sigma|$) structured as $\alpha\beta$ with the following conditions:

- α has length $\lfloor an \rfloor$ where a is a constant fraction to be chosen below,
- $\alpha < \beta$,
- β maximises the number of 231 patterns,
- α is structured in the same way recursively; that is, $\alpha = \sigma_{\lfloor an \rfloor}$.

We choose a to maximise the limiting value $\lim_{n\to\infty} d(1342, \sigma_n)$.

Suppose that p_n is the probability of picking 4 points from σ_n that define an occurrence of 1342 (that is, $p_n = d(1342, \sigma_n)$). Such an event occurs if either the 4 points chosen lie in α and define a 1342 subsequence, or one of them lies in α and the three others define a 231-subsequence of β. Therefore

$$p_n \geq a^4 p_{\lfloor an \rfloor} + 3a(1-a)^3 q_{n-\lfloor an \rfloor}$$

where $q_k = d(231, \beta)$. Let $p = \lim_{n\to\infty} p_n$ and $q = \lim_{n\to\infty} q_n$. Then $q = \delta(\{231\}) = 3 - 2\sqrt(3)$ ([4]) and, taking limits in the above inequality and rearranging, we have

$$p \geq \frac{3a(1-a)^3}{1-a^4}(2\sqrt(3) - 3)$$

The maximal value of the right hand side is $0.19657\ldots$ and since p is a lower bound on $\delta(1342)$ this completes the proof.

■

Proposition 3.2 $51/511 \leq \delta(2413) \leq 2/9$.

PROOF: The upper bound is proved in the same way as in the previous proposition. For the lower bound we consider permutations of the form $\sigma = \sigma_3\sigma_5\sigma_8\sigma_2\sigma_7\sigma_1\sigma_4\sigma_6$ where $\sigma_i < \sigma_{i+1}$ and $|\sigma_i| = |\sigma_{i+1}|$ for each i, and where each σ_i is recursively structured in the same way. The reason behind this construction is that the permutation 35827146 contains 2413 a relatively large number of times (17). Consider the probability p_n (where $n = |\sigma|$) of selecting 4 points from σ which constitute a 2413 pattern. We can do this by picking the 4 points from any one of the segments σ_i and, having done that we shall have a

2413 pattern with probability $p_{n/8}$; or we can pick any of the 17×24 ordered quadruples from distinct σ_i which give a 2413 pattern. Thus

$$p_n = (\frac{1}{8})^4 8p_{n/8} + \frac{17 \times 24}{8^4}$$

As $n \to \infty$, p_n approaches a limit p which satisfies the equation

$$p = \frac{p}{512} + \frac{51}{512}$$

Since $p_n \leq \delta_n(2413)$ we have $51/511 = p \leq \delta(2413)$ as required. ∎

For both 1342 and 2413 we believe that our lower bounds are closer to the packing density than the upper bounds. In the case of $\delta(1342)$ our belief stems from computing $\delta_n(1342)$ for $n \leq 16$ and noting the form of the optimising permutations which are close in form to the permutations we have used to derive the lower bounds. For $\delta(2413)$ we have examined a large number of patterns on which to base a recursive lower bound construction but have found none better than 35827146.

Finally we discuss $\delta(\Pi)$ for sets of permutations of length 3. By applying the usual permutation symmetries the only cases we need to consider are

$$\Pi = \{132\}, \{132, 213\}, \{132, 231\}, \{132, 213, 231\}, \{132, 213, 231, 312\}$$

The first of these is handled in [4]. In the second case we can appeal to the remark following Theorem 2.5.

Proposition 3.3 *If* $\Pi = \{132, 231\}$ *then* $\delta(\Pi) = 1/2$.

PROOF: Let $\sigma = \alpha n \beta$ be any permutation of length n. The Π-subsequences in σ occur either within $\alpha\beta$ or have their middle symbol as n. The latter type arise for every anb with $a \in \alpha, b \in \beta$ and so there are $|\alpha||\beta|$ of them. Hence if σ maximises the Π-subsequences we must have $|\alpha|$ and $|\beta|$ as equal as possible and also have $\alpha\beta$ maximising the number of Π-subsequences among permutations of length $n-1$. This proves that the maximal number f_n of Π-subsequences satisfies

$$f_n = f_{n-1} + \frac{n^2}{4} + o(n^2)$$

from which the result easily follows. ∎

Proposition 3.4 *If* $\Pi = \{132, 213, 231, 312\}$ *then* $\delta(\Pi) = 3/4$.

PROOF: Given a permutation $\sigma = s_1 s_2 \ldots s_n$ we define an edge colouring of the complete graph on vertices $\{1, 2, \ldots, n\}$ by the rules:

- The edge between s_i and s_j is coloured red if $i < j$ and $s_i < s_j$,
- All other edges are coloured blue.

Then every 123-subsequence s_i, s_j, s_k of σ corresponds to a red triangle (3 vertices joined by red edges) and every 321-subsequence corresponds to a blue triangle.

It is known that any 2-colouring of the edges of the complete graph on an even number of vertices must have at least $2\binom{n/2}{3}$ monochromatic triangles [3]. But, if n is even, the permutation

$$n/2, n/2-1, \dots, 2, 1, n, n-1, \dots, n/2+1$$

has exactly this number of 321-subsequences (and no 123-subsequences). So this permutation has the least possible number of 123 and 321-subsequences. Therefore it has the greatest number of Π-subsequences and the result now follows easily. ■

Proposition 3.5 *If* $\Pi = \{132, 213, 231\}$ *then* $\delta(\Pi) = 3/4$.

PROOF: Clearly $\delta(\{132, 213\}) \leq \delta(\Pi) \leq \delta(\{132, 213, 231, 312\})$. However, Remark 2.6 together with the previous result shows that equality holds throughout. ■

4 Open problems

We have seen that with respect to sets Π of layered patterns, the structure of Π-maximal permutations is reasonably well understood. However, there still remain a number of open problems which we have not addressed in this paper.

We stated earlier (Conjecture 2.9) a conjecture concerning a sufficient condition for a layered pattern to have maximisers with a bounded number of layers. Obviously the first problem is to confirm or refute this conjecture. Next, in the cases already known to be unbounded (beginning or ending with a singleton layer), there is the problem of determining whether or not there is a limiting optimal layering. For the two layers of sizes 1 and t this was part of Price's original analysis which established that in this situation there is a geometric decrease in the layer sizes. However, even for such simple sequences of layer sizes as $[1, 2, 1]$ or $[1, 3, 2]$ no such straightforward description is known.

A further interesting set of problems revolves around the situation where there are consecutive layers of size 1. In Proposition 2.4 we were able to determine the optimal layering for the pattern whose layers are of sizes 1, 1, and 2. Again, computational evidence suggests that pairs (or longer segments) of consecutive layers of size 1 are generally reflected in long sets of size 1 layers in the maximisers.

The question of determing the actual number of layers required in the bounded case, or better limits on that number is also open. By way of illustration consider the case of three layers, of sizes 2, 3, and 2. The bound given in Theorem 2.7 shows that at most 30 layers are required in a maximiser. The actual number of layers required seems to be 3, and if this is true, their relative lengths are clearly 2/7, 3/7, and 2/7. On the other hand for layers in the pattern of sizes 2, 1, and 2, there is a similar bound on the number of layers required by estimates like those of Theorem 2.7 but this time the actual answer seems to be 4. If this is true, then the optimal layer proportions are a, $(1-a)/2$, $(1-a)/2$, and a where a is the real root of:

$$40x^3 - 32x^2 + 9x - 1 = 0$$

(equalling approximately 0.3829).

Finally, the reader will already have noticed that the unlayered case is still very much terra incognita. Our lower bound in Proposition 3.1 goes a little beyond the most straightforward idea of structuring a permutation recursively in the same way as the pattern which we are trying to maximise, but is still fundamentally based on it.

References

[1] M. Bóna: Exact enumeration of 1342-avoiding permutations, A close link with labeled trees and planar maps, Journal of Combinatorial Theory, Series A, 80 (1997) 257–272.

[2] S. A. Burr, V. Rosta: On the Ramsey multiplicities pf graphs—Problems and recent results, J. Graph Theory 4 (1980), 347–361.

[3] A. W. Goodman: On sets of acquaintances and strangers at any party, Amer. Math. Monthly 66 (1959) 778–783.

[4] A. Price, Packing densities of layered patterns, Ph.D. thesis, University of Pennsylvania, Philadelphia, PA, 1997.

[5] R. Simion, F. W. Schmidt: Restricted permutations, Europ. J. Combinatorics 6 (1985), 383–406.

[6] Z. E. Stankova: Classification of forbidden subsequences of length 4, European J. Combin. 17 (1996), no. 5, 501–517.

[7] W. Stromquist: Packing layered posets into posets, Unpublished typescript, 1993.

[8] J. West: Generating trees and the Catalan and Schröder numbers, Discrete Math. 146 (1995), 247–262.