

DENSE PACKING OF PATTERNS IN A PERMUTATION

HENRIK ERIKSSON, KIMMO ERIKSSON, SVANTE LINUSSON, AND JOHAN WÄSTLUND

ABSTRACT. We study the problem of finding the shortest permutation containing all patterns of a given length k . An upper bound of $3k^2/4$ is established. We also prove that as $k \rightarrow \infty$, there are permutations of length $k^2/4 + o(k^2)$ containing almost all patterns of length k .

RÉSUMÉ. Nous étudions le problème de trouver la permutation la plus courte contenant tous les motifs d'une longueur donnée k . Une borne supérieure de $3k^2/4$ est établie. Nous montrons aussi qu'il y a des permutations de longueur $k^2/4 + o(k^2)$ contenant presque tous les motifs de longueur k .

This paper was born at FPSAC'01 in Arizona, where Richard Stanley presented the mathematical work of the late Rodica Simion. Among many other things, she was a pioneer in the field of pattern avoiding permutations. In his talk, Stanley mentioned that Rodica Simion and Frank Schmidt [2] gave a formula for the number of permutations of length n that do not avoid *any* pattern of length three. There is little hope of finding a similar formula for arbitrary pattern length k , considering how complex the theory of pattern avoiding permutations becomes for larger k . But listening to Stanley, the four present authors became curious about a related question:

The pattern packing problem: What is the length L_k of the shortest permutation containing all patterns of length k ?

We later found that this pattern packing problem was posed already in 1999 by Arratia [1], but only trivial bounds on L_k have been given so far. The pattern packing problem is of course reminiscent of other dense packing problems, such as that of finding a shortest bit sequence containing every binary k -word as a contiguous subsequence. For the latter problem there is a well-known solution, the so called *de Bruijn sequences*, which contain every k -word exactly once. We cannot hope for such an efficient solution to the pattern packing problem. For instance, the shortest possible permutation containing all patterns of length $k = 2$ is evidently of length $L_2 = 3$ (there are four of them: 132, 213, 231, 312). But such a permutation contains $\binom{3}{2} = 3$ subsequences of length 2, of which at most two can have different patterns.

In spite of our efforts, the problem is still unsolved. We offer the conjecture that $L_k \sim k^2/2$ asymptotically.

Our story is one of partial results, to be developed in nine sections as follows.

- (1) The minimal length L_k of a permutation containing all k -patterns lies between k^2/e^2 and k^2 .
- (2) We represent patterns by dot configurations on a square grid and we characterize compact representations in terms of ascents and inverse descents.
- (3) The pedestrian game is introduced to describe equivalent representations of a pattern.
- (4) Strong convergence of the game implies a unique terminal position.
- (5) The terminal configuration is in fact the compact representation of the pattern.

- (6) By a probabilistic argument, we show that any dot configuration can be played onto white squares of a k by $3k/2$ chessboard. This improves the upper bound to $L_k \leq 3k^2/4$, our main result.
- (7) *A variation:* If we relax our ambition to finding a permutation containing *almost all* k -patterns, we show that a length of $k^2/4 + o(k^2)$ suffices.
- (8) *A second variation:* If we are satisfied with a permutation containing each k -pattern or its inverse, then a length of $k^2/2$ is sufficient.
- (9) *Final twist:* Returning to the terminal positions of the pedestrian game, a simple counting argument gives a new proof of an identity of Carlitz.

1. ELEMENTARY BOUNDS ON L_k

In this section, we shall see that L_k is asymptotically proportional to k^2 and obtain the elementary bounds

$$\frac{k^2}{e^2} \leq L_k \leq k^2.$$

The lower bound is derived from the observation that the total number of k -subsequences in the permutation must be at least as big as the number of all possible k -patterns. In other words,

$$\binom{L_k}{k} \geq k!,$$

or equivalently,


$$(1) \quad L_k(L_k - 1) \cdots (L_k - k + 1) \geq (k!)^2.$$

The left-hand side is less than L_k^k . By Stirling's formula, the right-hand side is approximately $2\pi k(k/e)^{2k}$. Hence we obtain a lower bound

$$\frac{k^2}{e^2} < L_k.$$

This bound is mentioned in [1], where it is also conjectured to be sharp (contradicting our own conjecture).


The simple upper bound of $L_k \leq k^2$ is obtained as follows. Define a k^2 -permutation by arranging the numbers between 1 and k^2 into k sequences of length k , each decreasing by k in every step, the sequences taken in increasing order. For example, for $k = 3$ the permutation of length 9 is 741852963. The dot matrix of such a permutation (with values increasing upwards) will look like a slightly *tilted square*:




Clearly every k -pattern can be found in this permutation; simply take one dot in each row and column of the tilted square. For example, the pattern 231 is found in the subsequence 483.

2. MORE COMPACT PACKINGS

Since the rows in the tilted square actually correspond to increasing subsequences, while columns correspond to decreasing subsequences, it is possible to realize a pattern by a more compact subset of dots in the tilted square. For example, the permutation pattern 231 is realized by all the following dot patterns in the tilted square (and a few others).



There is a direct way of figuring out the most compact such dot representation of a pattern: Start with the dot matrix of the pattern. Draw a vertical line at every ascent, and a horizontal line at every descent in the inverse permutation (obtained by reflection in the line $y = x$). Adjust the dots so that they lie straight in the rows and columns induced by the lines. For example, the compact representation of 2614357 is obtained as follows:




A direct description of the compact position of a given permutation can be given in terms of descents of the permutation and its inverse. We write the permutation π as $(\pi_1 \pi_2 \pi_3 \dots \pi_k)$.

A pair π_i, π_{i+1} of adjacent symbols is a *descent* if $\pi_i > \pi_{i+1}$, otherwise it is an *ascent*. We let $d(\pi)$ be the number of descents in π , and we let $a(\pi)$ be the number of ascents. A descent in the inverse π^{-1} is a pair of symbols $x, x+1$ such that $x+1$ comes before (but not necessarily adjacent to) x in the permutation π . Such a pair is therefore called an *inverse descent*.

For each symbol π_i in π , we count the ascents *before* π_i , that is, pairs π_j, π_{j+1} with $j < i$ and $\pi_j < \pi_{j+1}$. We also count the inverse descents *below* π_i , by which we mean pairs of symbols $x, x+1$ such that $x < \pi_i$ and $x+1$ comes before x . The compact dot configuration in the tilted square is constructed by putting, for each symbol π_i in π , a dot in the site that has as many columns east of it as there are ascents before π_i , and as many rows south of it as there are inverse descents below π_i .

3. THE MELBOURNE PEDESTRIAN GAME

We will now introduce a game for transforming dot patterns. The setting is k pedestrians walking on a square grid in Melbourne. In Australian traffic, as a pedestrian you are supposed to look for any traffic to your right before you cross a street, and also watch out for any traffic to your left on the other side of the street. Thus, in our game a pedestrian is allowed to take a step forward to an unoccupied site on the grid if and only if there is nobody anywhere to her right before she takes the step, and will be nobody anywhere to her left after she has taken the step. "Forward" may be any direction on the grid (north, south, east or west).



Observe that moves are reversible. A *position* in the game is a placement of the k pedestrians on k distinct sites on the grid. Two positions are *game equivalent* if one can be reached from the other by a sequence of moves. Since moves can be reversed, this is an equivalence relation.

Proposition 3.1. *Two positions are game equivalent if and only if they realize the same pattern in the tilted square.*

Proof. The pattern realized by a position is invariant under the game: When a dot moves one step it changes its relative position only to dots that are to the right of the old site or to the left of the new site, but a move can be made only when no such dots are present. Hence, positions that are game equivalent realize the same pattern.

To prove the converse, we shall prove that any position can be played to a permutation (one dot in each row and column of a k by k square). By the game invariance of patterns, this permutation must be the unique permutation realizing this pattern. Hence any two positions that realize the same pattern can be played to each other via the permutation.

For any position, we can play the eastmost dot of the top row upwards as far as we want. Among the remaining dots we can again play the eastmost dot of the top row upwards as far as it does not reach the row of the first dot. Continuing in this way, we obtain a position with the dots on distinct rows. We then perform the same operation on columns, playing one dot at a time eastwards. Since this can be done without changing the row of any dot, we reach a position where the dots occupy sites on k distinct rows and k distinct columns.

Any empty rows or columns can be filled by playing in the dot from the row or column next to it. Hence we obtain a permutation. \square

Corollary 3.2. *There are $k!$ game equivalence classes with k pedestrians.*

The following theorem establishes a connection between the pedestrian game and the pattern packing problem:

Theorem 3.3. *L_k is the minimal size of a union of one representative from each game equivalence class in the pedestrian game with k pedestrians.*

Proof. A set S of sites which is a union of one representative from each game-equivalence class has the property that every position can be played so that all dots are moved into S . Hence the permutation that is represented by S contains each pattern of length k . Conversely, a permutation that contains each pattern of length k can be represented as a set of sites with the property that any position with k dots can be played into it. This set of sites therefore contains a representative of each game-equivalence class. \square

4. STRONG CONVERGENCE OF THE PEDESTRIAN GAME

In this section we consider the pedestrian game played on a k by k board. Proposition 3.1 still holds in this setting, for a position with k dots which contains several dots in the same rows or columns can be untangled within the board limits. The modified procedure is as follows.

First move the top dot to the top row, if necessary. Then make sure that there are at least two dots in the top two rows, possibly by moving a second dot to the north. In general, for every $m \leq k$, make sure that there are at least m dots in the top m rows. This also means that the bottom $k - m$ rows will contain at most $k - m$ dots. In particular, the bottom row contains at most one dot, and by moving the bottom dot to the south, we can make sure that there is exactly one. Repeating this for each row, we get exactly one dot in each row, and to obtain this, we have not made any horizontal moves. This means that if the same operation is performed on columns, we reach a permutation, i.e. a position where each row and each column contains exactly one dot.

Hence each game-equivalence class on the k by k square has a unique representative with one dot in each row and one dot in each column. We now show that there is another natural representation of game-equivalence classes.

Proposition 4.1. *Each game-equivalence class on the k by k board has a unique representative from which it is impossible to make a move in the south or west directions.*

Proof. By the *directed* pedestrian game, we mean the pedestrian game played on the positive quadrant, with the restriction that only moves in the south and west directions are permitted. A position where no such move is possible is called a *terminal* position. As we will show, the directed pedestrian game is *strongly convergent*, which means that from any given initial position, all move sequences lead to the same terminal position in the same number of moves. It is clear that this also proves the proposition. \square

A simple criterion for strong convergence is the *polygon property* [3]: In any position where two different moves, x and y , are legal, there are two play sequences of the same length and beginning with x and y respectively leading to the same position.

Lemma 4.2. *The directed pedestrian game on the k by k board is strongly convergent.*

Proof. We have to verify the polygon property. In this case, we show that two different moves from the same position actually *commute*, that is xy and yx are both legal and give the same result.

If one and the same dot can make a move in either of the directions south and west, then it is clear that after a move south, the move to the west will still be legal, and vice versa.

If two different dots can move in any of the two legal directions, it is easy to see that these moves cannot interfere. It follows that the directed pedestrian game is strongly convergent. \square




FIGURE 1. The six terminal positions in the directed pedestrian game for $k = 3$.

5. THE TERMINAL POSITION IS THE COMPACT REPRESENTATION

The unique terminal position of the directed pedestrian game provides a canonical dot representation for every k -pattern. We shall now see that it is the same compact configuration that we constructed from the ascents and inverse descents of the permutation.

Theorem 5.1. *The compact dot representation of π defined in Sec. 2 is the terminal position corresponding to π .*

Proof. First, we show that this position represents the permutation π . Consider any pair of symbols $x < y$ that occur in order, that is

$$\pi = (\dots x \dots y \dots).$$

Then there is at least one ascent between x and y , so the dot corresponding to y is in a column east of the dot corresponding to x . Since $x < y$, there are at least as many inverse descents below y as there are below x , so either the dots are in the same row, or the one corresponding to y is higher.

If the pair occurs in the wrong order, that is $\pi = (\dots y \dots x \dots)$ then there are at least as many ascents before x as there are before y , and at least one more inverse descent below y than below x . Hence in this case the dot corresponding to x will be in a row south of the dot corresponding to y , and in the same column or a column east of it.

This shows that every pair of dots has the correct relative location. Hence the position described represents π .

To show that the position we have defined is a terminal position, we first prove that no dot can move west. The dot corresponding to the symbol π_1 obviously cannot. If a pair π_{i-1}, π_i is a descent, then the dot corresponding to π_{i-1} will be above the one corresponding to π_i , in the same column. Hence the dot corresponding to π_i cannot move west. If on the other hand π_{i-1}, π_i is an ascent, then the dot corresponding to π_{i-1} will be in the column immediately to the west of the dot corresponding to π_i , and above it or in the same row. Hence the dot corresponding to π_i still cannot move west.

To show that no dot can move south either, we note that the symbol 1 is already in the bottom row. If a pair $x-1 \dots x$ is an inverse ascent, the corresponding dots will be in the same row so the x -dot cannot move south. And for an inverse descent $x \dots x-1$, the x -dot will have the other dot one row below, in the same column or to the east. In either case, it cannot move south. \square

6. THE MAIN RESULT

Our original motive for studying compact configurations was a desire to improve the upper bound $L_k \leq k^2$. We computed L_k for small k and the following table points to the existence of a tighter bound.

k	1	2	3	4	5
L_k	1	3	5	9	13

To our disappointment, the union of all compact configurations covers all of the $k \times k$ -board except a thin slice along the border which is asymptotically insignificant. A closer study of the optimal permutations found by the computer revealed that it may be easier to pack sparse configuration in a chessboard full of holes!

For $k = 3$ there exist only two permutations of minimal length ($L_k = 5$) containing all k -patterns. One of them is 41352, which can be interpreted in terms of chessboards. As shown in the figure, 41352 is given by the white squares of the 3 by 3 chessboard in our standard way of associating patterns with subsets of squares on the grid.




FIGURE 2. The white squares permutation 41352.

The same thing is true for $k = 5$; the thirteen white squares of a 5×5 chessboard define a permutation containing all 5-patterns. It is tempting to conjecture that the white squares permutation of any $k \times k$ -board contains all k -patterns, at least if k is odd, and in the first draft of this paper we gave in to that temptation. From our failed attempts to prove the false conjecture a probabilistic argument emerged that nonconstructively demonstrates the existence of a counterexample. Fortunately, a similar argument proves that a fifty percent wider rectangle *can* accommodate all patterns on its white squares.

Theorem 6.1. *The permutation given by the white squares of the $k \times \lfloor 3k/2 \rfloor$ chessboard contains all k -patterns, therefore*

$$L_k \leq \frac{3}{4}k^2 \text{ for } k > 1$$

Proof. For any given pattern, we start with the standard dot configuration in the k by k square to the west. If all dots are already on white squares, we are done; otherwise, some horizontal moves need to be made. Going from west to east, each time we encounter a dot on a black square, we move it together with all dots east of it one step. If the black dot is the second dot in a descent, this mass move goes west, otherwise it goes east. If we are lucky, the dots will stay within the stipulated rectangle, but we may as well have bad luck, as shown in the figure.

Returning to the starting position, we note that some vertical moves are possible. At every inverse ascent, we have the option of moving the upper dot, along with all dots above it, one step to the south. We now use a probabilistic argument. For every inverse ascent, we flip a coin. On heads, we do nothing; on tails, we move the upper dot and every dot above it to the south. The dots obviously stay within the original k by k square. As before, we then make the necessary moves in the horizontal direction to put the dots on white squares, and we will show that there is now a nonzero probability of the dots staying inside the stipulated rectangle.




FIGURE 3. Two ways to move 3462517 to white squares

We estimate the expected number of moves to the east, counting a move to the west as -1 . The first dot will cause a move to the east if it is on a black square after the southward moves. Every other dot will cause a move if after the southward moves it is on a square of the color opposite to that of the previous dot. Otherwise it will automatically be pushed to a white square when the previous dot is moved to a white square. The potential move will be to the west if the dot is the second dot in a descent, and otherwise to the east.

The key point is that the probability of a dot being on a square of color opposite to that of the previous dot is $1/2$ as soon as there is at least one potential southward move that will change the color of the upper dot but not of the lower one. Since the southward moves are decided by independent coin flips, the probability that the two dots end up on squares of different colors is exactly $1/2$. Now, it is easy to see that two adjacent symbols in a permutation must have an inverse ascent in the interval between them unless it is a descent by one, like 54. In that special case, the probability of the dots occupying squares of different colors drops to zero.

To sum up, the first dot will cause at most one move to the east, and for the others, a descent by one will contribute zero moves to the east, while a descent by more than one will contribute $-1/2$ and an ascent $1/2$ to the expected number of moves to the east. Therefore the expected number of necessary moves to the east is at most

$$1 + \frac{1}{2}a(\pi).$$

Hence with one exception, the identity permutation $[123 \dots k]$ with $k - 1$ ascents (which is trivial), the expected number of moves to the east is at most $k/2$, so some lucky outcome

of the coin flips will need at most $\lfloor k/2 \rfloor$ extra columns in order to play all the dots to white squares. \square

7. VARIATION 1: PERMUTATIONS CONTAINING ALMOST ALL k -PATTERNS

We show that a permutation need not be much longer than $k^2/4$ in order to contain almost all patterns of size k .

Theorem 7.1. *Let f be a function such that $f(k)$ tends to infinity as k does. Then there is a sequence $\sigma(k)$ of permutations of length less than $k^2/4 + f(k)k^{3/2}$ that contains almost all patterns of size k in the sense that the fraction of patterns of length k that are not contained in $\sigma(k)$ tends to zero as k tends to infinity.*

Proof. Let π be a random permutation of length k . We can write $d(\pi) = d_{\text{odd}} + d_{\text{even}}$, where d_{odd} and d_{even} are the number of descents in odd and even position, respectively, that is, d_{odd} is the number of descents (π_i, π_{i+1}) for which i is odd, and d_{even} is the number of such descents for which i is even. The both d_{odd} and d_{even} are binomial distributed. Their distributions can therefore be approximated by normal distributions with mean $k/4$ and standard deviation $O(\sqrt{k})$.

Now let g be a function such that $g(k) \rightarrow \infty$ as $k \rightarrow \infty$. Then the probability that d_{even} or d_{odd} deviates by more than $g(k)\sqrt{k}/2$ from $k/4$ tends to zero as k tends to infinity. Hence the probability that $d(\pi) < k/2 - g(k)\sqrt{k}$ tends to zero. Similarly, the probability that $a(\pi^{-1}) < k/2 - g(k)\sqrt{k}$ tends to zero.

It follows that with high probability, the permutation π can be played into a square of side $k/2 + g(k)\sqrt{k}$. Hence there is a permutation of size $(k/2 + g(k)\sqrt{k})^2$ that contains almost all patterns of size k .

If we put

$$(k/2 + g(k)\sqrt{k})^2 = k^2/4 + f(k)k^{3/2},$$

then it is clear that g tends to infinity if and only if f does. The theorem follows. \square

8. VARIATION 2: PERMUTATIONS CONTAINING PATTERN OR INVERSE

We can construct a permutation of length $\binom{k}{2}$ which in a slightly weaker sense contains every pattern of length k . This permutation consists of the elements on and below the diagonal in the tilted square. For example, $T_2 = 312$ and $T_3 = 641523$.

Theorem 8.1. *For every pattern τ of length k , the triangular permutation T_k contains either τ or τ^{-1} .*

Proof. A dot in the terminal position of a permutation will be outside T_k only if the sum of the number of ascents before the symbol and the number of inverse descents below it is at least k . An inverse descent in τ is of course a descent in τ^{-1} , so the total number of ascents and inverse descents is $a(\tau) + d(\tau^{-1})$. If that total is less than k , the permutation T_k will certainly contain τ . Analogously, if the total $a(\tau^{-1}) + d(\tau)$ is less than k , the permutation T_k will certainly contain τ^{-1} .

But there are $k - 1$ positions which are either descents or ascents, so we have

$$d(\tau) + a(\tau^{-1}) + d(\tau^{-1}) + a(\tau) = 2(k - 1).$$

The theorem follows from the fact that either τ or its inverse must have the property that the total number of ascents plus the total number of inverse descents is at most $k - 1$. \square

9. FINAL TWIST: ENUMERATION OF GAME POSITIONS

We consider the pedestrian game played on an m by n board. We wish to count the representations of a given permutation, that is, the number of positions in a given game-equivalence class.

If a certain permutation π can be represented at all on an m by n board (m rows and n columns), then there is a terminal position P_0 of π in the south-west directed pedestrian game. A different representation P of the same permutation π can now be described by labeling each dot with the two coordinates of the distance it has to move from P_0 to P . Since P_0 is the terminal position, each dot has to move north and east to reach P . Moreover, if a dot x is west of a dot y , then y has to move at least as many steps east as x , since any move to the east by x will bump all other dots one step east.

The terminal position P_0 occupies $a(\pi) + 1$ columns and $d(\pi^{-1}) + 1$ rows. Therefore, the number of times we can allow the eastmost dot to be bumped without being bumped off the m by n board is $n - a(\pi) - 1$. Hence the sequence of horizontal distances that the dots, taken from west to east, have to travel to get from P_0 to P forms a weakly increasing sequence of k nonnegative integers that is bounded from above by $n - a(\pi) - 1$. An elementary theorem in enumerative combinatorics tells us that the number of such sequences is

$$\binom{n - a(\pi) - 1 + k}{k} = \binom{n + d(\pi)}{k}.$$

Similarly, the number of sequences of vertical distances from P_0 to P of the dots taken from bottom to top is

$$\binom{m + a(\pi^{-1})}{k}.$$

Since we get a representation of π for each pair of such integer sequences, we see that the number of representations of a permutation π on an m by n board is

$$\binom{m + a(\pi^{-1})}{k} \binom{n + d(\pi)}{k}.$$

From this we obtain, by summing over all permutations in S_k , a classical identity, probably due to Carlitz.

$$\sum_{\pi \in S_k} \binom{m + a(\pi^{-1})}{k} \binom{n + d(\pi)}{k} = \binom{mn}{k},$$

since the total number of positions of k dots on an m by n board is $\binom{mn}{k}$.

REFERENCES

- [1] R. Arratia, On the Stanley-Wilf conjecture for the number of permutations avoiding a given pattern, *The Electronic Journal of Combinatorics* **6** (1999), #N1.
- [2] R. Simion and F. W. Schmidt, Restricted permutations, *European Journal of Combinatorics* **6** (1985), no. 4, 383–406.
- [3] K. Eriksson, Strong convergence and a game of numbers, *Eur. J. Combinatorics* **17** (1996), 379–390.

NADA, KTH, SE-100 44 STOCKHOLM, SWEDEN

E-mail address: **henrik@nada.kth.se**

IMA, MÄLARDALENS HÖGSKOLA, SE-721 23 VÄSTERÅS, SWEDEN

E-mail address: **kimmo.eriksson@mdh.se**

MAI, LINKÖPINGS UNIVERSITET, SE-581 83 LINKÖPING, SWEDEN

E-mail address: **{linusson, jowas}@mai.liu.se**